International Conference on Mathematics, Computational Methods & Reactor Physics (M&C 2009)
Saratoga Springs, New York, May 3-7, 2009, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2009)

A WEIGHT-WINDOW GENERATOR FOR ELECTRON-PHOTON TRANSPORT IN THE INTEGRATED TIGER SERIES CODES

Brian C. Franke, Martin J. Crawford, and Ronald P. Kensek
Sandia National Laboratories
P.O. Box 5800, MS 1179
Albuquerque, NM 87185
bcfrank@sandia.gov; mjcrawf@sandia.gov; rpkense@sandia.gov

Jason A. Kraftcheck
Department of Engineering Physics
University of Wisconsin
1500 Engineering Drive
Madison, WI 53706
kraftche@cae.wisc.edu

ABSTRACT

While the weight-windows method for biasing Monte Carlo calculations has existed for more than 25 years, only recently has it been applied to coupled electron-photon calculations. We present a weight-window generator intended for such problems in three spatial dimensions and designed to automatically arrive at an effective biasing map. Automation is achieved through importance tallies adaptively refined (or coarsened) across multiple adjoint Monte Carlo tallies.

Key Words: Weight windows, electron-photon, radiation transport

1. INTRODUCTION

To facilitate biasing of electron-photon transport in the Integrated Tiger Series (ITS) codes [1], we have implemented a weight-window generator based on the adjoint Monte Carlo transport capability. Among our goals were: (1) to use the same geometry model in the generator as in the biased Monte Carlo calculation, (2) to avoid the need for geometry-model modification to facilitate biasing, (3) to avoid the need for user intervention in setting weight windows, and (4) to bias any electron-photon problem, including those with rapidly varying electron importance near surfaces. To accomplish these, we use: (1) adjoint Monte Carlo calculations, (2) an adjoint-flux/weight-windows tally structure distinct from the material geometry, and (3) an adaptive adjoint-flux tally structure. Tallies are performed on a tree data structure that has been implemented in the Mesh-Oriented datABase (MOAB) [2]. While the adaptivity could encompass all seven transport dimensions, the current work adapts only three spatial dimensions.

Weight-windows generation and biasing have existed for some time [3, 4], but have only recently been applied to coupled electron-photon problems [5, 6]. Dionne [5] explored use of deterministic adjoint methods to bias three-dimensional forward Monte Carlo calculations. Ueki [6] focused on using forward calculations to determine particle importance. Ueki demonstrated

in one-dimensional calculations the effectiveness of using weight-windows biasing for the type of electron-emission calculations that we will be examining in this paper. Here we attempt to address the issue of adequately representing the weight-windows biasing map in a multi-dimensional problem.

Adaptive Monte Carlo tallies have been described by Booth [7, 8] and Lai and Spanier [9]. Our method is limited to only refining the phase-space division, like Ref. [7]. But like Refs. [8] and [9], our approach directly iteratively refines the importance map of a problem. Those works demonstrated that adaptivity of importance tallies could be used to converge toward zero-variance Monte Carlo calculations. Those approaches estimated importance from forward Monte Carlo calculations and used the resulting importance map to bias the next forward calculation, iteratively improving the importance map. Our approach would appear to be advantageous in the early portion of the importance map evolution when detector tallies are very difficult to achieve. We also expect our approach to be more reliable, if more expensive, because using unbiased adjoint calculations to build the importance map is unlikely to over-bias by omitting important low-probability physics paths. (Perhaps these concerns are alleviated when an initial deterministic global importance-map evaluation is used).

In Section 2, we discuss our spatial-adaptivity algorithm used to refine the adjoint-flux tally distribution. In Section 3, we discuss our implementation of the weight-windows biasing method. In Section 4, we present data on the efficiency of our method for a test problem. We offer some concluding remarks in Section 5.

2. SPATIAL ADAPTIVITY OF ADJOINT-FLUX TALLIES

We use “volume-source” tallies in adjoint calculations. These are equivalent to track-length adjoint-flux tallies but normalized to yield detector-response values for volume-distributed sources. The tally routines were modified to allow particle tracks to be apportioned into a tree data structure.

The spatial distribution of the tallies is stored in a binary space-partitioning tree data structure. Because it is a binary tree, each node of the tree can only be divided into two smaller nodes. For simplicity, we have chosen to constrain the tree cut planes to be axis aligned, which results in a “kd-tree.” We require that the root node of the tree is an axis-aligned bounding box, so every node in the tree will also be an axis-aligned bounding box. This tree has been implemented in MOAB [2]. Tallies are only made on the leaves of the tree, where a “leaf” is any node on the tree that is not further divided. Since this is a space-partitioning tree, each point in space within the root node is contained by one and only one tree leaf.

