The Monte Carlo Method: Versatility Unbounded In A Dynamic Computing World
Chattanooga, Tennessee, April 17–21, 2005, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2003)

A FUNCTIONAL EXPANSION METHOD FOR MONTE CARLO EIGENVALUE CALCULATIONS

David P. Griesheimer, William R. Martin, James Paul Holloway
Department of Nuclear Engineering and Radiological Sciences
University of Michigan
Ann Arbor, MI 48109
dgrieshe@umich.edu; wrm@umich.edu; hagar@umich.edu

ABSTRACT

A new fission source convergence methodology has been developed for Monte Carlo criticality calculations. The new method uses functional expansion tallies (FET) to create functional approximations for the shape of the fission source. New fission sites can then be obtained by sampling directly from the functional form of the source. This eliminates the need to store individual fission sites in a fission bank. In addition, there is no longer a need for fixed length cycles during the source convergence and eigenvalue calculation. Instead additional, higher-order moments can be added to the source approximation as their statistical uncertainty reaches an acceptable level. This may allow faster convergence than is possible with the collision estimators that are used in current fission bank routines. A series of 1D benchmark tests were used to compare the FET source convergence method to the traditional fission bank method.

Key Words: Monte Carlo method, criticality, functional expansion tally, fission source

1 INTRODUCTION

Monte Carlo methods have long been an important tool for performing eigenvalue calculations for reactor analyses. Unfortunately, Monte Carlo simulations of large, loosely-coupled systems can have slow fission source convergence rates. Achieving a sufficiently converged fission source requires a large number of neutron generations (with many neutrons per generation) to be run. Simulations run with poorly converged fission sources can result in biased estimates of the system eigenvalue. The computational burden that results from this slow convergence has become a limiting factor in the use of Monte Carlo methods for whole core commercial reactor analysis. In order to address these issues a new Monte Carlo fission source convergence methodology has been developed and tested. The new method uses functional expansion tallies (FET) to create functional approximations for the shape of the fission source. It is hoped that this method will accelerate global fission source convergence in loosely-coupled systems.

2 THEORETICAL DEVELOPMENT

In Monte Carlo simulations the most commonly used method for estimating the eigenvalue in criticality calculations is the fission bank method. This method is actually an application of source iteration, a technique that is used in many deterministic transport solvers. The source iteration scheme iterates on the eigenvalue (k_{eff}) and eigenvector (fission source) for the system until convergence of these quantities is achieved. The method begins with an initial fission source and uses an appropriate transport solver to calculate the flux distribution in the system, due to this fixed source term. At the end of this fixed source calculation, a new eigenvalue and source distribution is computed from the resulting flux solution. The updated eigenvalue and source distribution is then used as a new source term for the system, and another transport calculation is performed. This process is repeated until both the eigenvalue and eigenvector have converged to within a preset threshold.

In deterministic methods, the implementation of the source iteration technique is fairly straightforward. Complex systems are subdivided into small nodes and only the average (or integrated) flux and source terms in each node are stored. The neutron transport equation can be rewritten in terms of these node average quantities, allowing the entire system to be expressed as a series of coupled linear equations that can be solved numerically.

Unfortunately, the extension of these ideas to Monte Carlo is not straightforward. Systems analyzed with Monte Carlo methods do not need to be subdivided into small, homogenous nodes. In fact, one of the major advantages of Monte Carlo methods is their ability to handle complex geometries exactly, without the approximations that often accompany mesh-based solutions. However, this ability also makes it difficult to estimate and store the shape of the fission source that results from a fixed source calculation. While Monte Carlo methods are extremely good at evaluating integral quantities (e.g. average flux in a volume, average current through a surface), they are not as well-suited for estimating the shape of a solution distribution. In Monte Carlo source iteration, both knowing the shape of the fission source and being able to sample from the source are key components.

2.1 Monte Carlo Eigenvalue Estimators

Estimating the system eigenvalue from a fixed-source Monte Carlo simulation is relatively simple. The eigenvalue is especially well suited for evaluation by Monte Carlo because it is an integral quantity,

$$k_{eff} = \frac{\int_x \int_E \nu(x,E)\Sigma_f(x,E)\phi_G(x,E)\,dx\,dE}{\int_x \int_E S_G(x,E)\,dx\,dE}, \tag{1}$$

where ϕ_G is the scalar flux distribution that results from a fixed source distribution S_G, Σ_f is the macroscopic fission cross section in the system, and ν is the average number of neutrons produced per fission event. Because k_{eff} is an integral quantity, any estimator for the eigenvalue needs only to know how many fission neutrons are produced during a fixed-source calculation, not the spatial distribution of those fission events. In practice, there are two main types of Monte Carlo estimators that are commonly used for criticality calculations: the collision and track length estimators.

As its name implies, the collision estimator scores the probability that a neutron will cause a fission event at each collision. Before formally defining this estimator, it is necessary to introduce some mathematical notation. We define N as the total number of independent particle histories in a single neutron generation, and i as a sequential index to describe these individual particles. The number of collisions that a single history undergoes is defined as C_i, and is indexed by the variable c. The total number of fissile isotopes in a problem is defined as J, with each individual isotope denoted by the identifier j. The macroscopic fission and total cross sections for an isotope j are represented by $\Sigma_{f,j}$ and $\Sigma_{t,j}$, respectively. The average number of neutrons released per fission event in isotope j is denoted by ν_j. With this notation, the collision estimator can be written as,

$$\hat{k}_{eff}^{C} = \frac{1}{N}\sum_{i=1}^{N}\sum_{c=0}^{C} w_{i,c}\sum_{j}^{J}\left[\frac{\nu_j\left(E_{i,c}\right)\Sigma_{f,j}\left(x_{i,c},E_{i,c}\right)}{\Sigma_{t,j}\left(x_{i,c},E_{i,c}\right)}\right], \tag{2}$$

where $w_{i,c}$, $x_{i,c}$, and $E_{i,c}$ and are the weight, position, and energy, respectively, of neutron i prior to collision c. At the end of a generation, Eq. (2) gives an estimate for the expected number of fission neutrons produced per starting neutron, or the effective multiplication for the system. The collision estimator is often considered the best choice for very large systems[1].

The track length estimator allows a neutron to contribute to the eigenvalue estimate along its entire path length. The estimator is derived directly from the fact that the scalar flux in a volume can be interpreted as the rate at which track length is generated by particles passing through the cell. The track length estimator is given by

$$\hat{k}_{eff}^{TL} = \frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} w_{i,c}\, d_{i,c}\sum_{j}^{J}\nu_j\left(E_{i,c}\right)\Sigma_{f,j}\left(x_{i,c},E_{i,c}\right), \tag{3}$$

where $d_{i,c}$ is the total distance traveled by particle i between events $c-1$ and c. The track length estimator is especially well-suited for optically thin systems where a limited number of collisions may cause the collision estimator to have a relatively large variance.

2.2 Monte Carlo Fission Source Sampling Routines

The previous section showed that estimating the system eigenvalue for each neutron generation can be done with simple, well-known, Monte Carlo tallies. Unfortunately, estimating the shape of the fission source distribution, and sampling from the distribution is not as easy.

Traditionally sampling the fission source has been handled by storing the sites of individual fission events that are produced by particles during independent neutron generations. Each neutron generation is an independent, fixed source simulation that uses a predefined number of neutron histories started from the same source distribution. At each collision event during the neutron histories, the Monte Carlo code stochastically determines how many, if any, secondary neutrons are produced from fission events at that collision site. The locations and energies of these fission neutrons are stored in a list, commonly referred to as the fission bank. After each fixed-source calculation has ended, the contents of the fission bank are used as the source points for the next fixed-source calculation. Thus, the actual shape of the fission source is never known, instead the fission bank stores samples from the source distribution. This method, however, raises important questions about the independence of source points within the fission bank. During each generation it is likely that some particle histories will produce multiple source point entries in the fission bank. Clearly, source locations that are generated by a single neutron history are closely related. This leads to the possibility that some of the fission bank source points for a generation may, in fact, be correlated with one another. In addition to the possibility of correlation among the source points for a single generation, there is also a large potential for correlation among the source points between generations. In systems where the mean free path of a neutron is small, neutrons will tend to produce fission events near their birth sites. In these systems, any perturbations in the source shape (e.g. over or under sampling source points in a localized region) will propagate from generation to generation.

