*Nuclear Mathematical and Computational Sciences: A Century in Review, A Century Anew*
Gatlinburg, Tennessee, April 6-11, 2003, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2003)

# A CASE STUDY IN MANUAL AND AUTOMATED MONTE CARLO VARIANCE REDUCTION WITH A DEEP PENETRATION REACTOR SHIELDING PROBLEM

**Herschel P. Smith**
Duke Energy Corporation
422 South Church Street
PO Box 1006
Charlotte, NC 28242
hpsmith@duke-energy.com

**John C. Wagner**
Oak Ridge National Laboratory
PO Box 2008
Oak Ridge, TN 37831-6370
wagnerjc@ornl.gov

## ABSTRACT

Certain reactor transients cause a reduction in moderator temperature, and hence, increased attenuation of neutrons and decreased response of excore detectors. This decreased detector response is of concern due to the credit assumed for detector-initiated reactor trip to terminate the transient. Explicit modeling of this phenomenon presents the analyst with a difficult problem due to the dense and optically thick neutron absorption media, given the constraint that precise response characteristics must be known in order to account for this phenomenon. The solution in this study was judged to be the use of Monte Carlo techniques coupled with robust variance reduction to accelerate problem convergence. A fresh discussion on the motivation for variance reduction is included, followed by separate accounts of manual and automated applications of variance reduction techniques. Finally, the results of both manual and automated variance reduction techniques are presented and compared.

*Key Words*: Monte Carlo, variance reduction, radiation transport, reactor pressure vessel.

## 1. INTRODUCTION

Increasing computational speed enables Monte Carlo analysts to simulate enough particle histories to consider increased computer run time in lieu of variance reduction. However, some problems remain challenging enough to require the effective use of variance reduction techniques in order to achieve reliable results in a timely manner so as to meet the needs of end-users. One such problem is described in this paper, along with the application of manual and automated variance reduction and the resultant success of each.

The problem involves the simulation of excore neutron detector response to changing conditions in regions inside the reactor vessel resulting from a transient, since some transients assume credit for automatic initiation of reactor trip from neutron flux. This problem affects the Steam Line

Break (SLB) from hot full power, where the decrease in moderator temperature causes (a) an increase in reactor power due to moderator temperature feedback (from a negative moderator temperature coefficient [MTC]) and (b) increased attenuation of neutrons and resulting decreased power indication by the excore neutron detectors. Although the phenomenon of transient-initiated changes in excore neutron detector response has been known to the nuclear power industry [1] since 1996, the industry response to this phenomenon has heretofore been primarily aimed at operational concerns. However, failure to account for this phenomenon in analytical space would result in non-conservative safety analysis results when termination of the transient depends upon automatic reactor trip initiated by excore neutron detector indication.

In addressing this phenomenon for nuclear installation safety analyses, an understanding of general response characteristics is not adequate; excore detector response must be quantitatively simulated and included in transient analysis calculations. The modeling of excore detector response presents a non-trivial problem for the analyst due to geometrical complexity, optically thick regions, dense neutron absorption media, and the presence of neutron absorbers (consequential poisons, such as fission products, and intentional poisons, such as chemical shim for suppression of excess core reactivity during beginning of cycle operations). The Monte Carlo method is considered to be the most accurate method available for solving complex radiation transport problems. However, due to its nature of simulating individual particle histories and inferring the average behavior of the particles in the system from the average behavior of the individually simulated particles, the Monte Carlo method is computationally intensive for deep penetration problems. The realism in transport simulations in the MCNP Monte Carlo computer code [2] makes it well-suited for this analysis, but robust variance reduction is an integral part of the analysis to ensure that all problem-significant phase space has been properly sampled and that the problem solution has converged. Hence, analysis of this problem has focused considerable effort on achieving good statistical performance. This work is described below, and results of manual and automated efforts to yield variance reduction parameters are compared.

## 2. MODELING APPROACH AND RESULTS

### 2.1 Geometry Setup & Tally Description

The MCNP geometry for this problem is specified as a quarter-core model to make use of problem symmetry, since there are four excore neutron detectors, each detector being centered 45° off the main axes of the core. The input description includes specific modeling of the fuel pellet stack, fuel pin gap and fuel cladding, guide tubes, instrument tubes, and interstitial moderator (light water) constituting a 17×17 pressurized water reactor fuel assembly. The fuel assembly description is filled into the fuel region inside the core baffle plates as a repeating lattice; the baffle plates are specified to as-built dimensions. In the axial direction, the geometry is truncated at the top and bottom of the fuel pins (i.e., the geometry does not include fuel assembly end fittings, upper or lower core plates or core upper internals). Additionally, as shown in Figure 1, the innermost region of the core is specified as a void. This phase space elimination was built into the input description for the purpose of assisting problem convergence. Thus, phase space elimination in the problem setup was seen to be a rudimentary but valuable variance reduction tool to avoid consideration of parts of the problem for which contribution to the tally was judged to be insignificant. The elimination of a central cylinder of the fuel region was done after careful consideration of prior analyses [3,4] showing the relative contributions to ex-vessel results from the various fuel assembly locations.

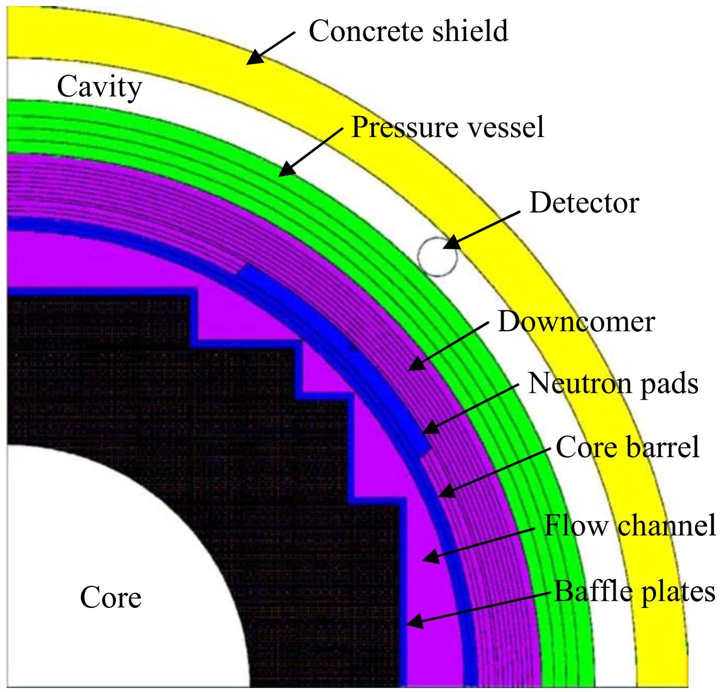


**Figure 1. MCNP model for McGuire Nuclear Station quarter-core geometry**

Beyond the core baffle plates is the flow channel region inside of the core barrel but outside of the fuel region. The core former plates (which are attached to the core barrel, and to which the baffle plates are bolted) are located in this region, and are included in the geometry input. Next, the core barrel and neutron pads are modeled as a set of concentric cylinders, as is the subsequent reactor coolant downcomer and reactor pressure vessel regions. The concrete shield wall behind the detector is included in order to simulate neutron backscatter to the detector. Finally, the detector is modeled as a cylindrical volume located between the reactor pressure vessel and the concrete shield wall. A track-length estimate of cell flux is performed for this volume with different moderator densities in the problem. The average particle flux in a cell can be written as:

$$\overline{\phi}_V = \frac{1}{V}\int dE \int dt \int dV \int d\Omega\, \psi\left(\vec{r}, \hat{\Omega}, E, t\right). \tag{1}$$

MCNP estimates $\overline{\phi}_V$ by summing $WT_l\,/V$ for all particle tracks in the cell, where $V$ is cell volume, and $W$ and $T_l$ are particle weight and track-length, respectively. The excore neutron detectors are uncompensated ion chambers that operate based on the $B^{10}(n,\alpha)Li^7$ interaction. Hence, the tally for this problem was modified (with a tally multiplier) to score the reaction of interest for the excore detector, specifically, $^{10}B$ absorptions.