We begin our suite of adjoint calculations with an initial tree that partitions the space of the problem. After each adjoint calculation, neighboring tree leaves are compared to determine whether the ratio of fluxes between them meets a test criterion (e.g., is the ratio less than a factor of three?). If the criterion is not met, then the leaves are split. If the boundaries of the shared faces of the neighboring leaves do not coincide in extent, then cuts will first be made to make the faces coincident. Otherwise, the leaves are split in half in the axial direction corresponding to the shared face. This can be thought of as an adaptive refinement technique driven by error

estimation. The error estimator is an infinity-norm assessed only on leaf boundaries using only leaf-average values. In these terms, there are clearly defects to the technique. The leaf-average value may mask important information about solution variation within the leaf, which cannot be assessed. Assessing the error only at leaf boundaries offers limited information. Clearly, we are not assured of achieving a solution within a factor of three of the true solution everywhere within the problem. However, this also suggests that, if necessary, a higher-order error estimator might be able to more rigorously enforce the desired infinity-norm error level.

3. WEIGHT-WINDOWS IMPLEMENTATION

Our implementation of weight windows follows common descriptions of the technique [10], with a few minor differences. The center weight of the window is stored on each leaf of the kd-tree on which the adjoint particle flux was calculated. The weight is the inverse of the adjoint flux, with all weights normalized by the same value. Ideally, the normalization factor will avoid splitting and minimize Russian roulette of source particles. (We have not implemented source biasing, which can alleviate these concerns to some degree.) The extent of the window is a user-prescribed setting for the ratio of the top to the bottom of the window, logarithmically centered about each weight. A particle with weight greater than the top of the weight window is split into an integer number of particles sufficient to bring the weight below the center weight. A particle with weight lower than the bottom of the window is Russian rouletted to bring the weight of a surviving particle to the center weight.

Weight-windows biasing is applied to photons at each interaction point and each time they cross a geometry boundary. Weight-windows biasing is applied to electrons at the beginning of each ITS step, which includes each boundary crossing into a unique material. For each application of the weight-windows settings, the kd-tree is interrogated to determine the weight window for the current particle phase-space.

As stated above, the center weight of a window is determined as the inverse of the adjoint flux. However, two special cases must be dealt with: adjoint-flux values with poor statistics and regions of phase space that obtained no adjoint-flux tallies. A lack of tallies in the unbiased adjoint calculation can occur for two reasons: a low adjoint-flux level and/or a small-volume tally region. Our approach is to determine the largest value of the product of flux and volume (equivalent to total track-length tallied in the region) with relative statistical uncertainty greater than 20%. We reset all smaller flux-times-volume values to this level. This is a conservative approach that systematically over-estimates the importance of particles in regions of phase space that lack adequate information.

4. RESULTS

To test and demonstrate our method, we use a complicated real problem that conveniently can be analyzed with the current constraints of our algorithm. That is, most of the geometry boundaries are axis-aligned, so many of the particle-importance contours will be nearly axis-aligned and can be efficiently represented using a tree structure with axis-aligned cutting planes. This problem involves modeling a validation experiment performed on Sandia's Saturn electron-accelerator facility. The geometry of the Saturn bremsstrahlung converter and cavity device is shown in Figure 1.

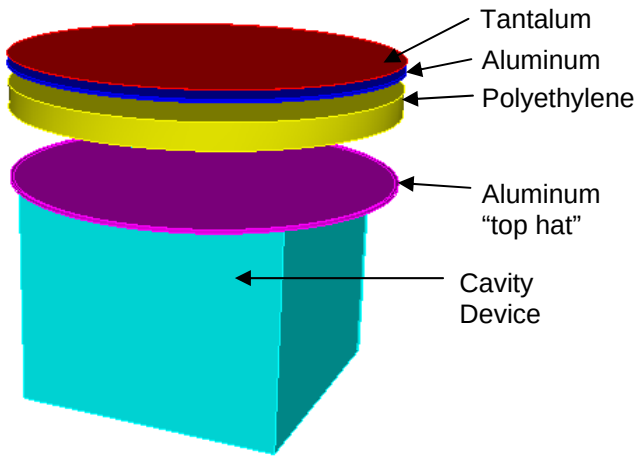

Figure 1. Geometry model with converter and cavity device.