One potential alternative to the traditional fission bank method for sampling the fission source is to use functional expansion tallies (FET) to estimate a functional approximation of the source shape between neutron generations. Functional expansion tallies are used in Monte Carlo simulations to estimate the moments of flux distributions with respect to a set of orthogonal basis functions[2-7]. These moments can then be used to reconstruct a series approximation for the true shape of the flux distribution. In order to apply the FET to Monte Carlo criticality calculations, we begin by writing the fission source distribution as a series expansion in Legendre polynomials,

$$S_G\left(\tilde{x}\right)=\sum_{n=0}^{M}\frac{2n+1}{2}a_n\,P_n\left(\tilde{x}\right), \tag{4}$$

where $P_n\left(\tilde{x}\right)$ is the n^{th} Legendre polynomial, a_n are the expansion coefficients, M is the truncation order of the series, and $\tilde{x}$ is a scaled spatial variable defined on the Legendre domain $\tilde{x}\in[-1,1]$. Using the orthogonality property of the Legendre polynomials, it is possible to solve Eq. (4) for the expansion coefficients

$$a_n=\int_{\tilde{x}}S_G\left(\tilde{x}\right)P_n\left(\tilde{x}\right)d\tilde{x}. \tag{5}$$

Finally, rewriting Equation (5) in terms of the flux distribution from the previous neutron generation, ϕ_{G-1},

$$a_n=\int_{\tilde{x}}\int_{E}\nu\left(\tilde{x},E\right)\Sigma_f\left(\tilde{x},E\right)\phi_{G-1}\left(\tilde{x},E\right)P_n\left(\tilde{x}\right)dE\,d\tilde{x}, \tag{6}$$

gives the an expression for the expansion coefficient in terms of a convenient flux moment integral at is suitable for estimation by Monte Carlo.

Unbiased Monte Carlo FET estimators for quantities of the form shown in Eq. (6) have previously been derived and published[2-12]. From these published results, it follows that the expansion coefficients can be estimated with either a collision or track length FET estimator, given in Eq. (7) and (8), respectively:

$$\hat{a}_n^C=\frac{1}{N}\sum_{i=1}^{N}\sum_{c=0}^{C}w_{i,c}\,P_n\left(\tilde{x}_{i,c}\right)\sum_{j}^{J}\left[\frac{\nu_j\left(E_{i,c}\right)\Sigma_{f,j}\left(x_{i,c},E_{i,c}\right)}{\Sigma_{t,j}\left(x_{i,c},E_{i,c}\right)}\right], \tag{7}$$

$$\hat{a}_n^{TL} = \frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C}\frac{w_{i,c}\, d_{i,c}}{\tilde{x}_{i,c}-\tilde{x}_{i,c-1}}\int_{\tilde{x}_{i,c-1}}^{\tilde{x}_{i,c}} P_n\left(\tilde{x}\right)d\tilde{x}\left[\sum_{j}^{J}\nu_j\left(E_{i,c}\right)\Sigma_{f,j}\left(x_{i,c},E_{i,c}\right)\right]. \tag{8}$$

Notice that both FET estimators for the fission source expansion coefficients appear similar to the eigenvalue estimators given in Eq. (2) and (3). In fact, in the case where $n = 0$, the FET estimators reduce exactly to the eigenvalue estimators, giving

$$\begin{aligned} \hat{a}_0^{C} &= \hat{k}_{eff}^{C}, \\ \hat{a}_0^{TL} &= \hat{k}_{eff}^{TL}. \end{aligned} \tag{9}$$

During each neutron generation, either Eq. (7) or Eq. (8) can be used to estimate a set of Legendre moments for the spatial source distribution. These moments can be used in Eq. (4) to give a functional approximation for the spatial shape of the fission source. During the next generation, source points can be sampled directly from this functional approximation, eliminating the need for the fission bank.

As with all Monte Carlo tallies, minimizing the statistical uncertainty in each of the Legendre moment estimates is very important for achieving an accurate result. Poorly converged moments can contaminate the functional source approximation with statistical noise, thus slowing the convergence rate of the source iteration. In order to avoid this problem, each neutron generation must contain a sufficient number of independent histories to reduce the statistical uncertainty in all of the expansion coefficients to an acceptable level. In addition to the statistical uncertainty, functional approximations produced with FET also contain truncation error due to approximating the true fission source by a low order series expansion. Therefore, any FET based source convergence scheme must use a sufficiently high order expansion to ensure that the source distribution is well represented by the functional approximation.

In the FET, these two sources of error are inversely related to each other. Low order expansion coefficients are easier to estimate with Monte Carlo and, therefore, have smaller statistical uncertainties. However, using too few coefficients in a series expansion will result in large truncation error and a low resolution approximation. On the other hand, keeping too many poorly converged coefficients will result in a final approximation that is heavily contaminated by statistical noise. In order to maximize the effectiveness of an FET source convergence method, an optimal balance must be found between these two sources of error, such that the total error in the approximation is minimized. To help with finding this optimal balance, a simple cost-to-benefit ratio has been developed[4,7] to help users determine which coefficient estimates should be included in the series expansion approximation. The metric is defined as

$$R_n^2 = \frac{2n+1}{2}\frac{\hat{\sigma}_{\hat{a}_n}}{\hat{a}_n}, \tag{10}$$

where $\hat{\sigma}_{\hat{a}_n}$ is the sample standard deviation for the coefficient estimate $\hat{a}_n$. The R_n^2 metric is simply the ratio of the increase in statistical error to the decrease in truncation error due to adding a nonzero coefficient $\hat{a}_n$ to the series expansion in Eq. (4).

This cost-to-benefit ratio provides a convenient test for determining how many expansion coefficients should be used for a given source approximation. Coefficients with values of $R_n^2 >> 1$ should not be included in the approximation because they are not well converged and do not add any useful information to the result. Terms with R_n^2 values less than, or close to 1, should be included in the approximation because they provide valuable information about the shape of the true function.

3 IMPLEMENTATION AND NUMERICAL RESULTS

In order to test the use of functional expansion tallies for eigenvalue calculations, the FET estimators given in Eqs. (7) and (8) were implemented in a modified version of MCNP4c. The set of Legendre polynomials was chosen as the expansion basis set. The initial version of the code is designed to estimate the first 20 spatial Legendre moments of the fission source. In this test code, both the collision and track length FET estimators operate concurrently with the original fission bank method. This means that each neutron generation produces three separate approximations for the spatial shape of the next fission source. The birth sites for the following generation can then be sampled from any of the three source approximations. Unfortunately, because the source approximations are implemented concurrently, it is not possible to obtain accurate run time comparisons between the methods.

For the FET source approximations, particle birth locations are sampled directly from the appropriate functional approximation using simple rejection sampling. To determine the functional approximation, the code considers all 20 spatial source moments estimated during the previous generation and rejects (sets to zero) those moments that have a cost-to-benefit ratio (R_n^2) greater than a user defined threshold. This filtering helps to reduce statistical noise in the functional approximation. The remaining coefficients are then normalized by the zeroth coefficient and used in the Legendre series expansion to give a normalized polynomial source approximation. For rejection sampling, sample points along the polynomial can be quickly and efficiently evaluated using the Legendre recursion relationships.