The track-length estimate of cell flux compared favorably with a point detector tally. The MCNP point detector tally is a next-event estimator, computing flux at an arbitrary point as follows:

$$\Phi(\vec{r}, E, t, \mu) = \frac{W p(\mu) e^{-\lambda}}{2\pi R^2}, \tag{2}$$

where

| | |
|---|---|
| W | = particle weight, |
| p(μ) | = value of probability density function at μ, the cosine of the angle between the particle trajectory and the direction to the detector, |
| λ | = total number of MFPs integrated over the trajectory from source or collision point to detector, |
| R | = distance from source or collision event to detector. |

The point detector tally was used as a means of variance reduction during the manual development and application of variance reduction parameters. This is discussed later in Section 2.3.3 of this paper.

## 2.2. Neutron Source Specification

Various options were considered for the neutron source specification for this problem. A criticality calculation could have been performed, and the resulting neutron direction, energy, and weight saved to a file for use in subsequent calculations. This option could have been used to write a source file at a given surface, such as the core barrel or reactor vessel, and this source could then be used as input to a transport calculation that tallied at the excore detector location. Another option considered was the use of the watt-fission spectrum in the fuel region as a fixed source neutron transport problem, along with allowing fissions to occur in the fuel. The option selected for this simulation was to use fixed-source neutron spectra from the CASMO computer code [5] in 40 energy groups (specifically computed by CASMO for this lattice design), while using the option in MCNP to turn off fissions, and treat the neutron transport as a shielding problem rather than modeling a critical system. This choice was deemed to be the most amenable to adjustments in source probability definition to address the effects of axial and radial power shape on the neutron attenuation factor and detector response, adjustments in the moderator density in the various regions, and to source energy biasing.

This choice was evaluated and justified by comparing the results of MCNP and CASMO calculations. For the MCNP calculation, the neutron source energy was specified to be the watt-fission spectrum, fissions were allowed in the fuel, and the energy-dependent flux was tallied on fuel cladding surfaces (using a track-length estimate) in the same energy group structure used in the CASMO calculation. The MCNP tally results were normalized and plotted against results from CASMO, and the comparison was deemed to be favorable, as shown in Figure 2 below.

Finally, two moderator density statepoints were evaluated for each of two source options as discussed above: (1) the watt-fission spectrum with fissions allowed, and (2) the CASMO multigroup energy spectrum with fissions turned off. The neutron attenuation factors for the two source options were compared and found to be equivalent (within statistical uncertainty of approximately one percent), thus demonstrating that the choice of the fixed source model would yield results comparable to the more computer-intensive solution with a fission system.

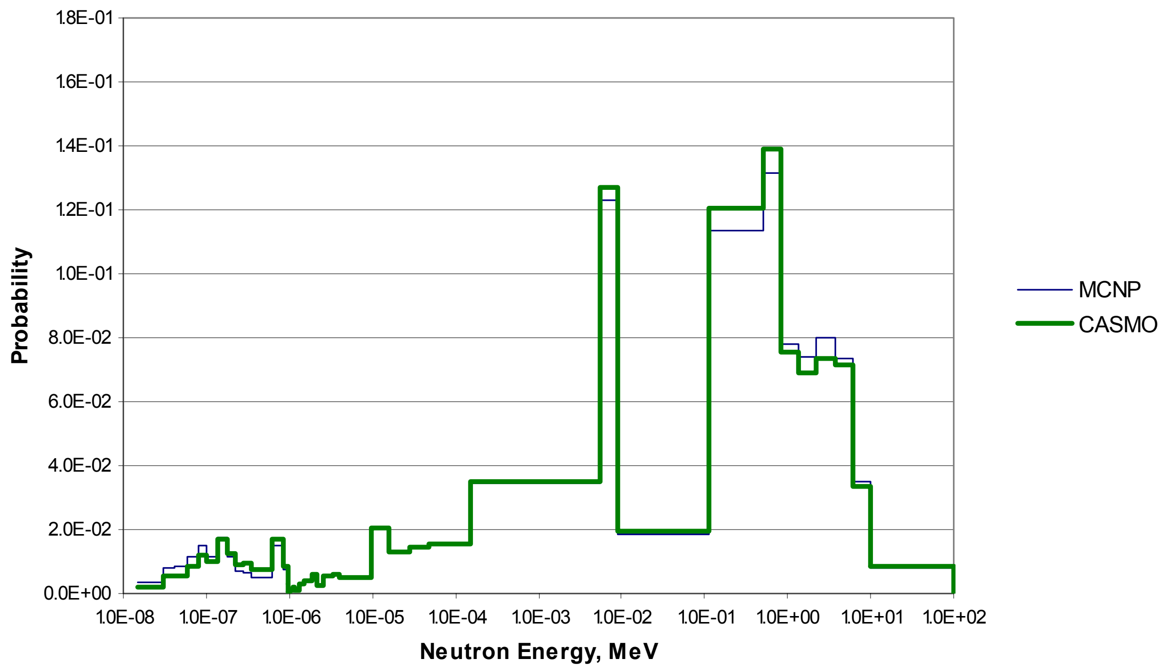


**Figure 2. Comparison of CASMO & MCNP neutron energy spectra**

Following the variance reduction efforts (described in the next section), the effects of variations in the radial and axial source distributions on the attenuation factor were analyzed to assess core power shape as an input to the attenuation factor. Various source shapes were considered (such as cosine axial power shape versus homogenous axial power shape, and realistic radial power shape based on core flux map results versus homogenous radial power shape). It was found that while the total tally mean could be affected by variations in the power shapes, the change in the mean at selected moderator density statepoints was not affected by the power density or shape assumptions.

## 2.3 Variance Reduction

### 2.3.1 Test for statistical efficiency

In all variance reduction efforts, the criteria used to judge success was the so-called Figure-of-Merit (FOM), which is defined as:

$$FOM = \frac{1}{RE^2 \times T}. \tag{3}$$

Where *RE* is the relative error for the sample mean, and *T* is the total computer time taken to simulate *n* histories.[2,6] As the $RE^2$ should be inversely proportional to *T*, the *FOM* value should be approximately constant for a given set of problem parameters. This quantity is deemed to be the appropriate means to ascertain the effectiveness of variance reduction efforts, since it includes consideration of both the resultant relative error of the sample mean and the computer time necessary to achieve this relative error. Variance reduction techniques will usually cause increased computer time per history, but the intention is that the associated reduction in $RE^2$ is greater than the increase in T, thus achieving greater statistical efficiency.

### 2.3.2. Variance reduction motivation

Most assessments of variance reduction motivation address the desired balance between user time to perform the variance reduction and computer time saved as a result of the user efforts. With increasing computer speed, it is tempting to take a diminutive view of variance reduction, especially since computer time is essentially free (with PC computing) and analyst time is expensive. In the abstract, such assessments seem compelling, but in practice can be moderately or even severely misleading. Several anecdotal justifications for robust variance reduction will be offered below, given from the perspective of a practitioner rather than a code developer, and then these observations will be summarized in three categories.

In lieu of a significant expenditure of analyst time to perform variance reduction, the approach sometimes taken to achieving reliable Monte Carlo results is simply to run the problem for longer time. However, without having achieved reliable results using variance reduction, it can be difficult to ascertain the existence of proper sampling of all significant phase space. Under-sampling can lead to pathological features of the problem, where passing statistics are seemingly achieved with the problem, until scores from un-sampled or under-sampled parts of the problem are tallied (low coverage rates are generally the result of too few large $x_i$ [scores, or observations] being observed, not too many [6]). Tallying these scores can result in dramatic increase in the problem variance, along with possible increase in the calculated mean. Poor statistical performance can cause the need to run the problem for a protracted period of time before these pathological features become manifest, thereby wasting not only analyst time to perform poor or mediocre variance reduction, but also computer time.

