Joint International Topical Meeting on Mathematics & Computation and Supercomputing in Nuclear Applications (M&C + SNA 2007)
Monterey, California, April 15-19, 2007, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2007)

TRANSITION AND COLLISION BIASING IN A VARIANCE REDUCTION SCHEME FOR CRITICALITY CALCULATIONS

Stavros Christoforou and J. Eduard Hoogenboom
Delft University of Technology, Department of Applied Sciences
Mekelweg 15, 2629 JB, Delft, The Netherlands
s.christoforou@tudelft.nl; j.e.hoogenboom@tudelft.nl

ABSTRACT

We use a simplified transport model in order to investigate the effect of transition and collision biasing applied to a Monte Carlo simulation. This is done by using the adjoint forms of the emission and collision densities, obtained from a deterministic calculation, in order to bias the transition and collision kernels, according to the zero-variance scheme we apply. The scheme is tested in two cases of 1D geometry, one with a simple 2-direction transport model and the other in a 3D transport model. Transition biasing is shown to decrease the variance significantly, but at a large computational cost, which reduces its effectiveness in improving efficiency of a Monte Carlo run. Collision biasing shows a significant increase in the efficiency of the simulation in both cases tested, and combined with the fact that it can be easily applied to a large range of geometries with little to no computational cost, makes it a much more attractive option for inclusion in production codes.

Key Words: Monte Carlo, Adjoint biasing, Variance reduction

1. INTRODUCTION

The statistical nature of Monte Carlo implies that there will always be an error associated with a calculation. It is therefore one of the main targets of the Monte Carlo approach to decrease this error as much as possible, without heavily taxing the computational time of the calculation. Towards this target, most Monte Carlo simulations today include some form of *biasing*. It is almost always beneficial to use biasing in shielding calculations, however, it is not straightforward how to reduce the variance in a criticality calculation, where the reactor serves both as the source and ”detector” of particles simulated.

The concept of biasing regards the process of encouraging particle histories which are more likely to contribute to the required response, while the particle histories less likely to contribute are given less importance. This is done by altering the probability density function (hereafter noted as *pdf*) that describes the process investigated. In order for the result to be valid, we compensate for this change by altering the statistical weight of each particle. It has been proven that in order to appropriately bias a Monte Carlo simulation, the adjoint equations should be used as weighting, or importance functions[1]. By importance $I(P)$ of a particle we hereafter mean the expected or probable contribution of the particle at point P to the detector response[2].

1.1. Formulation

In order to bias our transport kernel, we first need to define quantities that we can sample. Those are the collision density $\psi(P)$ of particles going into a collision at point P and the emission density $\chi(P)$ of particles coming out of a collision at P.

A particle will start a flight path which is sampled by the so called transition kernel T, where

$$T(\mathbf{r}' \rightarrow \mathbf{r}, \mathbf{\Omega}', E')dV$$

is the probability that a particle starting a flight path at $\mathbf{r}'$ will have its next collision in dV at $\mathbf{r}$. After the new collision point has been sampled, the collision process must be sampled. This is done by the collision kernel C, where

$$C(\mathbf{r}, \mathbf{\Omega}' \rightarrow \mathbf{\Omega}, E' \rightarrow E)d\Omega dE$$

is the probability that a particle will exit a collision in $\mathbf{r}$ with direction $\mathbf{\Omega}$ in $d\Omega$ and energy E in dE. Note that from the Boltzmann transport equation, we define $\chi(P)$ as

$$\chi(P) = S(P) + \int T(\mathbf{r}' \rightarrow \mathbf{r}, \mathbf{\Omega}', E')C(\mathbf{r}, \mathbf{\Omega}' \rightarrow \mathbf{\Omega}, E' \rightarrow E)\chi(P')dP \tag{1}$$

where $S(P)$ is source of the particles at P. We can see that if we start at a point P', we can sample emission density of the next point P by first sampling the transition kernel $T(\mathbf{r}' \rightarrow \mathbf{r}, \mathbf{\Omega}', E')$ in order to select a new collision point, and after that we sample the collision kernel $C(\mathbf{r}, \mathbf{\Omega}' \rightarrow \mathbf{\Omega}, E' \rightarrow E)$ in order to get the next sample of the emission density. If we want to sample the collision density $\psi(P)$, where