At birth, the initial direction of a source neutron is sampled isotropically in angle. The initial energy of the source neutron is sampled from the Maxwell fission energy spectrum with a mean temperature of 1.2895 MeV

$$P(E) = 0.77059\, e^{\left(-E/1.2895\right)} \sqrt{E}\,. \tag{11}$$

Note that the distribution in Eq. (11) does not take into account the energy of the neutron that causes the fission event. Unlike the fission bank method, the 1-D FET source convergence method does not retain any information about the parent neutrons for the source particles. This is because the source neutrons in the FET method are not directly associated with a single parent neutron. Rather they are sampled independently from the approximate source distribution created from all of the neutrons in the previous generation. This independent sampling strategy means that an unlimited number of source locations can be generated, without concern about intra-generation correlation between the samples.

For the fission bank source approximation, particle birth locations are taken directly from the list of stored fission sites during the previous generation. The internals of the fission bank source iteration scheme remained unchanged from the original distribution version of MCNP4c, with two minor exceptions. First, the source particle energies stored in the fission bank were not used. Instead, the initial energy of each source particle was sampled from the Maxwell fission spectrum given in Eq. (11). This resampling in energy allows a fair comparison between the FET and fission bank methods. Also, MCNP4c was modified to give a histogram edit for the contents of the fission bank after each generation. This change was made to allow easier visualization of the fission bank source shape.

For numerical verification of the FET source convergence method, the modified version of MCNP was run on a pair of one-dimensional homogeneous benchmark problems.

3.1 Fast Reactor Benchmark Problem

The first series of benchmark tests were performed on a simple one-dimensional representation of a homogeneous fast-spectrum critical assembly. The assembly was modeled as a 100 cm bare slab of pure uranium with a density of 10.97 g/cc. The isotopic composition of the uranium was 6.5% U^{235} and 93.5% U^{238}. In order to characterize this system a reference criticality calculation was performed using an unmodified version of MCNP4c. The reference calculation used 500 generations of 10,000 neutron histories. The first 50 generations of the simulation were discarded to avoid biasing the results with the initial (poor) source distribution. The simulation produced an eigenvalue estimate of 0.96674 with a relative standard deviation of 0.00021. The mean free path of neutrons followed during the simulation was 3.809 cm. The results from the reference calculation also confirmed that the system relies heavily on high energy neutrons to cause fission events (i.e. a fast system).

When compared with the diameter of the reactor, the mean free path of neutrons in the system is relatively large. This means that neutrons are able to spread from one side of the reactor to the other in only a few generations. Systems such as these are typically referred to as tightly coupled systems. In practice, source iteration schemes are very efficient for tightly coupled systems because local perturbations within the fission source shape are quickly dissipated, allowing rapid convergence of the global solution. Thus, it should be anticipated that both the fission bank and FET source convergence methods should provide fast convergence to the critical fission source distribution.

To compare the FET and fission bank source convergence methods, two separate k-code calculations were run. Each simulation used 3,000 source generations with 10,000 neutron histories per generation. Because the initial convergence of the fission source is of primary interest, no source generations were discarded. During each generation the collision FET was used to estimate the first 20 spatial Legendre expansion coefficients of the fission source. These coefficients were filtered using a cost-to-benefit ratio threshold of 0.95. Any coefficients with an R_n^2 value greater than 0.95 were set to zero. For sampling source points for the following generation, one simulation used sites stored in the fission bank, while the second simulation sampled directly from the FET functional approximation.

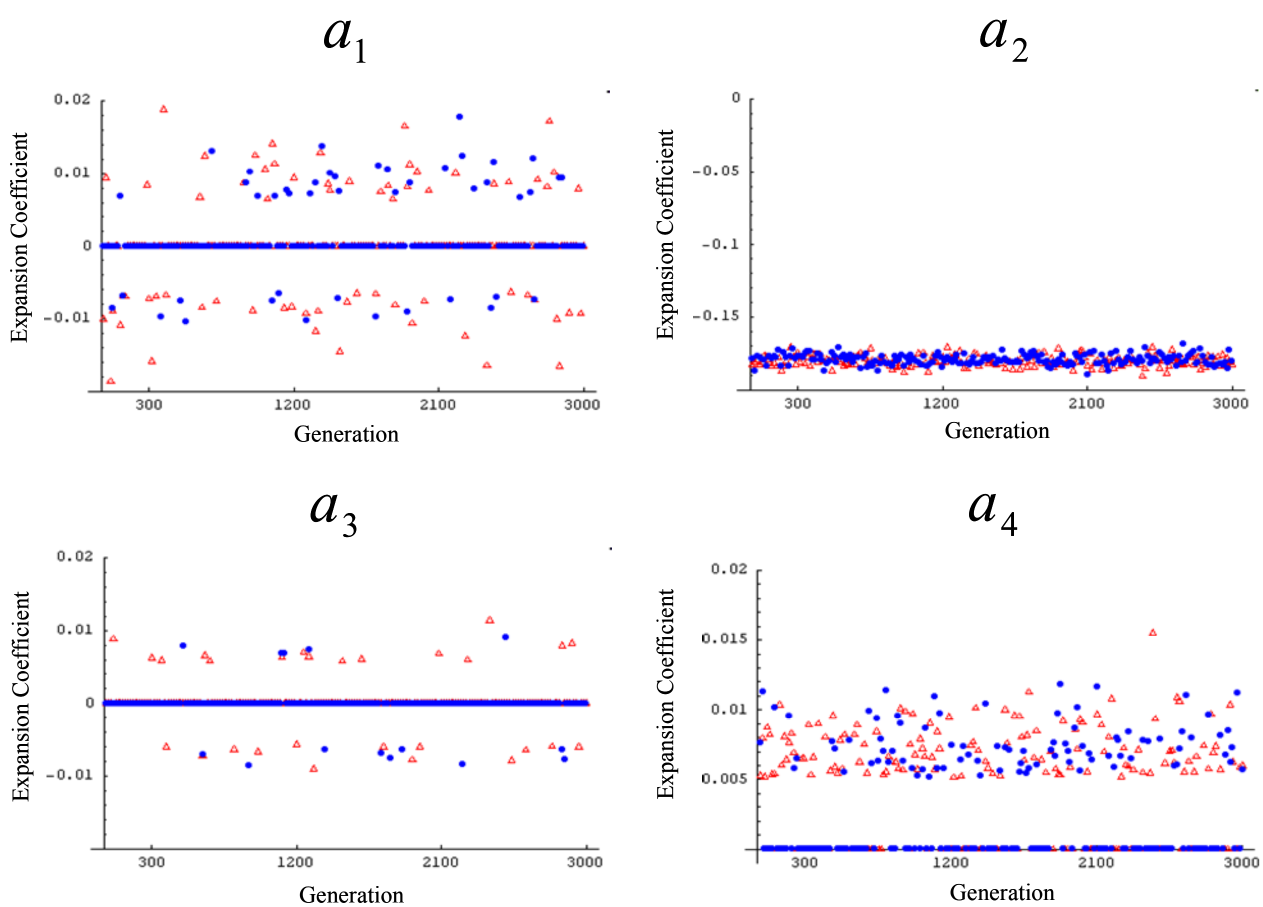

Figure 1. Low order Legendre moments of the fission source by generation for fast-spectrum critical assembly. Results are shown for both the fission bank (red) and collision FET (blue) source sampling routines. Both sets of results were generated in independent Monte Carlo criticality calculations using 3000 generations and 10,000 histories per generation. Legendre moments with a cost-to-benefit ratio greater than 0.95 were considered poorly converged and filtered out of the functional approximation.

For both sampling methods, as the fission source converges, the Legendre coefficients will each converge to a constant value. These constants are the source moments for the converged source distribution. As long as any of the expansion coefficients are showing systematic changes between generations the fission source is not converged and additional generations will be required in the source iteration process.

Figure 1 shows a plot of the first through the fourth Legendre coefficients for the fission bank and FET source convergence methods as a function of the source generation number. As expected, the coefficients for both convergence methods appear to converge immediately to the same values and remain nearly constant over all of the generations. The first and third coefficients appear randomly distributed about zero, while the second and fourth coefficients appear randomly distributed about -0.18 and 0.007, respectively. The empty gaps that appear in the first, third and fourth moments are a result of the R_n^2 filtering used to reduce statistical uncertainty in the functional approximation. Data points that would normally appear within these gaps contain large statistical uncertainties and therefore were filtered out and set to zero.