Moreover, even if the analyst has access to powerful computing resources and is willing to simply run the problem until achieving passing statistics, this tradeoff of analysis time versus computing time can cause unintended consequences. In regulated nuclear activities (i.e., within

Department of Energy, Nuclear Regulatory Commission, Environmental Protection Agency, or Department of Transportation jurisdictions), internal corporate or governmental quality assurance (QA) procedures usually require a peer review of analyses, and these peer reviews often result in questions and issues that can only be addressed by (sometimes an unforeseen number of) iterations and permutations on input, assumptions, problem boundary conditions and modeling choices. For example, in the case study that is the subject of this essay, the results for various moderator density statepoints demonstrated that the attenuation factor computed for use in transient analysis calculations is not a linear function with Reactor Coolant System temperature. This finding necessitated evaluations at multiple moderator density statepoints, each set of moderator density statepoints being evaluated for different modeling choices (such as permutations on core radial and axial power shape). Unintended consequences of poor or mediocre variance reduction efforts can be expenditure of the same protracted computer time for each unforeseen computer run, thus multiplying the computational inefficiency.

Finally, experienced and seasoned practitioners have often had to amend and revise existing analyses to address new regulatory concerns, expand the scope of application of analyses, answer new questions and concerns, and correct minor errors unintentionally introduced into original analyses in spite of painstaking preparation and review. The legacy turned over to successors in the form of input development and documentation, models, and calculation files is an important aspect of traceability, reproducibility and maintainability. At Duke Energy, the cornerstone philosophical tenet of nuclear calculation QA is supplying sufficient documentation and calculational tools to allow a future analyst to comprehend and reproduce the subject analyses, without reference to, or discourse with, the original analyst. In the case of Monte Carlo calculations, the experienced analyst will desire to turn over a legacy of statistical efficiency to his successors in order to facilitate future use.

These observations may be summarized in the following motivations for variance reduction:

1. Confidence in results. Unless a Monte Carlo problem can be considered to be properly sampled and converged, the results cannot be trusted and used, especially as regards nuclear safety related analyses.

2. Timeliness of analyst response to requests for information. Protracted computer analysis time may be economically inexpensive, but end-users rarely are in a position to wait for a protracted amount of time to apply analysis results. Early knowledge of results usually means extended time to address permutations and make decisions on mode(s) of application of the results.

3. Legacy. Ability to efficiently modify inputs, assumptions and boundary conditions and adapt and adjust models to new problem variables facilitates use of the models in the future.

Based on the above discussion and the computationally challenging nature of the problem considered herein, effective variance reduction was deemed to be critical to this analysis. Consequently, an earnest effort was made to effectively utilize the variance reduction techniques available in MCNP by manually developing the required variance reduction parameters and iteratively applying the cell-based and mesh-based weight-window generators. At the completion of this “manual” effort, where additional efforts yielded diminishing returns in terms of computational efficiency and reliability, a separate effort was initiated that utilized a recently developed code [7] that automatically generates variance reduction parameters for MCNP, based on three-dimensional (3-D) deterministic adjoint functions. The two efforts and approaches, and the subsequent results, are described in the following sections.

**2.3.3. Manual application of variance reduction techniques**

Many of the variance reduction techniques available to the user in the MCNP code [2] were used in this analysis. The specific techniques are briefly discussed below, roughly in the order in which they were applied. During the process of applying each variance reduction technique, a number of MCNP calculations were performed to incrementally evaluate effectiveness, adjust the required parameters, and subsequently, develop effective values for the required parameters. To enable appropriate evaluation, these incremental MCNP calculations were each allowed to run for approximately 30 CPU minutes.

Geometry splitting and Russian roulette were used in parts of the problem outside the fuel region in order to start the process of moving particles toward the tally region. To effectively utilize this technique, many of the geometry cells were subdivided into numerous smaller cells. These subdivisions were preserved in the geometry input description throughout the variance reduction efforts in order to provide tools to assess the relative worth of decisions and success of efforts. The importance factors were adjusted by trial and error with multiple MCNP runs to generate a particle-track profile that decreased "gracefully" towards the outer regions of the problem geometry, with no dramatic variations in the number of particle tracks in adjacent geometry cells.

The exponential transform increases particle walks to move toward a preferred direction by artificially reducing the macroscopic cross section in the preferred direction and increasing the cross section in the opposite direction according to:

$$\Sigma_t^{\,*} = \Sigma_t\left(1 - p\mu\right), \qquad (4)$$

where

- $\Sigma_t$* = artificially adjusted total cross section,
- $\Sigma_t$ = true total cross section,
- p = the exponential transform parameter used to vary the degree of biasing,
- μ = cosine of the angle between the preferred direction and the particle's direction.

The problem described in this essay is a "deep penetration" problem. It was difficult to achieve good sampling at the periphery of the problem geometry (where the tally is being performed) due to the optical thickness, the dense absorption media, and the high absorption cross sections associated with the shielding and absorption media (i.e., $UO_2$ fuel matrix, moderator and stainless steel pressure vessel). The exponential transform proved to be a powerful and effective tool during the initial stages of variance reduction due to (1) the ability to move particles towards the tally region through dense media and thereby generate scores and (2) the ability to properly sample geometrical phase space, thereby supplying the necessary information for the automated variance reduction features used later (i.e., the MCNP weight window generators). Use of the exponential transform alone caused an increase in the FOM by more than an order of magnitude. The exponential transform was used only in the fuel region. The exponential transform parameter (so-called "stretching parameter") was finally set through trial and error and iterative MCNP runs until the highest FOM was achieved. Initial estimates of the optimum value for this parameter proved to be fairly accurate.

Next, source energy biasing was used, although the source parameters were merely rudimentary estimates, and thus only marginally effective. After geometric splitting, exponential transform, and source energy biasing, the MCNP cell-based weight window generator, which is a stochastic automated variance reduction generator, was used to generate weight window values. During the random walk simulation, the weight window generator estimates particle importance (with respect to a specified tally) in a given space-energy region as the ratio of the total tally score from all particles entering the region and the total weight entering the region. The weight window values are then calculated inversely proportional to the importance estimates. To obtain an importance estimate for a given region, it is necessary for particles to enter that space-energy region and subsequently contribute to the tally of interest. The weight window technique is a space- and energy-dependent facility by which splitting and roulette are applied. An advantage of the weight window technique over geometry splitting and Russian roulette is that the particle weight games are played at both boundary crossings and collision sites, whereas splitting/roulette are only played at boundary crossings. The importance factors for splitting/roulette and the cell-based weight windows were adjusted by trial and error through iterative MCNP runs to supply a track distribution that had no dramatic step changes between adjacent regions. This ensured that sufficient sampling of phase space occurred for the code to generate meaningful and efficient cell-based weight window parameters.

A point detector is a deterministic estimate (from each event) of the flux at a point in space, using equation 2 above. Contributions to the point detector tally are made at source and collision events throughout the random walk.[2] Being a deterministic estimation of flux makes this tally a useful variance reduction tool, since contributions to the tally will be made from all parts of the system in which collisions occur and/or source particles are started. This increases the contributions to weight window estimates by the weight window generator, as compared to the use of the track-length estimate of cell flux. Therefore, during all phases of the manual application of variance reduction, a point detector was used. The final solution (to develop neutron attenuation factors) used a track-length estimate of cell flux after development of the variance reduction parameters, which proved to increase the final FOM by a factor of approximately 2 over the point detector.