$$\psi(P) = S_1(P) + \int C(\mathbf{r}, \mathbf{\Omega}' \rightarrow \mathbf{\Omega}, E' \rightarrow E)T(\mathbf{r}' \rightarrow \mathbf{r}, \mathbf{\Omega}, E)\psi(P')dP \tag{2}$$

we start by sampling the collision kernel (since we are entering a collision, we already know the position of the collision) and then sample the transport kernel for the next collision as already shown above. Note the use of $S_1(P)$ instead of $S(P)$, as this is the source of first collisions, not the true particle source.

As stated, we can express the adjoint forms of our equations in terms of importance. We can label

$$\psi^*(P) = \eta_\psi(P) + \int L(P \rightarrow P')\psi^*(P')dP' \tag{3}$$

as the importance of a neutron entering a collision at P. Note that since the path is that of the adjoint particle, we are using the adjoint transport kernel, with the arguments P and P' interchanged, and the source density is actually the detector response function of the forward simulation

$$\eta_\psi = \frac{\eta_\phi}{\Sigma_t} \tag{4}$$

Similarly for $\chi^*(P)$,

$$\chi^*(P) = \eta_\chi(P) + \int K(P \rightarrow P')\chi^*(P')dP' \tag{5}$$

is the importance of a neutron leaving a collision/source at P and starting a new flight path.

1.2. Multiplying Systems

In a multiplying system, we can express the eigenvalue k_{eff} as the ratio of source neutrons of two successive neutron generations:

$$k_{eff} = \frac{\int S_{n+1}(P)dP}{\int S_n(P)dP} \tag{6}$$

Starting with a normalized source $S_n(P)$ we can regard k_{eff} as the detector response

$$R = k_{eff} = \int S_{n+1}(P)dP = \int \frac{\nu\Sigma_f}{\Sigma_t}\psi_{n+1}(P)dP \tag{7}$$

In the equation adjoint to eq. 7, the source term in adjoint equation becomes $\eta_\psi = \frac{\nu\Sigma_f}{\Sigma_t}$. Therefore, by running an adjoint shielding calculation with the above source term, we can obtain the adjoints and then use them in order to bias our forward criticality problem.

2. BIASING TECHNIQUE

We can separate biasing the simulation of a particle's path into 3 stages: Biasing the energy and direction of a particle at generation time (ie biasing the fission source), biasing the path of a particle until its next collision and biasing the energy and direction after a collision. These forms of biasing can be applied independently of each other, since the statistical weight of the particle is corrected after the biasing to ensure an unbiased simulation at every stage.

2.1. Fission Source Biasing

In the scheme, when using a multigroup treatment, the energy of neutrons generated by fission at a point $P(\mathbf{r}, g, \mathbf{\Omega})$ is selected from the fission spectrum χ_f biased by the direction averaged adjoint function $\chi^*(\mathbf{r}, g)$. In general terms:

$$\overline{S}(P) = S(P) \times \frac{\chi^*(P)}{R} \tag{8}$$

where R is the expected detector response, which we set to an estimate for the first particle generation and is automatically recalculated after each generation.

The particle direction is then selected from the isotropic distribution biased by the directional adjoint function, in a way similar to collision biasing, shown later. As stated earlier, after biasing we need to alter the statistical weight of the particle. In this case, it needs to be set inversely proportional to the adjoint function for the selected energy group and direction. This requires appropriate normalization of the weights of all particles in a new batch. In practice, when using a multigroup treatment, the initial weight is set equal to

$$w(\mathbf{r}, g, \mathbf{\Omega}) = \frac{\sum_{g"} \chi_{fis,g"}\chi^*_{g"}(\mathbf{r})}{\chi^*_g(\mathbf{r}, \mathbf{\Omega})} \tag{9}$$