The Saturn "converter" consists of tantalum, aluminum, and polyethylene layers to convert source electrons into bremsstrahlung photons, which impinge from above onto the cavity device. A source of electrons is modeled in three concentric rings. The source electrons are 1.6 MeV and normally incident on the tantalum. An aluminum "top hat" sits over the cavity device. The cavity device is simplified as a square shell of brass wrapping around the four sides of a hollow graphite box. A portion of the inside surface of the top of the graphite box is covered in a thin layer of gold. The goal of the ITS calculation is to calculate the distribution of electrons emitted into the cavity, differential in the five dimensions of surface area, angle, and energy. Most of the electrons entering the cavity are photo-electrons from the gold layer. However, electrons also enter the cavity due to photon interactions in the graphite. Plus, the graphite is sufficiently thin to allow electrons from the aluminum top hat and brass shell to transport into the cavity.

We begin by performing a calculation of the adjoint particle flux due to the detector response of electrons emitted into the cavity. Figure 2 shows the evolution of the importance map for the highest-energy electron flux. Despite our objective to avoid the need for user intervention in constructing the weight-windows settings, there are several "knobs" in the algorithm during the adjoint-flux adaptivity. First, the user must decide on an initial tree to begin the iteration process. In Figure 2, the first tree used a uniformly refined tree. A better choice is to begin with a tree that incorporates information about material boundaries, but this may be more labor intensive for the user. For each iteration of the tree, the user must choose the adjoint-flux tallies. Since our algorithm only partitions spatial dimensions, the user must choose the energy and angular binning of the flux tallies. In the example problems, five energy bins and one angle bin

were used for both electrons and photons. Also for each iteration of the tree, the user must select the number of histories to be simulated. In the third flux map shown in Figure 2, corresponding to the 20th iteration on the tree, the poorer statistical results are shown by the dark blue cells indicating zero flux. Another feature of the results worth noting is that the importance inside the cavity has been over-estimated compared to a true adjoint calculation of electron emission into the cavity. The adjoint source was intentionally modified to accomplish this so as to prevent the tree-refinement algorithm from refining the importance map within the cavity, where the forward calculation will not track electrons.

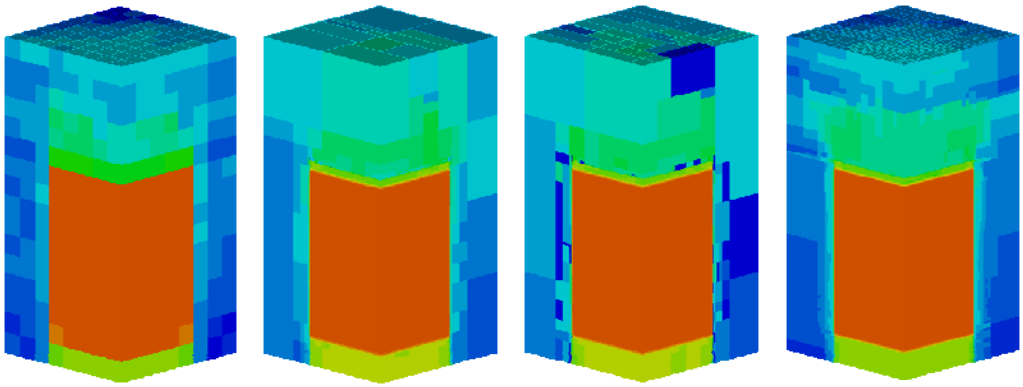

Figure 2. Adjoint-flux distribution for the highest-energy electrons after 1, 10, 20, and 30 Monte Carlo calculations.

In Figure 3, line-outs of the converged electron and photon adjoint fluxes (i.e., unnormalized importances) are shown on log-log plots as a function of distance upward from the gold-cavity interface. These results do not include the "fix up" for poor statistical results, which are visible in the low-energy electron results. Indeed, the lowest-energy results for electrons and photons received no tallies beyond the gold material of the first 12.5 μm. However, the importance map has been resolved over almost six orders of magnitude. The map has resolved both the rapid variation of importance of the lowest-energy electrons over almost six orders of magnitude in importance across less than 1 μm and the slower variation of the highest-energy electrons over four orders-of-magnitude in importance across 10 cm. We observe that above the 10^{-9} flux level in the line-out data, the tree has spatially resolved factors of three within each energy level (with one exception in the 10-50 keV photon flux in the converter region). However, deficiencies in the importance map remain. From the data in Figure 3, we can observe that the coarse energy binning causes significant differences in importance from one energy bin to the next at some locations. This data does not indicate how much variation in importance might exist as a function of angle, which has been suppressed into a single angle bin. Also not shown in Figure 3, the tree has not resolved importance gradients where contours are not axis aligned, but this would require a very large number of leaves using axis-aligned cutting planes.