It was intuitively predicted that the concrete behind the excore neutron detector would be significant in terms of backscatter and contribution to the final tally. This intuition proved to be correct, and forced collisions were used to ensure proper sampling of the concrete. The forced collision method is a variance reduction scheme that increases sampling of collisions in specified cells, splitting particles into collided and uncollided parts. The collided part is forced to collide within the current cell. The uncollided part exits the current cell without collision with weight

$$W = W_o e^{-\Sigma_t d}, \qquad (5)$$

where

$W_o$ = particle weight before forced collision,
d = distance to cell surface in the particle's direction,
$\Sigma_t$ = macroscopic total cross section of the cell material.

The collided part has weight $W = W_o(1 - e^{-\Sigma_t d})$, and collision distance *x* is

$$x = -\frac{1}{\Sigma_t} \ln\left[1 - \xi(1 - e^{-d\,\Sigma_t})\right]. \qquad (6)$$

Without the use of forced collisions the concrete was poorly sampled, leading to pathological problems with the solution (erratic error estimates).

Finally, the MCNP mesh-based weight window generator was used repeatedly to develop variance reduction parameters (weight windows) for the arbitrary rectangular mesh that had been specified. The mesh specifications were developed by iterations on mesh size and number of energy groups (i.e., mesh sizes were reduced until FOM no longer increased). It proved difficult to sample each mesh cell adequately to generate viable variance reduction parameters, and the weight window file was manually modified with user-assisted weight window parameters. The FOM decreased on all further attempts to achieve increased efficiency with more than two energy groups. The final variance reduction parameters included:

1. Exponential transform for the fuel region;
2. Forced collisions for the concrete behind the excore neutron detector;
3. Source energy biasing;
4. Mesh-based weight windows for two energy groups; and
5. Implicit capture was turned off in the transport calculation.

This last feature (turning off implicit capture) did *not* undesirably kill particles (in the thermal regime) before they participated in reactions important to the problem, given that the source definition was not specified for a critical system.

The variance reduction parameters yielded by manual efforts were used in MCNP runs that were judged to be protracted enough to demonstrate that there were no pathological features remaining in the problem.

**2.3.4 Manual variance reduction results**

The final optimized model achieved an increase in the FOM of a factor of ~ 6500 when compared to an analog calculation (no variance reduction used). Table I below outlines the approximate gains during the step-wise process of applying the variance reduction features discussed above. Note that the step increases in FOM relate to the progressive application of each variance reduction feature, and thus are not necessarily indicative of the increase in FOM associated with each individual variance reduction feature in the final optimized model.

**Table I. Manual Variance Reduction Results**

| Variance Reduction Feature | Step Increase in Figure of Merit (FOM) | Total Increase in Figure of Merit (FOM) |
|---|---|---|
| No variance reduction | N/A | 1 |
| Roulette / Splitting | 12 | 12 |
| Exponential Transform | 17 | 204 |
| Cell-Based Weight Windows + Source Energy Biasing | 8 | 1632 |
| Forced Collisions | 2 | 3264 |
| Mesh-Based Weight Windows + Implicit Capture Turned Off | 2 | 6500 |

For comparison purposes, an analog (unbiased) MCNP case was run for 800 minutes, yielding the following results: relative error = 0.8153, histories = 1,627,806. The relationship between number of histories, initial relative error and desired relative error is given by:

$$N_T = \left( \frac{RE_I \sqrt{N_I}}{RE_T} \right)^2 , \tag{7}$$

where

$N_T$ = target number of particles histories,
$N_I$ = initial number of particle histories,
$RE_I$ = initial relative error,
$RE_T$ = target relative error.

The relative error criterion for the problem was ≤ 1%. Using equation 7, it would require a run with approximately $1.082\times10^{10}$ particle histories to achieve a relative error of 0.01 with an unbiased calculation. Noting that it required 800 computer minutes (on a Pentium III, 1000 MHz processor) to achieve 1,627,806 histories, it would therefore require ~ $5.32\times10^{6}$ minutes to achieve the relative error criterion with the unbiased problem (or $8.86\times10^{4}$ hours, 3693 days, or 10.1 years).

### 2.3.5 Application of automated variance reduction

At the completion of the "manual" variance reduction effort, a separate effort was initiated to evaluate an automated variance reduction approach for this problem. In this section we briefly review the automated variance reduction methodology and discuss the application of a recently developed code [7], ADVANTG (Automated Deterministic VAriaNce reducTion Generator), that automates the generation of variance reduction parameters (source biasing and mesh-based weight window parameters) for MCNP based on 3-D deterministic adjoint functions.

It is well-known that the adjoint function (i.e., the solution to the adjoint form of the Boltzmann transport equation) has physical significance [8] as a measure of the importance of a particle to some objective function (e.g., the response of a detector) and that this physical interpretation makes the adjoint function well suited for biasing Monte Carlo calculations. Accordingly, recent trends in Monte Carlo code development have reflected a recognition of the benefits of using deterministic adjoint (importance) functions for Monte Carlo variance reduction.[9] Even though manually applied variance reduction by experienced Monte Carlo practitioners can yield increases in computational performance on the order of thousands for difficult problems (as shown in the previous section), automated variance reduction based on a deterministic importance function is expected to yield equal or superior computational performance and convergence reliability, while significantly reducing the requirements for user time and expertise. To evaluate these expectations, the ADVANTG code was applied to this problem and the results were compared to those achieved by manual application of variance reduction techniques, as discussed in the previous section.

#### 2.3.5.1 Automated variance reduction methodology

The variance reduction approach in ADVANTG is based on the CADIS (Consistent Adjoint Driven Importance Sampling) methodology [10], which provides consistent relationships for calculating source and transport biasing parameters based on importance sampling. The methodology is utilized to calculate space- and energy-dependent source biasing parameters and weight-window values. The biased source distribution, $\hat{q}(\vec{r},E)$, is given by the following relation

$$\hat{q}(\vec{r},E) = \frac{\phi^{+}(\vec{r},E)\,q(\vec{r},E)}{R} = \frac{\phi^{+}(\vec{r},E)\,q(\vec{r},E)}{\int_V \int_E q(\vec{r},E)\phi^{+}(\vec{r},E)\,drdE}, \tag{8}$$

where $\phi^{+}(\vec{r},E)$, $q(\vec{r},E)$, and $R$ are the scalar adjoint function, the unbiased source, and the detector response, respectively. The numerator is the detector response from space-energy element *($d\vec{r}$, dE),* and the denominator is the total detector response, $R$. Therefore, the ratio is a measure of the relative contribution from each space-energy element to the total detector response. Although the methodology is directly applicable to angular-dependent biasing by simply including angular dependency in the above equation, angular-dependency was not included.

For transport biasing, the weight window technique is employed. The weight-window technique provides a means for assigning space- and energy-dependent importances and applying geometric splitting/roulette and energy splitting/roulette, while at the same time controlling weight variations. The weight-window technique requires weight window lower bounds $w_\ell$, and the width of the window is controlled by the input parameter $c_u$, which is the ratio of upper and lower weight-window bounds ($c_u = \frac{w_u}{w_\ell}$). The space- and energy-dependent weight window lower bounds $w_\ell$ are given by [10]

$$w_\ell(\vec{r},E) = \frac{w}{\left(\frac{c_u+1}{2}\right)} = \frac{R}{\phi^+(\vec{r},E)\left(\frac{c_u+1}{2}\right)}, \tag{9}$$

where $w$ is particle weight. Because the calculational efficiency has been observed to be fairly insensitive to small deviations in the $c_u$ parameter, the MCNP default value of five was employed throughout this work. Note that because the source biasing parameters and weight window lower bounds are consistent, the source particles are started with statistical weights ($w(\vec{r},E) = \frac{q(\vec{r},E)}{\hat{q}(\vec{r},E)}$) that are within the weight windows, as desired. This is an important aspect of the CADIS methodology because it eliminates the incompatibility between source and transport biasing that has been problematic in other approaches due to poor calculational efficiency and/or false convergence.[11] For example, if the statistical weights of the source particles are not within the weight windows, the particles are immediately split or rouletted in an effort to bring their weights into the weight window. This results in unnecessary splitting/rouletting and a corresponding degradation in computational efficiency.