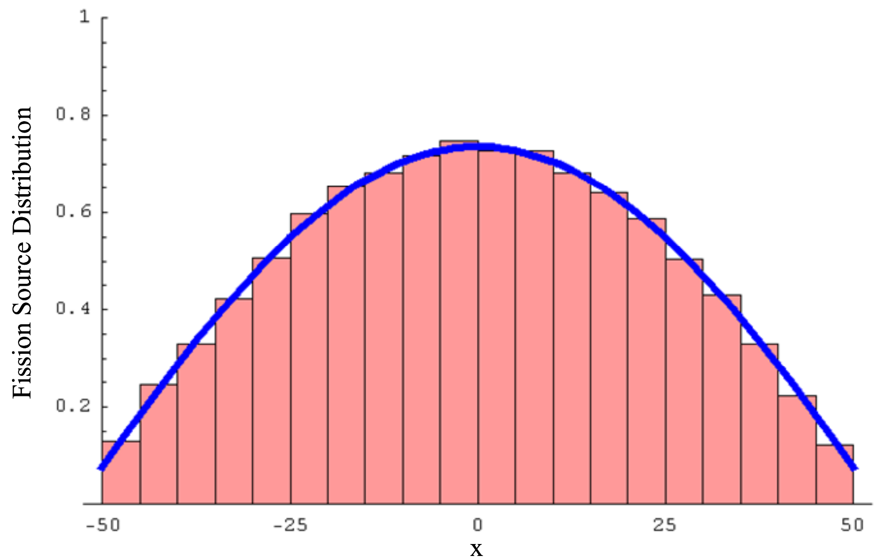

Figure 2. Comparison of spatial source distribution produced by the fission bank and collision FET source sampling routines for the fast-spectrum critical assembly benchmark. The blue line shows the 20th order Legendre approximation to the source shape used by the FET source convergence routine. The red bars show a histogram representation of the source points stored in the fission bank. Results were produced in independent Monte Carlo criticality calculations using 3000 neutron generations with 10,000 histories per generation. Data shown are for the fission source produced after generation 3000.

The width of these gaps can be reduced by either raising the R_n^2 filter threshold or increasing the number of histories per generation.

Based on the coefficient plots, both the fission bank and FET convergence methods produce nearly identical results for this benchmark problem. This is confirmed by a visual comparison between the FET fission source approximation and a histogram representation of the fission bank data taken after cycle number 3000, shown in Figure 2. This comparison highlights the excellent agreement between the fission source shapes produced by the two methods. These results verify that the new FET based method is implemented correctly and that it can be used to accurately estimate the converged fission source distribution for this problem.

3.2 Thermal Reactor Benchmark Problem

The second series of benchmark tests were conducted on a homogenous, one-dimensional, thermal-spectrum critical assembly. The assembly was modeled as a 100 cm bare slab of a homogenized water and uranium mixture with a density of 10.97 g/cc. The ratio of hydrogen to uranium atoms in the mixture was 900:1. The isotopic composition of the uranium was 40% U^{235} and 60% U^{238}. The system was again characterized by a reference criticality calculation using an unmodified version of MCNP4c. As with the previous benchmark, the reference calculation used 500 generations of 10,000 neutron histories, with the first 50 cycles being discarded. The estimated eigenvalue from the simulation was found to be 0.96531 with a relative standard deviation of 0.00010, and the mean free path of neutrons in the simulation was 0.113 cm.

The mean free path of neutrons in the water/uranium assembly is much smaller than the mean free path in the pure uranium assembly. The smaller mean free path means that it will take neutrons in the system many generations to spread between local regions in the assembly. A system like this is commonly referred to as being loosely coupled. In general, source iteration techniques will converge slowly for loosely coupled systems. In fact, the convergence rate depends heavily on how loosely coupled the system under consideration is. The problem with loosely coupled systems is that separate local regions of a single global system often start to converge independently of one another. For example, during the early source generations, an extremely wide reactor may behave neutronically like two, or more, smaller reactors placed side by side. Only after many generations (typically thousands) do these local sources begin to coalesce into the global solution. Unfortunately, for large scale problems, such as commercial reactor cores, running a sufficiently large number of neutron generations to achieve source convergence with current Monte Carlo codes is computationally very expensive. It is hoped that the FET source convergence method can help to accelerate the initial fission source convergence rate for loosely coupled systems, thus reducing the overall computational cost of the eigenvalue calculations.

To test this hypothesis, a series of Monte Carlo eigenvalue calculations were conducted on the water/uranium system. Each calculation used 1000 generations with either 10-, 50-, or 100,000 neutron histories per generation. During each generation both the collision and track length functional expansion tallies were used to estimate sets of 20 Legendre expansion coefficients of the fission source. These coefficients were filtered using a cost-to-benefit ratio threshold of 0.6. Source points for the next generation were either taken directly from the fission bank or sampled from the collision FET or track length FET source approximations. The combination of the three source sampling routines and the three different generation sizes gave a total of nine independent criticality simulations used for this study.

As with the fast-spectrum system, the convergence of the fission source was assessed by examining the convergence of the individual Legendre moments as a function of the generation number. Figure 3 through Figure 5 shows the first through sixth expansion coefficients for all three source convergence methods in the 10-, 50-, and 100,000 histories per generation calculations, respectively. The convergence behavior of the expansion coefficients in the thermal system is strikingly different than that observed in the fast spectrum benchmark. In the fast system the moments appeared to change by a large margin between generations, but they always appeared to be approximately randomly distributed about a mean value. In the thermal system the values of individual moments change relatively little between generations. Furthermore, in the thermal system the moments are obviously not converged during the early cycles. Rather, each moment shows a steady trend from generation to generation. This trending clearly illustrates the convergence of the source shape. Statistical noise in the source shape of the thermal system appears as slow drifts over time in each of the moments, instead of large oscillations about a central value. These slow drifts make it difficult to assess when, and where a particular moment has converged.