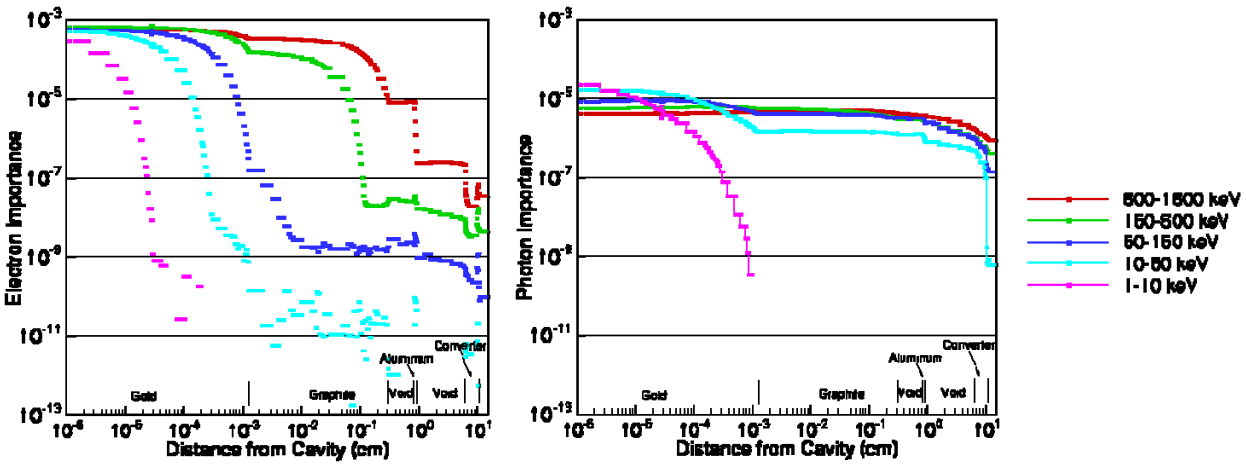

Figure 3. Line-outs of the adjoint-flux distributions for electrons (left) and photons (right).

In Figure 4, we show the "figure of merit ratio" as a function of the electron-emission tally. The figure of merit is defined as

$$FOM = \frac{1}{\sigma^2 t}, \tag{1}$$

where σ^2 is the variance of the tally and t is the computational time. The figure-of-merit ratios are given relative to a manually biased calculation. The tallies were differential in 7 surfaces, 128 angle bins, and 40 energy bins. The manual biasing used three techniques: (1) scaling the cross section to increase bremsstrahlung production in the converter, (2) raising electron-cutoff energies in some regions, and (3) trapping electrons to terminate electron transport when appropriate. (It may be worth mentioning that the manual biasing originally used with this problem included forcing of interactions near the cavity boundaries. Anomalies in the emission tallies indicated that the calculation was being over-biased, a common danger with manual biasing techniques that can lead to false confidence in erroneous results.) The "automatic" results used the algorithm previously described in this paper. The "assisted" results differ from the "automatic" results only in that they began with a tree that had been divided along material interfaces near the surface of the cavity. From Figure 4, we conclude that the weight-windows biasing scheme has achieved an efficiency improvement of approximately an order-of-magnitude over a manually biased calculation for this test problem. Some improvement in efficiency was achieved when the adaptive tree was initialized with key information about geometry boundaries. Based on calculations using coarse binning structures (five energy bins and a single angular bin), the manually biased calculation was estimated to be about four orders-of-magnitude more efficient than a completely unbiased calculation.