The CADIS methodology has been implemented in the ADVANTG code, which is a deterministic weight window generator (WWG) for MCNP that also generates consistent source biasing parameters. The input for using ADVANTG is very similar to that of the MCNP mesh-based stochastic WWG [2], and like the MCNP mesh-based WWG, ADVANTG outputs weight window values to a formatted file (i.e., the MCNP WWINP file) that may be read and utilized by the standard (unmodified) version of MCNP4C. However, unlike the stochastic MCNP WWG, ADVANTG also produces consistent source biasing parameters, does not require repeated applications to iteratively develop the weight window values, and does not require user modification of the weight window values. As indicated in the flowchart shown in Figure 3, ADVANTG automatically generates input files for material cross-section processing based on the GIP code [12] and 3-D (x-y-z or r-θ-z) discrete ordinates adjoint calculations with the TORT code.[12] Following the GIP and TORT calculations, ADVANTG (1) reads the standard TORT binary output file and the MCNP unbiased source, (2) calculates the source biasing and weight parameters, and (3) outputs the parameters for use with MCNP4C.

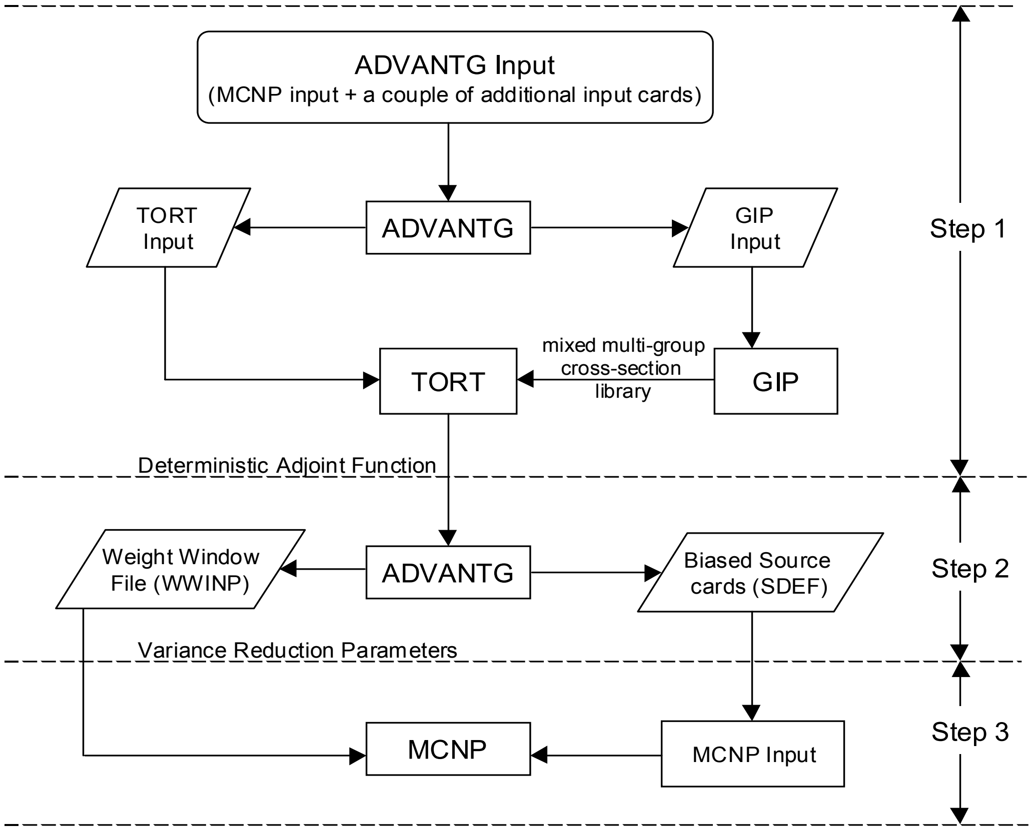


**Figure 3. Automated variance reduction process with ADVANTG**

### 2.3.5.2 Application of ADVANTG

Although many of the geometric structures in the problem are cylindrical, and the MCNP and ADVANTG codes support either cylindrical or rectangular mesh-based weight windows, previous experience [7,11] has shown that the implementation (in MCNP) of mesh-based weight windows in cylindrical geometry is less efficient than for rectangular geometry. Hence, rectangular mesh-based weight windows were used here. The weight window mesh boundaries were selected to be consistent with the problem material boundaries, and the resolution of meshes between material boundaries was varied to evaluate the impact on efficiency. Table II summarizes the mesh characteristics and computer time required for the discrete ordinates (TORT) calculations for selected cases. Note that the CPU times listed for the TORT calculations are considerably less than the computational effort associated with manually developing the variance reduction parameters, as discussed in the previous section. The spatial mesh distribution and material assignments used by TORT are shown in Figure 4. The TORT calculations used the 47-group SAILOR96 library [13] and an $S_8$ quadrature. Although it would seem preferable to use a multi-group cross-section library with a fewer number of energy groups (to minimize the time required for the TORT calculations), subsequent studies with the 22-group CASK cross-section library [14] did not show improved overall efficiency. As the reaction of interest in the excore detector is $^{10}B$ absorption, the cross section for this reaction is used as the source spectrum in the adjoint TORT calculation.

**Table II. Automated Variance Reduction Results**

| Case | TORT | | | MCNP | |
|---|---|---|---|---|---|
| | x-y-z mesh | Total number of meshes | CPU time (minutes) | FOM | Speedup (FOM/ unbiased FOM) |
| Unbiased | n/a | n/a | n/a | 0.0018 | 1 |
| 1 | 20×20×09 | 3,600 | 4.2 | 52 | 28,889 |
| 2 | 26×26×13 | 8,788 | 9.9 | 143 | 79,444 |
| 3 | 32×32×13 | 13,312 | 21.6 | 156 | 86,667 |
| 4 | 36×36×13 | 16,848 | 21.1 | 105 | 58,333 |
| 5 | 40×40×21 | 33,600 | 44.4 | 82 | 45,556 |
| 6 | 50×50×21 | 52,500 | 90.5 | 46 | 25,556 |

Because the ADVANTG code and associated methodology are still *relatively* new, sensitivity studies are routinely performed by the code author to assess the sensitivity of the major input parameters (e.g., spatial mesh, multi-group library, and quadrature order) and contribute to the development of guidance for future users. The capability to generate two-dimensional plots of the spatial mesh and material assignments for any (and all) axial planes (see Figure 4) has proven quite useful for these studies. Although not required, it is also instructive to visualize the adjoint function being utilized. ADVANTG optionally prints the adjoint data in a format suitable for visualization with TecPlot®. If, for example, the deterministic results included problems with "ray-effects," which is not the case for this problem, it would be evident by visualization of the results. Figure 5 provides plots of the Case 3 adjoint function for selected energy groups. These figures illustrate the expected behavior; dramatic decline in particle importance (especially for the lower energies) as one moves from the detector toward the core center. To facilitate comparisons, the adjoint function plots in Figure 5 use a consistent scale. Note, however, that refinement of the scale reveals the importance of scattering from the periphery of the reactor pressure vessel and inner region of the concrete shield. Finally, if desired, one can examine the generated weight window values overlaid on the MCNP geometry with the superimposed mesh plotting capability in MCNP4C.