10,000 Histories per Generation

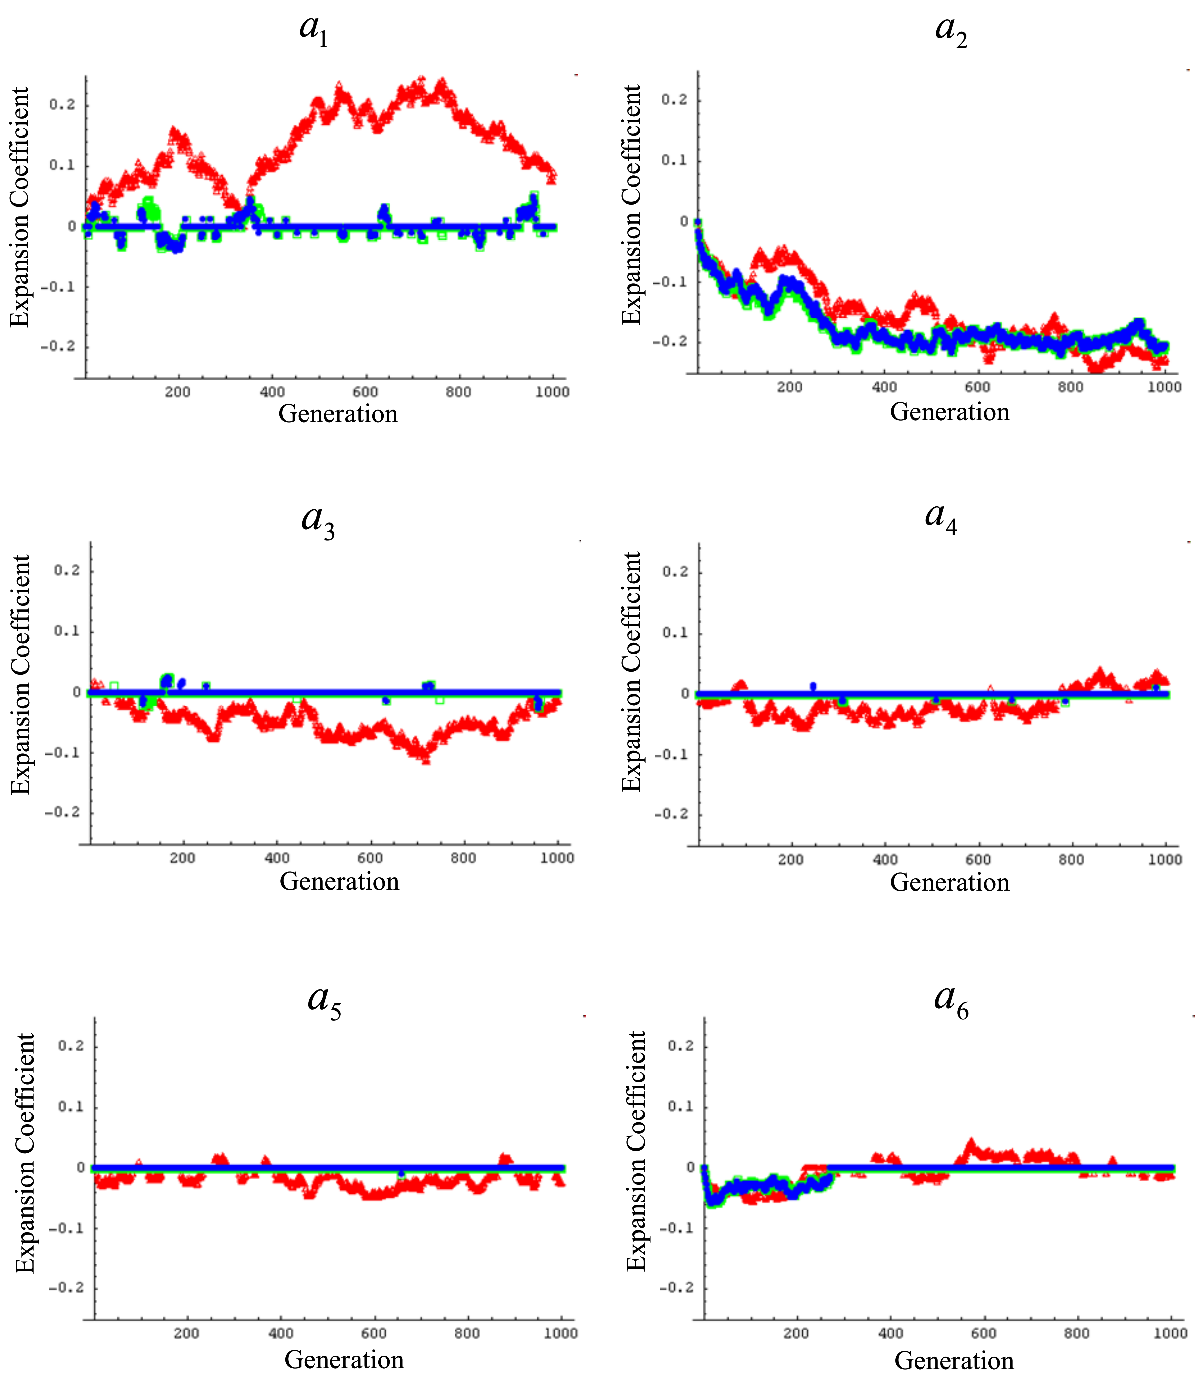

Figure 3. Low order Legendre moments of the fission source by generation for thermal-spectrum critical assembly. Results are shown for the fission bank (red), collision FET (blue), and track length FET (green) source sampling routines. All results were generated in independent Monte Carlo criticality calculations using 3000 generations and 10,000 histories per generation. Legendre moments with a cost-to-benefit ratio greater than 0.95 were considered poorly converged and filtered out of the functional approximation.