2.2. Transition Biasing

For biasing the transition kernel to select a new collision site, the scheme requires the adjoint function ψ^* for particles entering a collision at P as the biasing function. The new collision site is therefore selected from the biased transition kernel

$$\overline{T}(\mathbf{r}' \rightarrow \mathbf{r}, E, \Omega) = \frac{T(\mathbf{r}' \rightarrow \mathbf{r}, E, \Omega)\psi^*(\mathbf{r}, E, \Omega)}{\int T(\mathbf{r}' \rightarrow \mathbf{r}", E, \Omega)\psi^*(\mathbf{r}", E, \Omega)dV"} \tag{10}$$

In practice, the new collision site needs to be selected from the normalized probability

$$\xi = \frac{\int_0^s T(\mathbf{r} \rightarrow \mathbf{r} + s'\mathbf{\Omega}, E, \mathbf{\Omega})\psi^*(\mathbf{r} + s'\mathbf{\Omega}, g, \mathbf{\Omega}_m)ds'}{\int_0^{s_{max}} T(\mathbf{r} \rightarrow \mathbf{r} + s'\mathbf{\Omega}, E, \mathbf{\Omega})\psi^*(\mathbf{r} + s'\mathbf{\Omega}, g, \mathbf{\Omega}_m)ds'} \tag{11}$$

where ξ is a random number, uniformly distributed between 0 and 1, $\mathbf{\Omega_m}$ is the direction cosine and g is the current energy group of the particle. After selecting ξ, we can now calculate the value of $\mathbf{r}$ using a root-finding iterative method, such as the regula falsi method, until $\mathbf{r}$ converges.

2.3. Collision Biasing

In order to apply the biasing scheme to the collision kernel, we have to bias it by $\chi^*(x)$, the importance of the particles that are exiting a collision at x. The biasing is done in two steps - initially, we select the new energy group g after the collision with probability

$$\overline{p}(g|\mathbf{r}, g') = \frac{\Sigma_s(\mathbf{r}, g' \rightarrow g)\chi^*(\mathbf{r}, g)}{\sum_{g''} \Sigma_s(\mathbf{r}, g' \rightarrow g'')\chi^*(\mathbf{r}, g'')} \tag{12}$$

After that, we need to select an outgoing direction, with probability

$$\overline{p}(\mathbf{\Omega}|\mathbf{r}, g', g) = \frac{\Sigma_s(\mathbf{r}, \mathbf{\Omega}' \rightarrow \mathbf{\Omega}|g', g)\chi^*(\mathbf{r}, g, \mathbf{\Omega})}{\int \Sigma_s(\mathbf{r}, \mathbf{\Omega}' \rightarrow \mathbf{\Omega}''|g', g)\chi^*(\mathbf{r}, g, \mathbf{\Omega}")d\mathbf{\Omega}''} \tag{13}$$

2.4. Fission Selection Biasing

No adjoint biasing is performed during the fission process. However, in order to avoid a diverging number of particles because of a too low or too high initial estimate of k_{eff}, we regulate the number of neutrons at each generation through the following factor F, applied during the selection of the number of new fission neutrons:

$$F = \frac{N_{init}}{k_{est} \times N_{curr} \times w_{avg}} \tag{14}$$

where N_{init} is the initial number of neutrons, N_{curr} is number of neutrons in the current generation, k_{est} is the estimate of k_{eff} from the previous generation and w_{avg} is the average initial weight of particles, calculated from eq. 9. The number of new particles produced by fission is then selected from

$$N_{new} = w \times \frac{\nu\Sigma_f}{\Sigma_t} + \xi \tag{15}$$

where w is the statistical weight of the particle, ξ a random number in $[0, 1]$, $\nu\Sigma_f$ the expected number of particles produced by a fission and Σ_t the total cross section.

3. IMPLEMENTATION

In order to test the scheme, we have chosen two test cases, one of a homogeneous slab reactor and another of a loosely coupled system consisting of a 3-zone slab where water separates two fissionable zones.

In the first case, as validation for the scheme, we have used the two-direction model, proposed by Hoogenboom[3]. The model limits the transport of particles to the $\pm x$ direction, and although it remains a true transport model, the equations describing the particle transport become differential, diffusion type equations. This means that it is possible to obtain analytical solutions for both the forward and adjoint equations, which makes the model an excellent one for verifying computational biasing schemes.