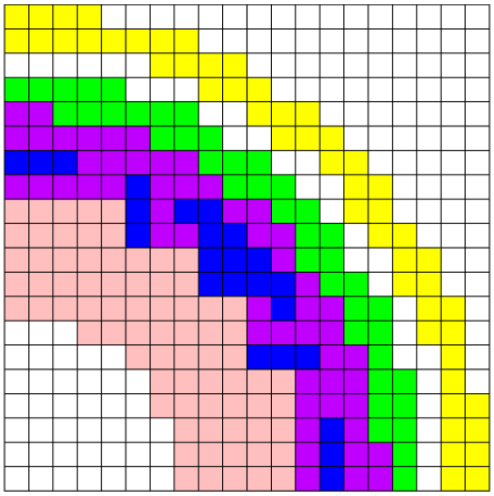

Case 1, 20×20×09

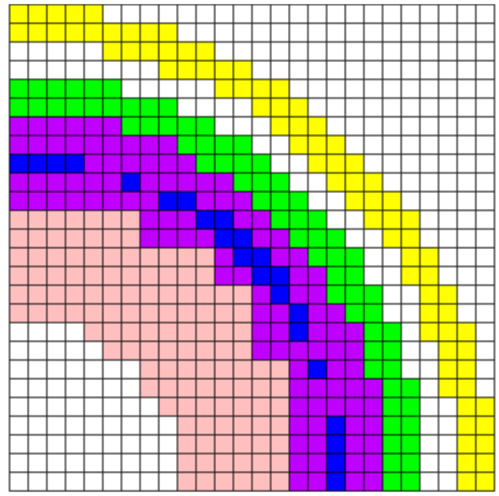

Case 2, 26×26×13

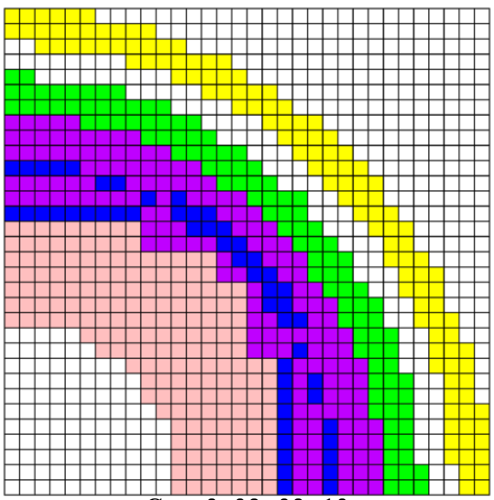

Case 3, 32×32×13

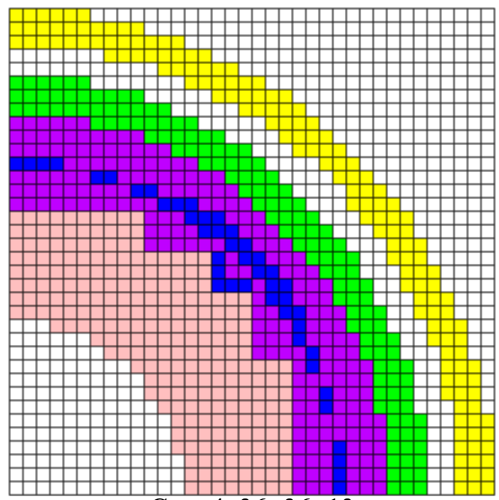

Case 4, 36×36×13

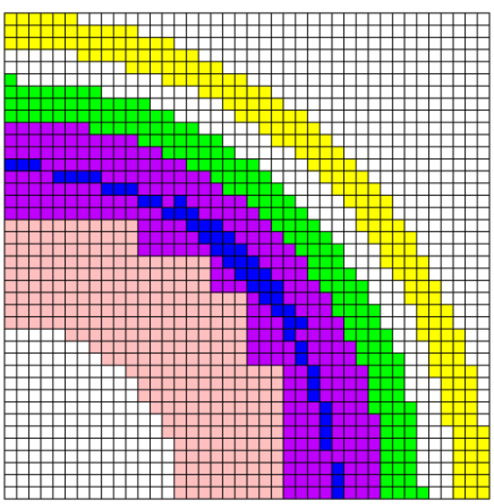

Case 5, 40×40×21

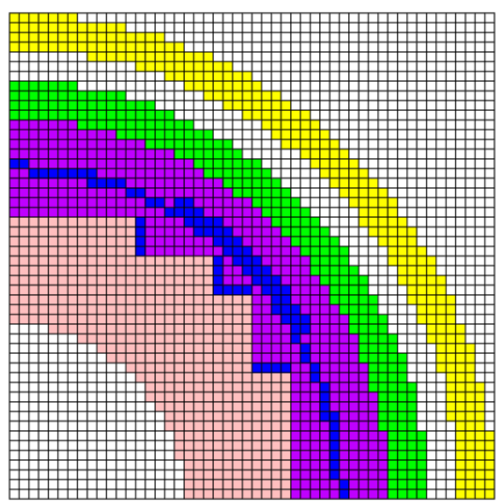

Case 6, 50×50×21

**Figure 4. Spatial mesh and material assignments for the various cases**

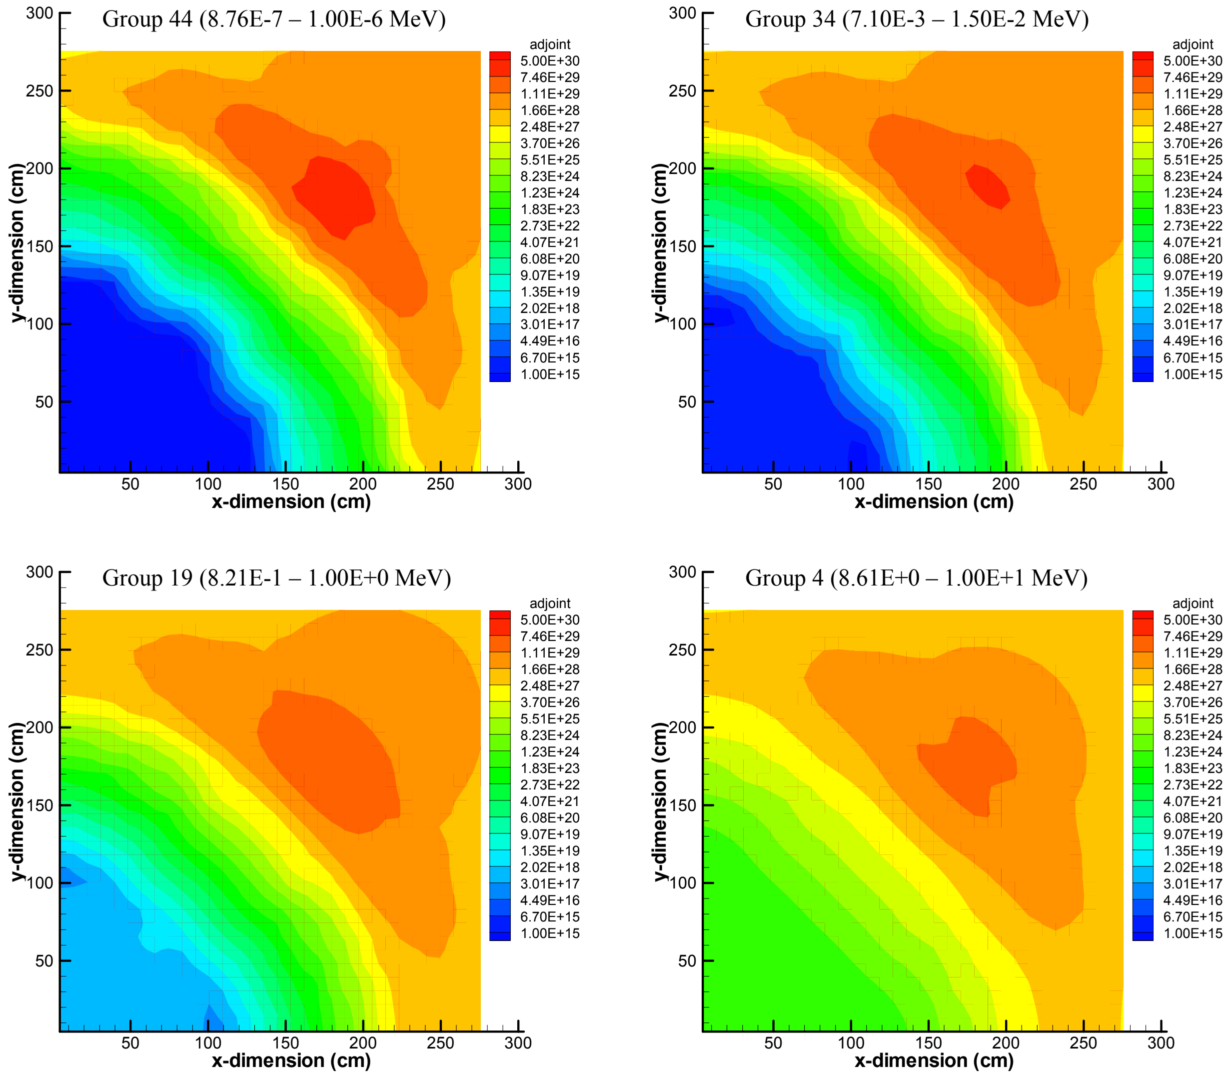


**Figure 5. Adjoint functions from Case 3 for selected energy groups**

For the reference model, a source spatial probability was defined for each quadrant of each assembly and the source energy spectrum was defined in terms of 40 energy groups. As discussed previously, the source in the inner assemblies does not contribute significantly to the excore detectors, and thus was neglected. With this source definition, ADVANTG was used to calculate a biased probability for each quadrant of each assembly and corresponding spatially-dependent biased energy spectra. The original (unbiased) and biased spatial source probabilities for case 5 are compared in Figure 6. As intuition would predict, the source is heavily biased toward the core periphery. To illustrate the importance of spatially-dependent biased energy spectra, the original (unbiased) spectrum is compared in Figure 7 to biased spectra from an assembly quadrant nearest to the detector and an assembly quadrant farthest from the detector. The biased spectra show the importance of the higher energy neutrons in the source regions and the insignificance of the lower energy neutrons, particularly as the thickness of fuel region between the source and detector locations increases.