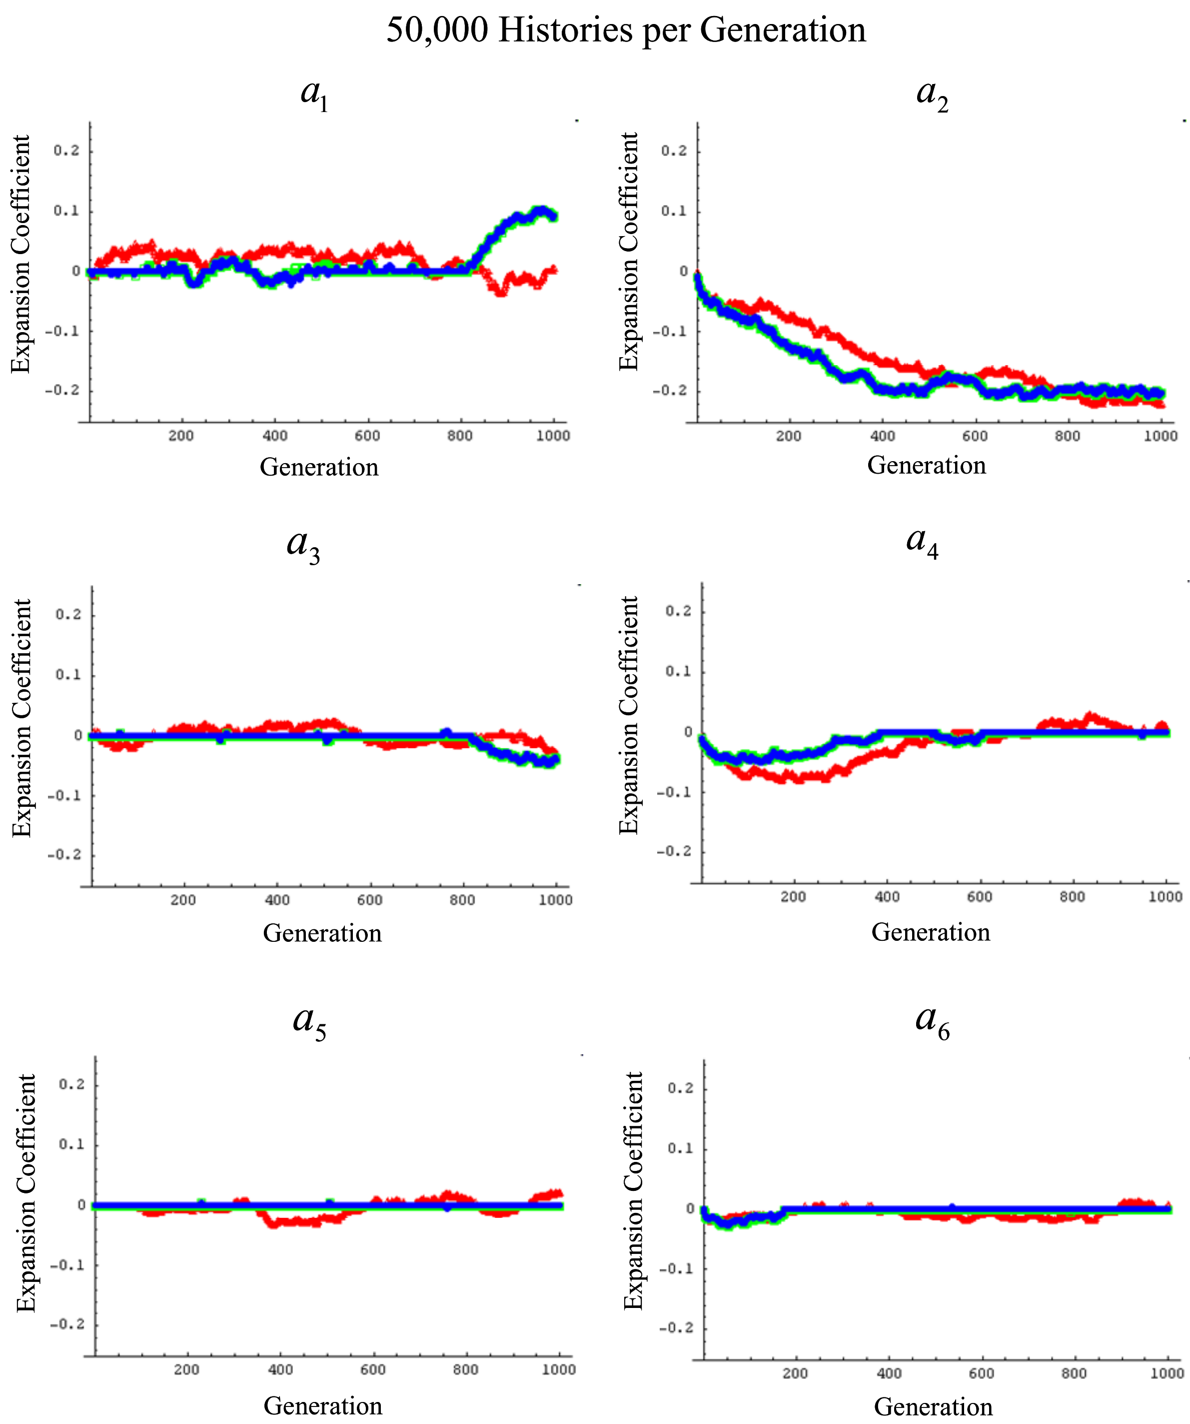

Figure 4. Low order Legendre moments of the fission source by generation for thermal-spectrum critical assembly. Results are shown for the fission bank (red), collision FET (blue), and track length FET (green) source sampling routines. All results were generated in independent Monte Carlo criticality calculations using 3000 generations and 50,000 histories per generation. Legendre moments with a cost-to-benefit ratio greater than 0.95 were considered poorly converged and filtered out of the functional approximation.

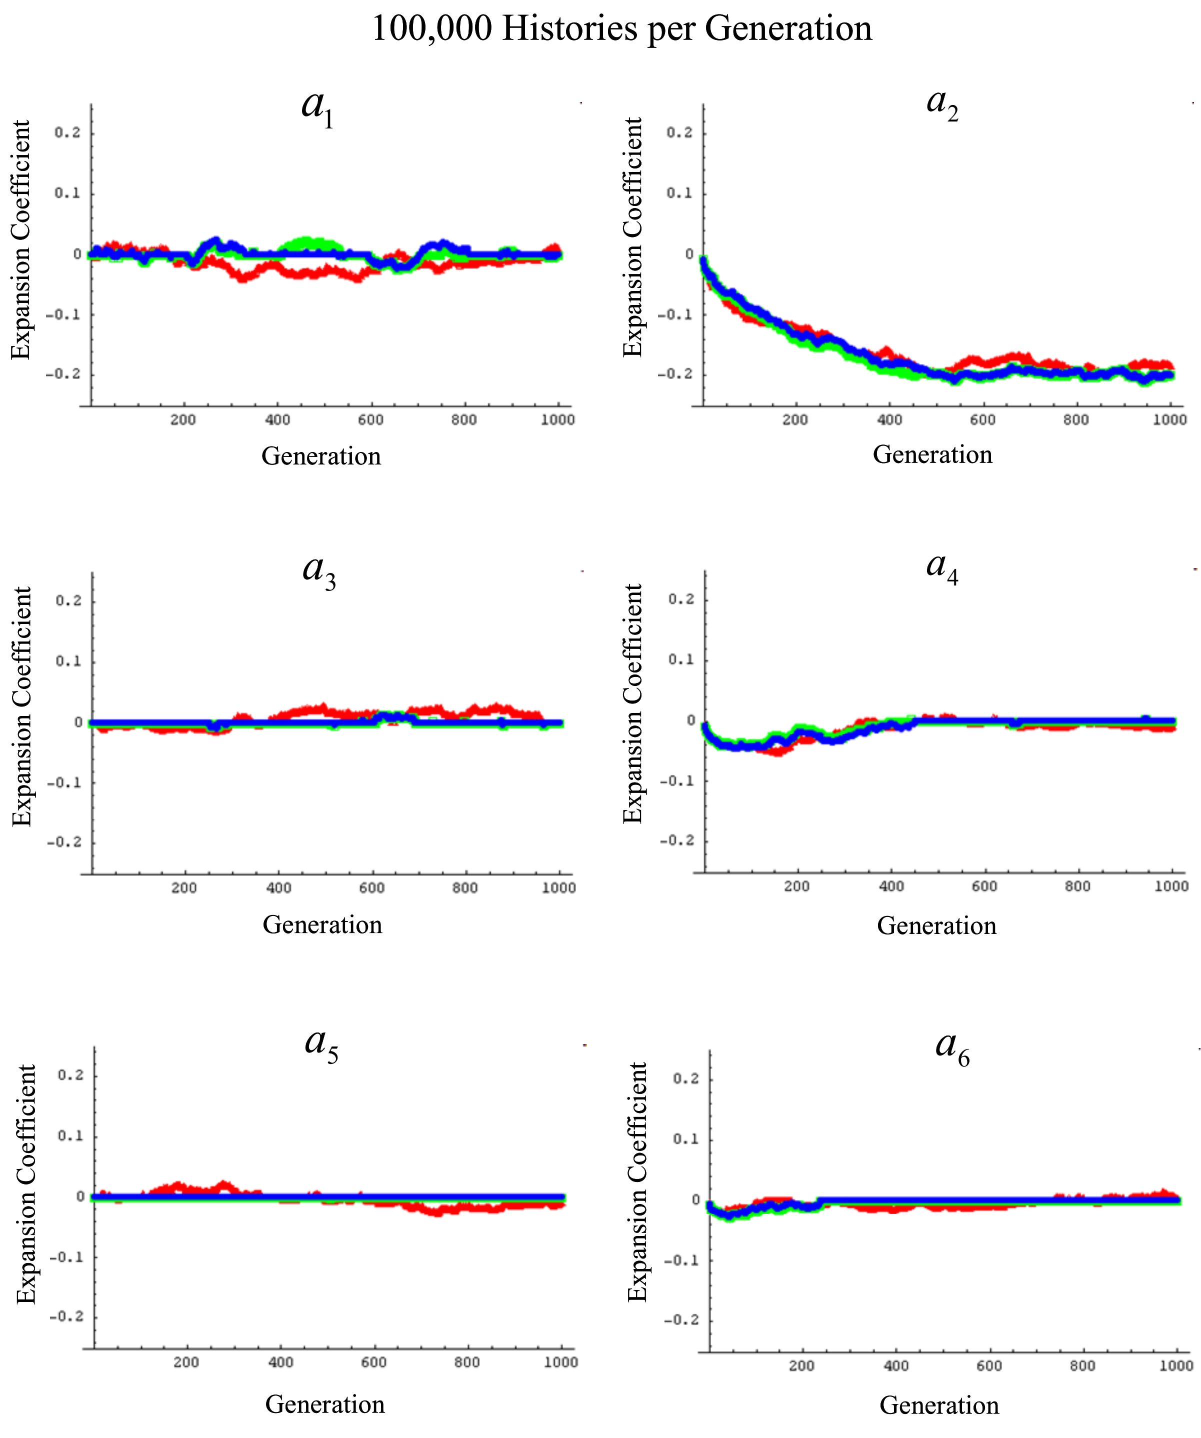

Figure 5. Low order Legendre moments of the fission source by generation for thermal-spectrum critical assembly. Results are shown for the fission bank (red), collision FET (blue), and track length FET (green) source sampling routines. All results were generated in independent Monte Carlo criticality calculations using 3000 generations and 100,000 histories per generation. Legendre moments with a cost-to-benefit ratio greater than 0.95 were considered poorly converged and filtered out of the functional approximation.

Table 1. Exact Legendre expansion coefficients for a cosine fission source distribution across a one-dimensional slab geometry.