By limiting our transport in the $\pm x$ directions, the transition and scattering kernels become quite simplified. The path between two collisions becomes $s = |x - x'|$, while the direction cosine Ω can only take values of $+1$ if the particle is moving to the right or -1 if it is moving to the left.

In order to test the performance of transition and collision biasing in the 2-direction model, we used a homogeneous 1-D infinite slab reactor of half-width $a = 10$ cm, containing an artificial fissionable material. The energy dependence is treated in a 3-group approximation, in order to allow comparison to the theoretical solutions we have obtained. The total cross section for each group was chosen equal in order to have the same (converged) spatial distribution for each group, which allows an analytical calculation of k_{eff}. The 3-group scattering matrix allowed only downscattering. Vacuum boundary conditions were imposed at both boundaries of the system.

The adjoint function was obtained from the one-dimensional S_N code XSDRN[4]. A cross section library for this code was prepared with the same cross sections as for the Monte Carlo runs. Each calculation was carried over 30 batches of 1000 particles each. The initial particle source was uniformly distributed, with the particle weight appropriately modified.

The cases we studied included:

- No biasing throughout the reactor
- No collision biasing, transition biasing if $|x| > 0.6 * a$
- No collision biasing, full transition biasing
- Collision biasing, no transition biasing
- Full biasing

The reasoning behind the second case is that, when calculating the adjoints analytically for our case, we noticed that they are rather flat in the central part of the geometry. Therefore, in such cases, the transition biasing effect should be negligible and therefore we should notice a significant speedup of the calculation without a large corresponding variance penalty.

In order to further test the effectiveness of collision biasing we examined the case of a loosely-coupled system. The case we tested comes from the OECD/NEA benchmarks on source convergence, and more specifically case 4 in benchmark 3[5]. It is composed of a one-dimensional infinite slab, as shown in fig.1. A water slab of thickness 30 cm separates Unit1 and Unit 2, two Uranyl solutions of thickness 20 and 12 cm respectively. A vacuum boundary condition is imposed on both sides of the slab.

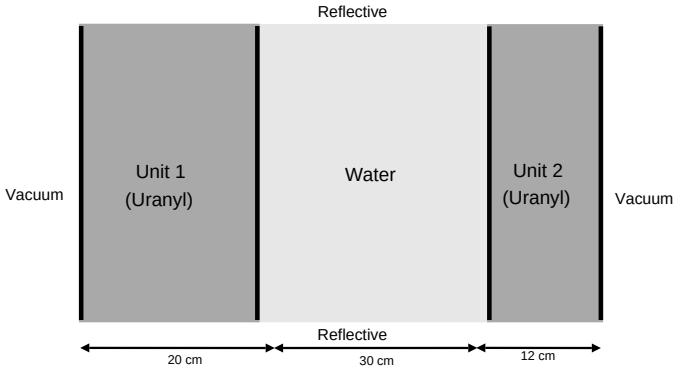

Figure 1. Geometry of the infinite slab geometry used with the 3D transport model

Using the compositions given in the benchmark specification, a reduced set of zone-mixed cross-sections for the 3 zones was prepared through the use of the SCALE code system. For simplicity, only the P0 scattering matrix was used.

During the initial simulations, we found out that the 2-direction model exhibited excessive leakage (because of the ± 1 direction cosines), which resulted in a much lower than expected k_{eff}. We therefore decided, for a better approximation, to use full 3-D transport for the system. The directional adjoints were input from a TORT[6] calculation for the same system, using 16 directions. It is recommended in the benchmark specification to run the simulations with batch size of 2000 particles, but we chose a smaller batch size, 500 particles, in order to examine whether undersampling occurred and whether the biasing had a positive or negative effect on the sampling. A larger number of batches, 60, was selected in order to allow the source to converge from the uniform initial source distribution, while the 20 first batches were discarded for the final statistics. We used a Russian roulette weight threshold of 10^{-3} in order to allow for sufficient interactions of a particle before it is destroyed.