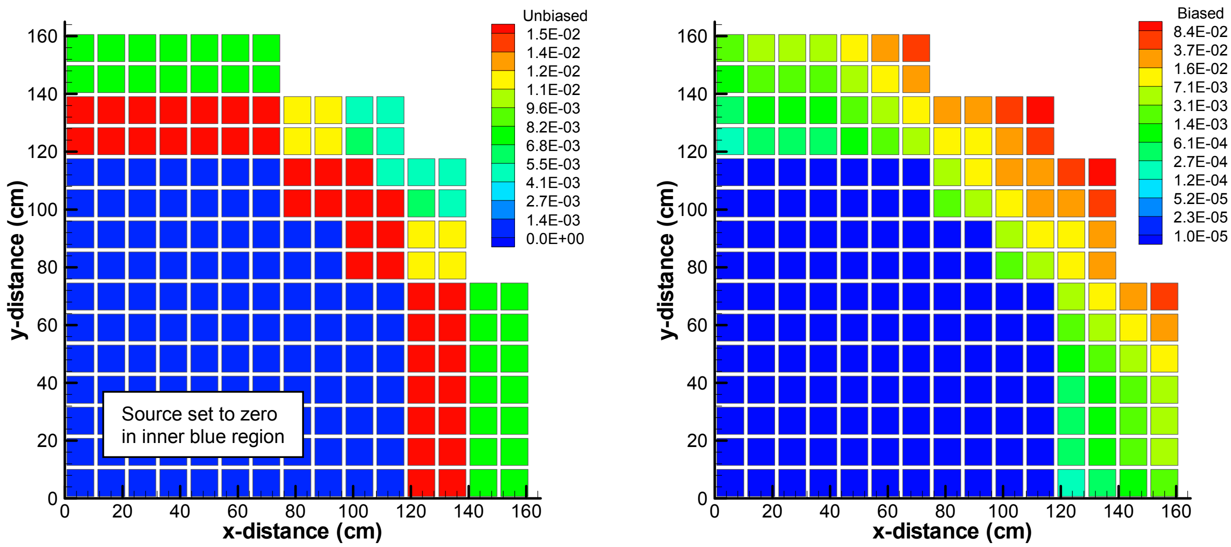


**Figure 6. Comparison of original (unbiased) and biased spatial source probabilities**

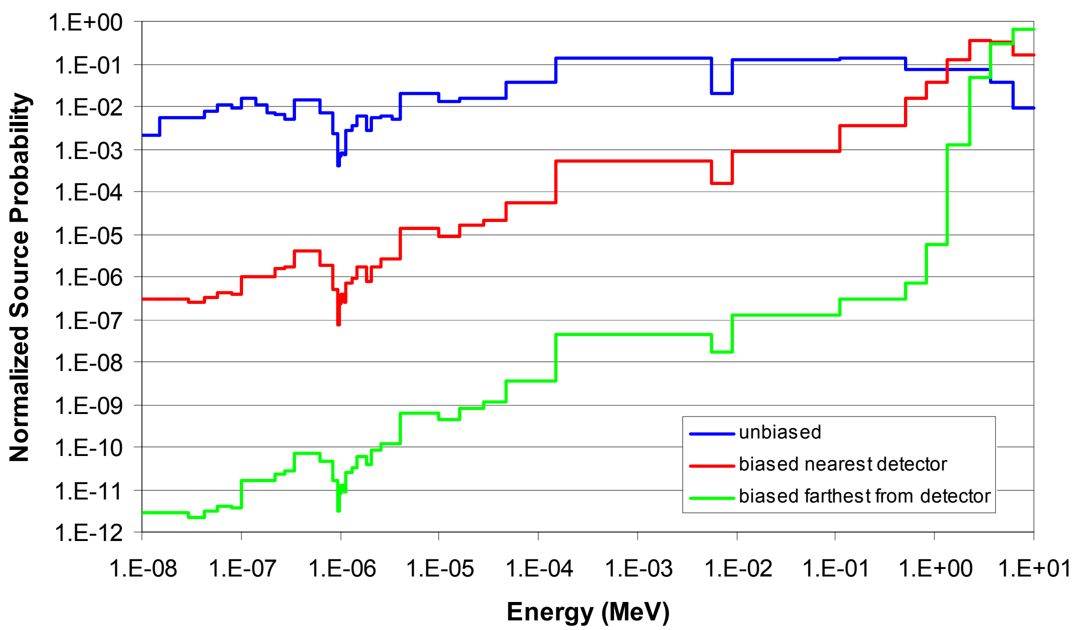


**Figure 7. Comparison of original (unbiased) and biased energy spectra**

Table II includes a summary of the FOM values for the selected cases and shows FOM speedups in the range of 25,000 – 87,000, with respect to an unbiased case. Additionally, comparison of the FOM values in Tables I and II indicates speedups in the range of ~ 4 – 13, with respect to the best manually-optimized case. In all cases (i.e., manually optimized and ADVANTG cases), good statistical convergence behavior was achieved, and thus the comparisons made herein between manual and automated variance reduction are valid and meaningful. While the improvement in computational efficiency with ADVANTG is significant, the overall efficiency is substantially greater, as compared to the manually-developed variance reduction, when one considers the user-time required by an experienced Monte Carlo practitioner to develop the variance reduction parameters, as well as the associated computer time required to develop the parameters. Note, the variance reduction parameters produced by ADVANTG were used "as is" (i.e., the weight window file and source biasing parameters were not manually modified in any way).