Legendre Order (n)	Expansion Coefficient (a_n)
1	0
2	-2.60
3	0
4	0.00903
5	0
6	-0.000161

For this example it is useful to compare the numerical results with an analytical approximation to the source shape, in order to verify that the methods are converging to the correct shape. Because this benchmark problem is a bare, homogenous, critical slab it is reasonable to expect that diffusion theory should apply and the fission source should be roughly a cosine shape over much of the slab. For comparative purposes the first 6 exact Legendre moments for a cosine over the slab are listed in Table 1. We note that the continuous energy Monte Carlo solutions should not be expected to converge to exactly the cosine shape. However, the fact that the dominant second order expansion coefficient appears to converge to a value close to the simple one-group diffusion approximation provides an indication that the methods are working as predicted.

In the simulation with 10,000 histories per generation (Figure 3), both FET source sampling routines clearly outperform the fission bank. The fission bank solution has a large 1st moment, which indicates that the distribution is tilted towards one side of the assembly. The FET source method, by contrast, is able to quickly converge the odd moments to zero, which is physically correct for a symmetric system. In the simulation with 50,000 histories per generation (Figure 4), the FET appears to show faster convergence than the fission bank for the second coefficient (the dominant term). In fact, it appears as though the FET estimates for a_2 reach convergence near generation 400, several hundred cycles before the fission bank estimates. At 100,000 histories per generation (Figure 5) all three source sampling routines appear to show roughly the same performance. However, even with this large number of histories per generation the FET sampling methods seem to do slightly better on the odd moments. For all of these test cases, there appears to be no significant difference between the collision FET and track length FET source convergence methods.

Note that, in some cases, the results appear to converge, but may suddenly drift away from convergence (i.e. get a large first moment component) and then drift slowly back towards the original converged value. This behavior is observed for both the FET and the fission bank

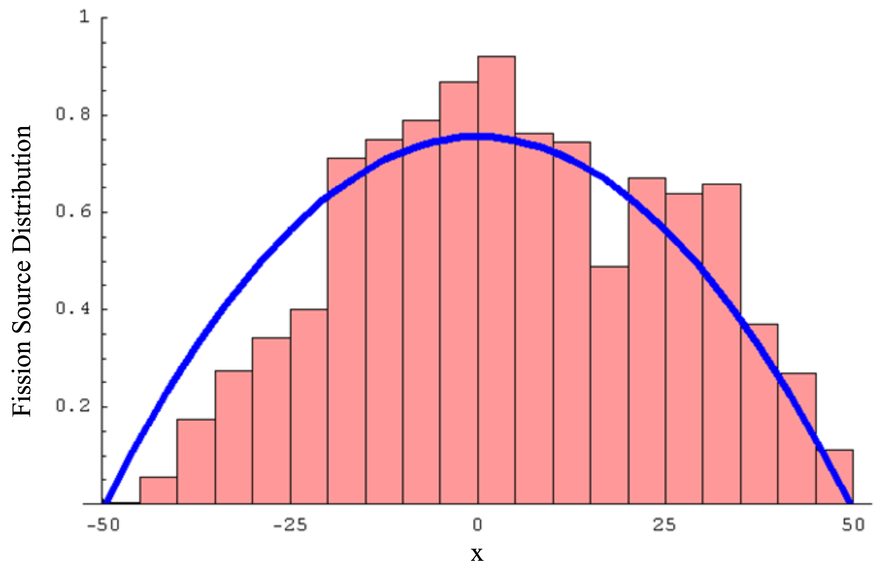

Figure 6. Comparison of spatial source distribution produced by the fission bank and collision FET source sampling routines for the thermal-spectrum critical assembly benchmark. The blue line shows the 20th order Legendre approximation to the source shape used by the FET source convergence routine. The red bars show a histogram representation of the source points stored in the fission bank. Results were produced in independent Monte Carlo criticality calculations using 3000 neutron generations with 10,000 histories per generation. Data shown are for the fission source produced after generation 3000.

methods. The magnitude of this drift is, in some cases, surprisingly large. The causes of these drifts are not well understood. One possible explanation is that the drifts are due to a temporary false convergence towards a higher (non-stable) eigenmode. It is clear that analysis of the individual Legendre moments can provide valuable information about the convergence of the fission source shape.

For smaller numbers of histories per generation, the fission bank has more noise in the source shape than the FET based method. This difference is illustrated by a comparison of the functional approximation and a histogram representation of the fission bank data, shown in Figure 6. To a large extent this reduction in noise is due to the filtering of moments with high statistical uncertainty. In order to study the effect of different filtering thresholds on the source convergence a series of three additional criticality calculations were conducted using the collision FET source sampling method. Each of these additional calculations used 1000 generations with 10,000 neutron histories per generation. The filtering threshold for each of the three simulations was set to 0.3, 0.6 and 0.95, respectively. Plots of the first through sixth expansion coefficients are shown in Figure 7. With the largest threshold, 0.95, the FET allows a large amount of fluctuation in the first and third expansion coefficients, but generally does a good job in quickly converging the important second and fourth moments. Reducing the threshold to 0.6 eliminates the fluctuation in the odd moments, but maintains the rapid convergence of the even terms. In the final simulation, the threshold was further reduced to 0.3. The results show that this threshold is set too low, and as a result, the FET filtered out important a_6 coefficients during the early generations. Without these initial a_6 terms, the fission source actually takes longer to converge than with the other filter threshold values.