3.1. Entropy

One thing to note is that k_{eff} convergence is not a reliable method of determining convergence of the source distribution, as k_{eff} can converge faster than the source. It is necessary to know when source convergence has been achieved, since only then should k_{eff} results be taken into account. In order to diagnose whether source convergence had been achieved in every case, we measured the Shannon entropy of the source distribution[7]. It is defined as

$$H(S) = -\sum_{i=1}^{N} p_i \times log_2(p_i) \tag{16}$$

where N is the number of spatial bins we divide our geometry into and $p_i = \frac{\text{\# of source particles in bin i}}{\text{total \# of source particles}}$. Convergence of the Shannon entropy in time indicates convergence of the fission source.

4. RESULTS AND DISCUSSION

The results of the k_{eff} calculation for the homogeneous slab are shown in table I. We can see that, as expected, transition biasing suffers from a large performance hit, since the continuous recalculation of the biased kernel during the root-finding method costs significant CPU time. Therefore, although it results in the lowest variance of all possible methods, it cannot be considered a viable method unless computational techniques in order to speed it up are applied. This effect would be even more highlighted in a more complicated 3D geometry. when all geometrical elements would have to be taken into account when normalizing the transition kernel. Limiting transition biasing to the outer ranges of our reactor improved the performance of transition biasing significantly, however it is still nowhere near the performance of collision biasing.

Table I. Results from the different biasing methods in the 2-direction model

Type of biasing	k_{eff} result	Total time (min)	Figure of merit
No biasing	1.12947 ± 0.00349	0.36	807
Transition if $\|x\| > 0.6 * a$	1.12566 ± 0.00139	3.35	215
Transition only	1.13046 ± 0.00108	8.91	104
Collision only	1.12845 ± 0.00147	0.49	1388
Full biasing	1.12880 ± 0.00123	9.28	88
Theoretical Value	1.12865		

Collision biasing, on the other hand, performed very well; its computational overhead compared to an unbiased calculation is minimal, and the reduction in variance big enough to allow it a higher figure of merit than the unbiased case. Since collision biasing only requires the group and energy dependant adjoint at a point in space, it is also much easier to implement in an existing code than full transport biasing, and suffers from no drawback when used in a full 3D geometry.

Of course, it is possible to improve the performance of transition biasing by using different computational methods. For example, although using cell-averaged adjoints means that we no longer use exact biasing and therefore the principle of the zero-variance scheme is not valid anymore, the computational gain is tremendous, since the code does not need to recalculate the adjoints during the selection of a new collision site. A trade-off then needs to be made between the size of the cells within which the adjoint is averaged and the accuracy we wish to have in our calculation. However, as we can already see, the simplified method of obtaining an adjoint at a point (interpolation in our case, or possible cell averaging), affects the transition biasing, leading to results which, including statistics, diverge slightly from the expected value

obtained through theoretical calculations. As the biased collision kernel is normalized by adjoints calculated for the same cell/point, this is not an issue in collision biasing, again apparent in the results.

We can see the source convergence achieved by each method of biasing in fig.2. In such a simple geometry, and starting from a flat source distribution, we can see that the source converges very fast even without biasing, therefore the biasing schemes do not provide a significant advantage. It is possible that in such cases the relative entropy can provide a better measure for the convergence of the system, which will be investigated in the future.