Consistent with previous findings [7,9,15], the calculational efficiency for this problem was not overly sensitive to the spatial-mesh resolution of the adjoint function. To illustrate this point, Figure 8 plots the speedup (FOM ratio) as a function of the total number of meshes (used for both the TORT calculation and the weight windows). This behavior is considered to be desirable, as different users will inevitably take different approaches toward defining mesh resolution. In contrast, the mesh-based WWG in MCNP has been found to be fairly sensitive to mesh resolution, and the temptation to define meshes that are too small to be properly sampled can lead to inadequate estimates of the weight windows. The results in Figure 8 also show that the speedups achieved with the track-length estimator and point detector are quite comparable, providing an indication that the methodology is effective for both estimators. Regarding the mesh distribution, it is important to capture the bulk characteristics of the problem geometry (e.g., material locations and thicknesses) in order to capture the physics characteristics of the problem. Failure to do so can manifest itself in poor computational efficiency and/or convergence behavior. However, because the process is automated, creating and utilizing different mesh distributions is a simple matter.

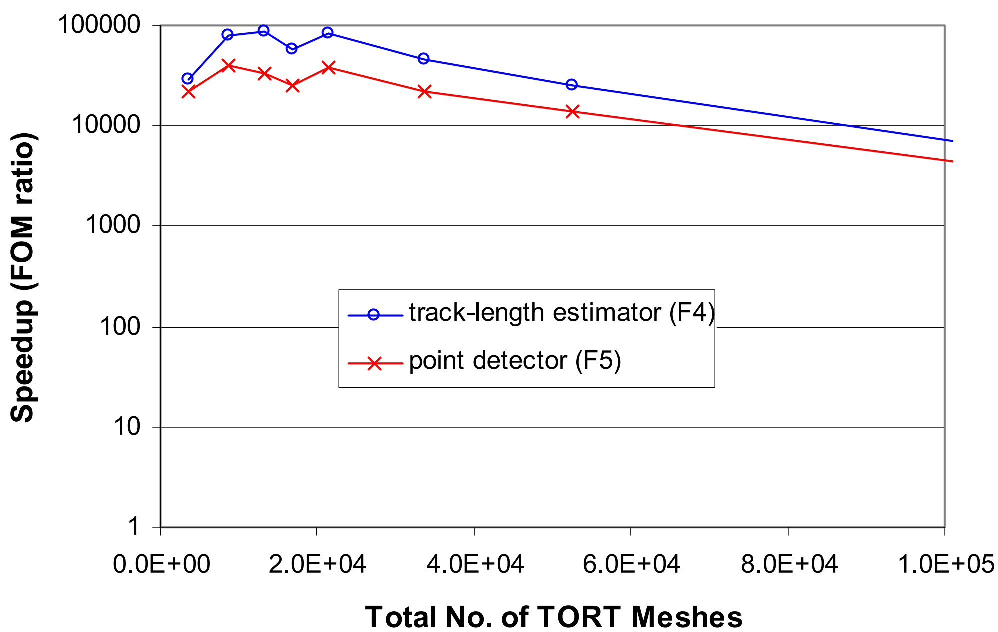


**Figure 8. Speedup as a function of mesh refinement**

## 3. CONCLUSIONS

Detailed Monte Carlo analysis of the excore neutron detector response to changing conditions in the core resulting from a transient requires effective use of variance reduction techniques to accelerate and ensure proper problem convergence. Consequently, a considerable expenditure of time and effort was made to effectively utilize the variance reduction techniques available in MCNP by manually developing the required parameters and iteratively applying the cell-based and mesh-based weight-window generators. Although this process was user-time intensive, a computational speedup of 6500, with stable convergence characteristics, was ultimately achieved. Unfortunately, neither the total CPU time nor the total user-time for this process was recorded, and thus it is not possible to accurately quantify this effort. However, due to the many sensitivity calculations required to completely study the phenomenon of interest and issues related to calculational confidence and legacy, the process represented a worthwhile effort to the overall analysis. Recognizing the importance of verifying problem convergence for this safety-related analysis and the potential for further increase in computational efficiency, a separate effort was undertaken to utilize the recently developed ADVANTG code for automated variance reduction based on 3-D deterministic adjoint functions. Application of the automated variance reduction capability (1) yielded stable statistical convergence behavior, (2) confirmed the problem convergence achieved with the manually-developed variance reduction parameters, and (3) resulted in a maximum computational speedup of ~ 87,000, with respect to an unbiased case, and ~13, with respect to the best manually-optimized case.

To assess the effects of decreased moderator temperature on excore detector response, calculations were performed using the variance reduction parameters generated via ADVANTG with variations in the moderator density in the fuel, flow channel, and downcomer regions. The tally results from the various moderator density statepoints were utilized to generate an attenuation factor, in percent power indication per degree Fahrenheit, for subsequent use in time-dependent multi-node transient analysis calculations.

The MCNP result for each moderator density statepoint was converged to less than or equal to 1% relative error. Finally, the MCNP results were used to perform a curve fit of attenuation functions versus moderator density with the TableCurve 2D v5.00® software. The resultant curve fit had a correlation coefficient of approximately 0.9995. Therefore, good statistical performance was achieved and the results were deemed to be appropriate for application in nuclear installation safety analyses.

## REFERENCES

1. Nuclear Regulatory Commission Information Notice 96-41, “Effects of a Decrease in Feedwater Temperature on Nuclear Instrumentation.”
2. Judith F. Briesmeister, Ed., “MCNP™ - A General Purpose Monte Carlo N-Particle Transport Code,” Version 4C, LA-13709-M, December 2000.
3. Electric Power Research Center Topical Report, “Validation & Verification of the Dynamic Reactivity Measurement of Rod Worth Technique,” North Carolina State University, College of Engineering, 2001.

4. J.C. Wagner, “Monte Carlo Transport Calculations and Analysis for Reactor Pressure Vessel Neutron Fluence,” *M.S. Thesis* The Pennsylvania State University, The Graduate School College of Engineering, December 1994.
5. “CASMO-4: A Fuel Assembly Burnup Program,” STUDSVIK/SOA-95/1.
6. S. P. Pederson, R. A. Forster, T. E. Booth, “Confidence Interval Procedures for Monte Carlo Transport Simulations,” *Nucl. Sci. Eng.*, **127**, 54-77 (1997).
7. J.C. Wagner, “An Automated Deterministic Variance Reduction Generator for Monte Carlo Shielding Applications,” wagner1m.pdf, Paper Log 131 in Proc. of American Nuclear Society 12th Biennial RPSD Topical Meeting, April 12-14, 2002, Santa Fe, NM.
8. G.I. Bell and S. Glasstone, Nuclear Reactor Theory, Van Nostrand and Reinhold, New York, 1970.
9. A. Haghighat and J.C. Wagner, “Monte Carlo Variance Reduction with Deterministic Importance Functions,” *Progress in Nuclear Energy*, **42(1)**, 25-53 (2003).
10. J.C. Wagner and A. Haghighat, “Automated Variance Reduction of Monte Carlo Shielding Calculations Using the Discrete Ordinates Adjoint Function,” *Nucl. Sci. Eng.*, **128**, 186 (1998).
11. C.N. Culbertson and J.S. Hendricks, “An Assessment of MCNP Weight Windows,” Los Alamos National Laboratory, LA-13668 (1999).
12. *DOORS 3.2: One, Two- and Three-Dimensional Discrete Ordinates Neutron/Photon Code System*, Available from Radiation Safety Information Computational Center at Oak Ridge National Laboratory as CCC-650.
13. *BUGLE-96: Coupled 47 Neutron, 20 Gamma-Ray Group Cross Section Library Derived from ENDF/B-VI for LWR Shielding and Pressure Vessel Dosimetry Applications*, Available from Radiation Safety Information Computational Center at Oak Ridge National Laboratory as DLC-185
14. *CASK-81: 22 Neutron, 18 Gamma-Ray Group, $P_3$ Cross Sections for Shipping Cask Analysis*, Available from Radiation Safety Information Computational Center at Oak Ridge National Laboratory as DLC-23.
15. J.C. Wagner, “Acceleration of Monte Carlo Shielding Calculations with an Automated Variance Reduction Technique and Parallel Processing,” *Ph.D. Thesis*, The Pennsylvania State University, The Graduate School College of Engineering, December 1997.