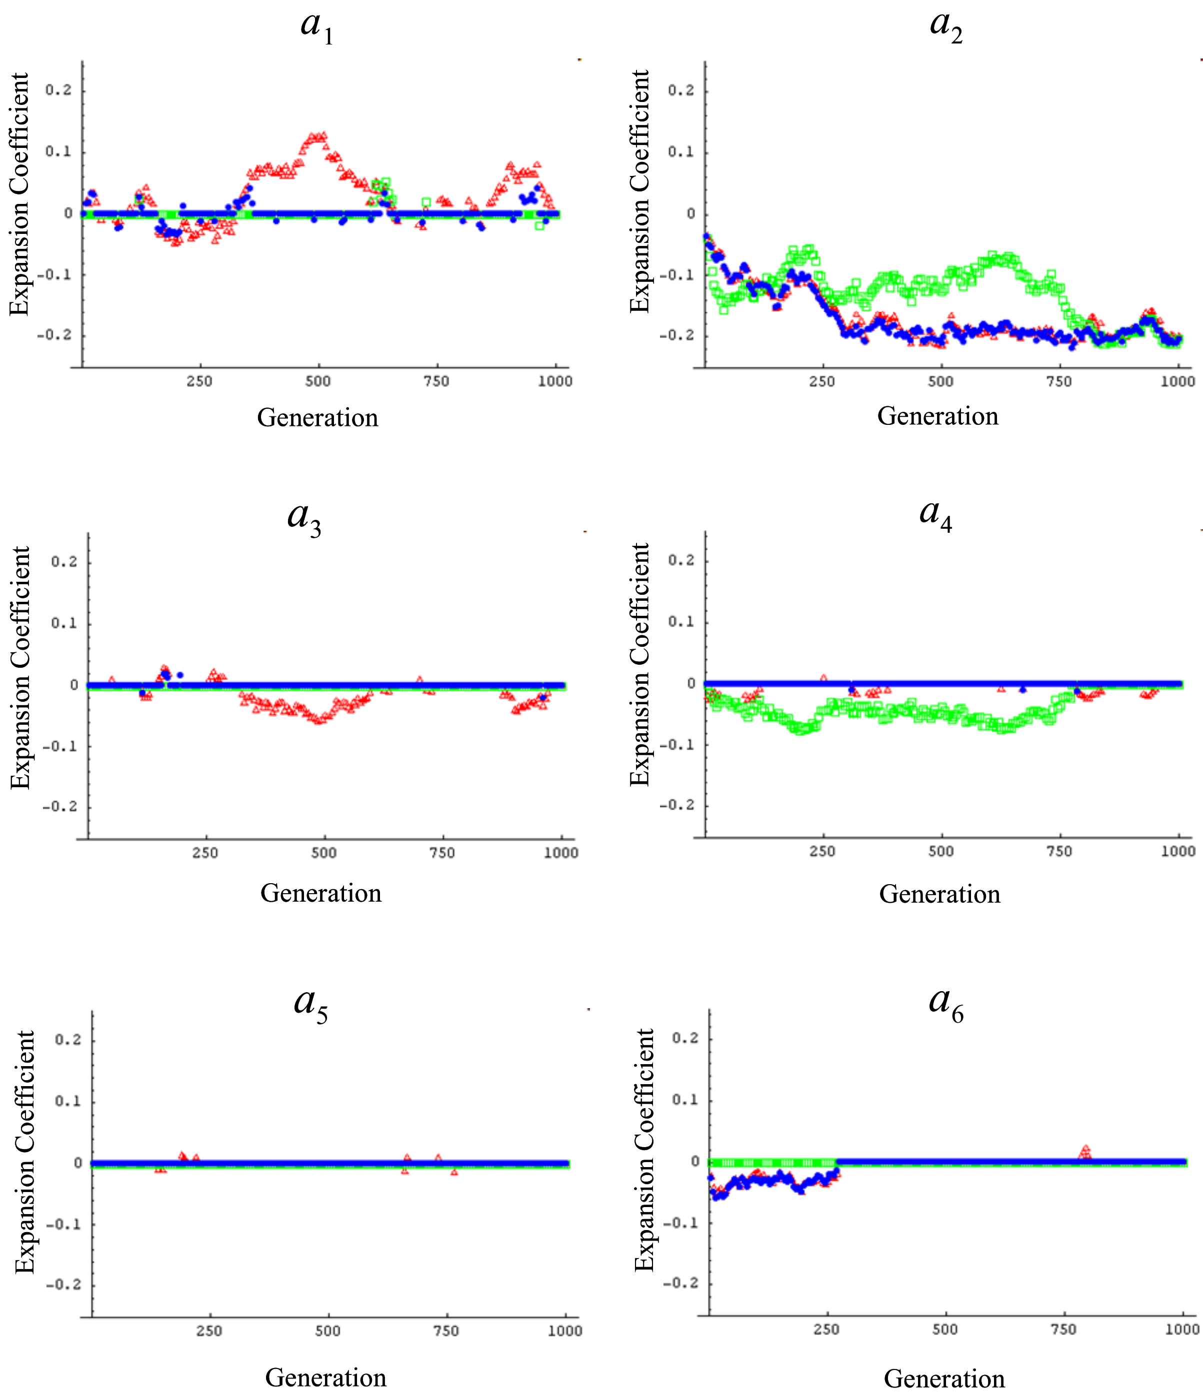

Figure 7. Filtered low order Legendre moments of the fission source by generation for thermal-spectrum critical assembly. Results are shown for the cost-to-benefit ratio filter threshold values of 0.95 (red), 0.6 (blue), and 0.3 (green). All results were generated in independent Monte Carlo criticality calculations using collision FET source sampling with 1000 generations and 100,000 histories per generation.

4 CONCLUSIONS

A functional expansion source convergence method for Monte Carlo criticality calculations has been developed and tested in 1-D. Initial results indicate that the FET based method is a viable alternative to the traditional fission bank approach. The numerical results presented here indicate that the FET source convergence method may even outperform the fission bank method in 1-D homogeneous, loosely coupled systems when small numbers of neutron histories per generation are used. The filtering capability of the method appears to aid in this increased convergence by dampening out much of the statistical noise that is present in the fission source from each generation. Future work on this method will focus in extending it to more realistic (multidimensional and heterogeneous) systems as well as taking into account the energy dependence of parent neutrons on the fission spectrum with multi-dimensional expansions in space and energy.

The use of the FET based source convergence method offers several advantages over the traditional fission bank approach. As noted previously, the data can be filtered to eliminate some of the statistical noise present in the results. Also, the new method is not limited to a fixed number of source points per generation like the fission bank. As many independent samples as desired can be taken from the functional approximation fission source without concern over statistical correlation that can arise from reusing fission bank samples in a single generation. This property of the FET based method means that it is possible to use a variable number of neutron histories per generation. In fact, one promising idea is to allow each generation to run until a number of expansion coefficients have passed some user defined convergence test, before starting the next generation. These advantages of the FET source convergence method make it an attractive option for accelerating global source convergence in early generations of very large problems.

Aside from accelerating the source convergence, tallying the expansion moments of the source distribution also provides a new and interesting method for studying source convergence behavior. The results presented in this paper have illustrated some interesting features of the source convergence in both the fission bank and FET based methods. One of these features, the slow drift of the source away from and back to an apparently converged shape, is not yet understood. However, it appears that significant insight into this behavior can be obtained by examining the generation-to-generation behavior of the individual expansion coefficients.

5 ACKNOWLEDGMENTS

The first author would like to acknowledge the educational support provided by the Naval Nuclear Propulsion Fellowship Program sponsored by Naval Reactors Division of the U.S. Department of Energy. The authors would also like to thank Natalie Durgin for running many of the Monte Carlo simulations as a part of her NASA SHARP summer research apprenticeship.

6 REFERENCES

1. J. F. Briesmeister, *MCNP - A General Monte Carlo N-Particle Transport Code*, LA-13709-M, Los Alamos National Laboratory, Los Alamos (2000).

2. J. Spanier, "Monte Carlo Methods for Flux Expansion Solutions of Transport Problems," *Nucl. Sci. Eng.*, **133**, pp. 73-79 (1999).

3. D. Legrady and E. J. Hoogenboom, "Visualization of Space-Dependency of Responses of Monte Carlo Calculations Using Legendre Polynomials," *Proceeding of Physor 2004-The Physics of Fuel Cycles and Advanced Nuclear Systems: Global Developments*, Chicago, Illinois, April 25-29, CD-ROM, (2004)

4. D. P. Griesheimer, *Functional Expansion Tallies for Monte Carlo Simulations*, Ph. D. Dissertation, University of Michigan, Ann Arbor. (2005).

5. D. P. Griesheimer and W. R. Martin, "Estimating the Global Scalar Flux Distribution with Orthogonal Function Expansion," *Trans. Am. Nucl. Soc.*, **89**, pp. 462-464 (2003).

6. D. P. Griesheimer and W. R. Martin, "Two Dimensional Functional Expansion Tallies for Monte Carlo Simulations," *Proceeding of Physor 2004-The Physics of Fuel Cycles and Advanced Nuclear Systems: Global Developments*, Chicago, Illinois, April 25-29, CD-ROM, (2004)

7. D. P. Griesheimer, W. R. Martin, and J. P. Holloway, "Estimation of Flux Distributions with Monte Carlo Functional Expansion Tallies," *Proceeding of Tenth International Conference on Radiation Shielding (ICRS-10)*, Madeira, Portugal, (2004).

8. D. P. Griesheimer and W. R. Martin, "Monte-Carlo Based Angular Flux Response Functions," *Trans. Am. Nucl. Soc.*, **89**, (2003).

9. B. L. Beers and V. W. Pine, "Functional Expansion Technique for Monte Carlo Electron Transport Calculations," *IEEE Trans. Nucl. Sci*, **23**, pp. 1850-1856 (1976).

10. A. Noel and H. S. Wio, "A New Series-Expansion Approach in Monte Carlo: Application to Neutron Shielding," *Ann. Nucl. Energy*, **11**, pp. 225-227 (1984).

11. W. L. Chadsey, C. W. Wilson, and V. W. Pine, "X-Ray Photoemission Calculations," *IEEE Trans. Nucl. Sci*, **22**, pp. 2345-2350 (1975).

12. F. B. Brown, D. P. Griesheimer, and W. R. Martin, "Continuously Varying Material Properties and Tallies for Monte Carlo Calculations," *Proceeding of Physor 2004-The Physics of Fuel Cycles and Advanced Nuclear Systems: Global Developments*, Chicago, Illinois, April 25-29, CD-ROM, (2004)