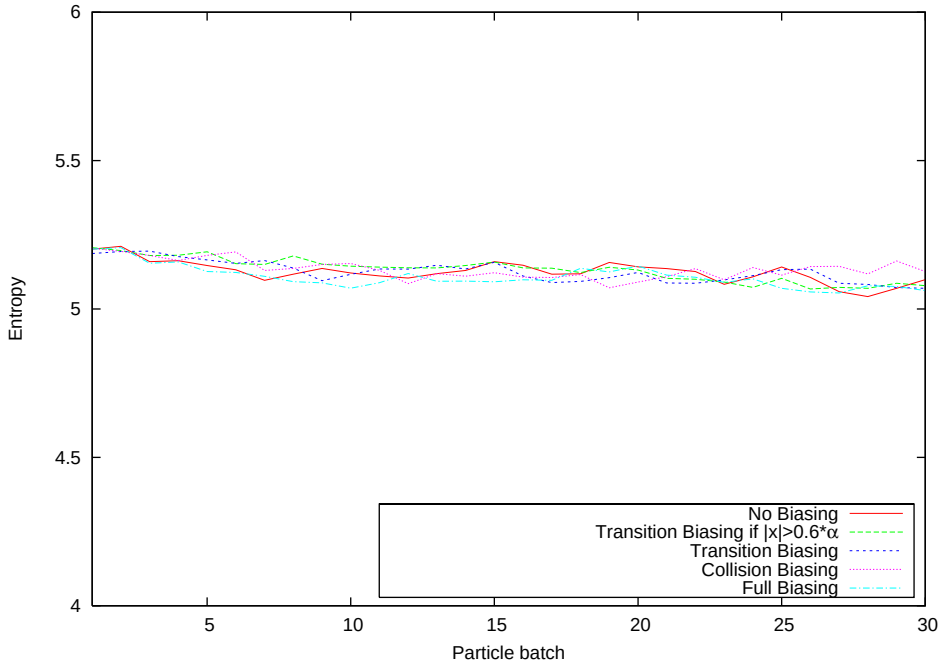

Figure 2. Shannon entropy of the source distribution over a full calculation

Since we are using implicit capture for our model, if the particle weight drops below a certain threshold, the particle has to survive a Russian roulette game in order to continue its path. This threshold directly affects the FOM of the simulation, especially the calculation time.

In fig.3 we can see the effect of changing the weight threshold of the Russian roulette game in the collision biasing and the unbiased case. The values are shown relatively to an unbiased calculation with a Russian roulette weight threshold of 0.1. It is shown that collision biasing remains the better choice regardless of the threshold used, and the ratio of FOM remains relatively constant. Both cases show a reduction in the FOM with lower threshold, since the computational time cost increases more than the gain in the number of scoring sampled. However, it must be noted that a lower threshold should provide an advantage in specific cases, since the longer flight path means more opportunities to score for a particle, which could be beneficial in cases of having to use a small statistical sample.

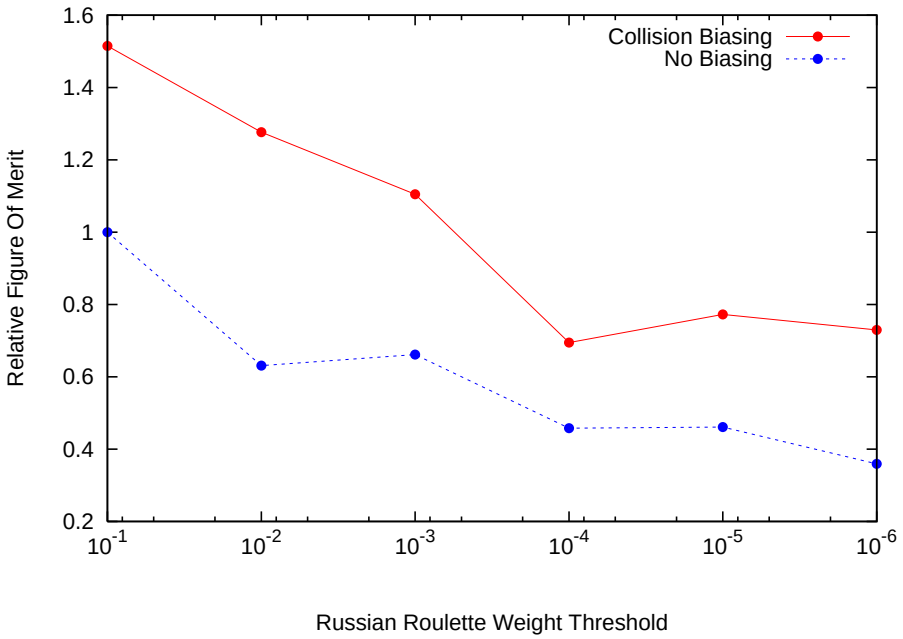

Figure 3. Relative figure of merit with lowering Russian roulette weight threshold

For the loosely-coupled system, the results are shown in table II. Note that because of the use of the $P0$ scattering matrix only, the k_{eff} of the system is higher than expected. We were able to replicate the high values found here by performing TORT, XSDRN and MCNP calculations without higher orders of scattering. Also, the calculation times here are longer, as the complexity of the system and the use of 3D transport (as well as a much more generalized model, to allow for the different zones) add CPU time compared to the homogeneous slab used in the previous case.