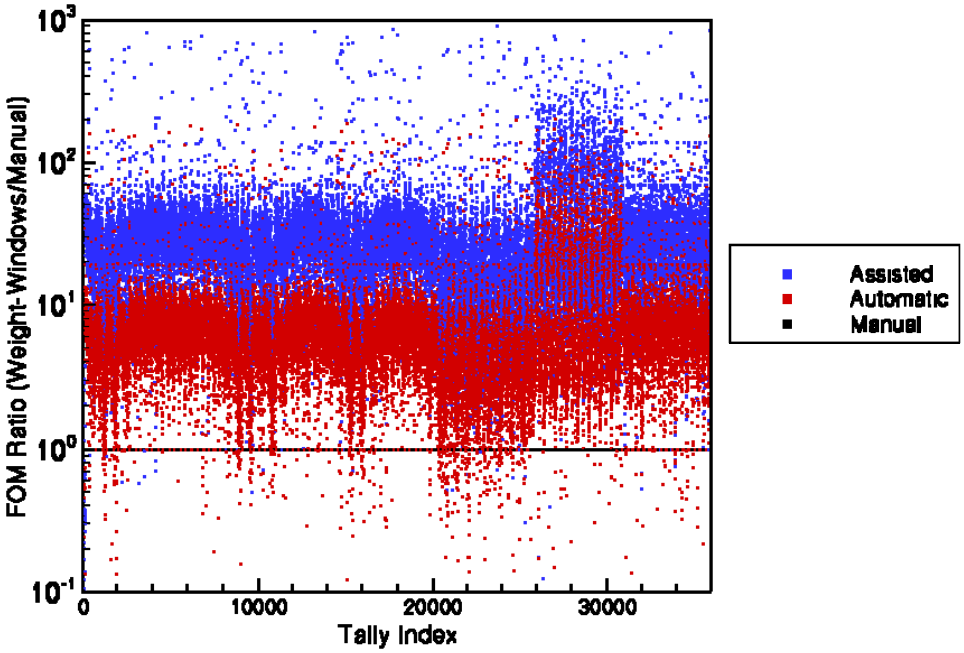

Figure 4. Figure-of-merit ratios for electron-emission tallies in three biased Monte Carlo calculations.

5. CONCLUSIONS

We have demonstrated the capability of our adaptive importance maps to achieve a sufficiently accurate importance map to result in an effective automated biasing method. For a realistic test problem, an efficiency improvement of approximately an order-of-magnitude was achieved over a manually biased calculation. Based on calculations using very coarse binning structures, the manually biased calculation was estimated to be about four orders-of-magnitude more efficient than a completely unbiased calculation. However, memory limitations encountered using the axis-aligned cutting planes of a kd-tree for particle importance maps severely limit the complexity and/or efficiency of problems that can be analyzed. To analyze more practical problems, a more general spatial-partitioning approach must be implemented.

ACKNOWLEDGMENTS

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000. The authors wish to acknowledge Tim Tautges for many helpful discussions.

REFERENCES

1. B. C. Franke, R. P. Kensek, T. W. Laub, "ITS Version 6: The Integrated TIGER Series of Coupled Electron/Photon Monte Carlo Transport Codes," Sandia National Laboratories, Technical Report SAND2008-3331 (2008).
2. T. J. Tautges, R. Meyers, K. Merkley, C. Stimpson, C. Ernst, "MOAB: A Mesh-Oriented datABase," Sandia National Laboratories, Technical Report SAND2004-1592 (2004).

3. J. S. Hendricks, "A Code-Generated Monte Carlo Importance Function," *Trans. Am. Nucl. Soc.*, **41**, pp. 307-308 (1982).
4. T. E. Booth, "Automatic Importance Estimation in Forward Monte Carlo Calculations," *Trans. Am. Nucl. Soc.*, **41**, pp. 308-309 (1982).
5. B. Dionne, "Automated Variance Reduction Technique for 3-D Monte Carlo Coupled Electron-Photon-Positron Simulation Using Deterministic Importance Functions," University of Florida, Ph.D. Dissertation (2007).
6. T. Ueki, "Monte Carlo Variance Reduction for Energy Profile Calculation at the Electron Emission Due to Photon Radiation," *Nucl. Sci. Eng.*, **157**, 2 pp. 119-131 (2007).
7. T. E. Booth, "Intelligent Monte Carlo Phase-Space Division and Importance Estimation," *Trans. Am. Nucl. Soc.*, **60**, pp. 358-360 (1989).
8. T. E. Booth, "Adaptive Importance Sampling with a Rapidly Varying Importance Function," *Nucl. Sci. Eng.*, **136**, pp. 399-408 (2000).
9. Y. Lai, J. Spanier, "Adaptive Importance Sampling Algorithms for Transport Problems," *Monte Carlo and Quasi-Monte Carlo Methods 1998, Lecture Notes in Computational Science and Engineering*, H. Niederreiter and J. Spanier, Eds., Springer-Verlag, New York, 273 (1999).
10. F. B. Brown, "Fundamentals of Monte Carlo Particle Transport," Los Alamos National Laboratory, Technical Report LA-UR05-4983 (2005).