Table II. Collision biasing performance in the 3-zone slab geometry

Type of biasing	k_{eff} result	Total time (min)	Figure of merit
XSDRN (S96)	1.0472		
No biasing	1.0443 ± 0.0095	31.6	3.32
Collision biasing	1.0438 ± 0.0085	8.2	13.25

In this system we observe multiple benefits of using collision biasing. Not only is the variance of the calculation reduced, but also because of the effect of the adjoints on the statistical weight of each particle after collision, the particles are destroyed by the Russian roulette in a smaller timeframe than the unbiased case, while at the same time providing the expected response (ie number of new fissions). Therefore the FOM in this case is significantly higher than in the unbiased calculation. Note that in the biased calculation the time required for the adjoint run of the S_N code is included, but it is in fact negligible when compared to the Monte Carlo run itself. Once again, the Shannon entropy did not indicate a significant improvement

of the biasing to the source convergence, therefor it needs to be determined whether the relative entropy is a better measure of the convergence of the source distribution.

5. CONCLUSIONS

It has been shown that biasing the transport kernel directly with the adjoint emission and collision densities can result in a significant reduction to the variance of the simulation. In terms of computational efficiency, collision biasing is always more efficient than transition biasing, since the latter requires calculation of integrals which is computationally taxing. However, transition biasing offers a higher absolute variance reduction gain, which is expected to increase in a more complicated geometry, where directional adjoints play a much more important role than in the simple cases examined here. Nevertheless, both biasing methods have a large window for improvement through computational techniques. Further investigation is needed in order to determine whether using cell averaged adjoints can improve the transition biasing enough to make it a feasible method of variance reduction, while the same techniques can be applied to collision biasing to further improve the efficiency of the method. Also, the effect of the biasing in convergence of the fission source needs to be further investigated, with an appropriate measure for the convergence used.

ACKNOWLEDGEMENTS

This research presented in this paper is supported by the European Commission under the NURESIM Integrated Project[8], within the 6th Framework Program, project number #516560.

REFERENCES

[1] S.Christoforou and J.E. Hoogenboom, “A Zero-Variance-Based Scheme for Variance Reduction in Monte Carlo Criticality” *Proceedings of PHYSOR 2006*, Vancouver, Canada (September 2006).

[2] J. Lewins, *Importance, The Adjoint Function*, Pergamon Press, New York, USA (1965)

[3] J.E. Hoogenboom, “The Two-Direction Neutron Particle Transport Model: A Useful Tool for Research and Education”, submitted to *Transport Theory and Statistical Physics* (2006).

[4] N.M. Greene and C.W. Craven Jr, ”XSDRN, A discrete ordinates spectral averaging code”, ORNL/TM-2500, Oak Ridge National Laboratory, Oak Ridge, TN, USA (1969).

[5] Nuclear Energy Agency NSC Working Party on Nuclear Criticality Safety (WPNCS) homepage, http://www.nea.fr/html/science/wpncs/.

[6] W.A. Rhoades and D.B. Simpson, ”The TORT three-dimensional discrete ordinates neutron/photon transport code”, ORNL/TM-13221, Oak Ridge National Laboratory, Oak Ridge, TN, USA (1997).

[7] T. Ueki and F.B. Brown, “Stationarity modeling and informatics-based diagnostics in Monte Carlo criticality calculations”, *Nuclear Science & Engineering*, **149**, pp.38-50 (2005).

[8] D. G. Cacuci et al., “NURESIM: A European platform for nuclear reactor simulation”, *Proceedings of FISA 2006, Conference on EU Research and Training in Reactor Systems*, Luxembourg (March 2006).