Self-adjusting Importances for the Acceleration of MCBEND

Edmund Shuttleworth

Abstract

The principal method of variance reduction in MCBEND is the use of splitting and Russian roulette under the control of an importance map. Theoretically, an ideal importance is a solution of the adjoint problem; in practice, an approximate adjoint is usually adequate and is often derived using diffusion methods. In some cases, the diffusion approximation introduces inaccuracies in the importance that lead to inefficiencies. This study considers the possibility of refining an initial importance during the course of a calculation.

The adjoint flux in a given space/energy cell may be interpreted as the expected score produced by introducing a particle into that cell. Adjoint fluxes may be estimated by recording the cells visited in a given particle history and, if the particle scores, accumulating appropriate contributions in these cells. This simple principle becomes complicated in realistic Monte Carlo calculations owing to branching events in the particle history and weight changes introduced by variance reduction. A practical implementation of the adjoint estimator and a strategy for merging the results with the initial importance map are described.

In practical problems, no finite scores at a detector will be produced in realistic execution times without an initial importance map. Although a solution of the adjoint diffusion equation meets this need in many cases, there are some problems (particularly those dominated by voids) for which the method is too inaccurate. An alternative has been derived that uses Monte Carlo tracking to obtain an initial adjoint. Particle histories are started in cells containing the detector then in cells that are progressively more remote. Tracking of each particle is restricted to the point where it enters a cell for which the adjoint has been estimated. The method is relatively slow and approximate but is valid in problems for which a diffusion solution would break down.

1. INTRODUCTION.

Application of the Monte Carlo method to practical problems of radiation transport requires biasing of the sampling process to concentrate effort on particles expected to make a useful score. A common acceleration technique uses the complementary processes of splitting and Russian roulette (S/R). The former subdivides particles that are travelling in favourable directions to enhance their probability of survival; the latter reduces the population of particles travelling in unfavourable directions. Appropriate changes to the particle weights are made to remove bias from the results.

In the Monte Carlo code MCBEND [1], the S/R mechanism is controlled by an importance function which varies with position and energy. (In some applications, direction and time dependence might be added). Each particle being tracked carries a weight that is intermittently compared with the local value of the importance function. When the ratio exceeds a control limit the particle is split; when it falls below a control limit it is subjected to Russian roulette.

A theoretical, ideal importance function is an adjoint solution to the defined transport problem - which is usually as difficult to obtain as the forward solution. Fortunately, an approximate solution of the adjoint, when used as an importance function, can dramatically improve the efficiency of a Monte Carlo calculation relative to execution in analogue mode with no acceleration.

This paper considers two aspects of acceleration: the generation of approximate importance functions and the possibility of improving them during the course of a calculation.

2. Estimating an improved importance function.

The adjoint flux in a given space/energy cell may be interpreted as the expected score produced by introducing a particle into that cell. An estimator based on this definition can be used during a forward calculation to score the adjoint. The results may be combined with an initial guess of the importance function to improve the efficiency of the calculation. Consider the following particle history.

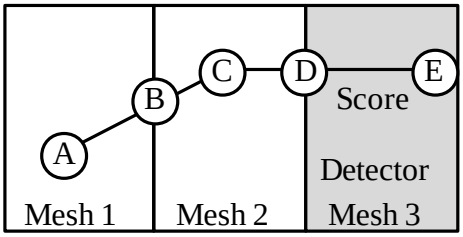

Fig.1. Particle history in a segment of an importance mesh.

A particle born at A crosses importance mesh boundaries at B and D and undergoes a collision at C. The final track segment between D and E is in an identified detector. The score may be added into an accumulator for estimating the adjoint in mesh 1 where the particle was born.

The portion of the particle history BCDE may be considered as that of a particle started at B so a contribution to the adjoint estimate for mesh 2 may also be made. This logic applies to all the events: the cell into which the particle emerges after an event may be considered as a starting point for the remainder of the history. A contribution to the adjoint estimate may be made in every cell of the importance map through which the particle has passed. After collisions, the cell is identified by the mesh volume in which the collision occurred and the importance energy group of the scattered particle.

It is clear from this simple example that a contribution to the adjoint estimator cannot be made as the particle passes through a given cell because the detector score will not be known at that point. A record must be preserved of all the cells through which a particle passes and what scores are made. When a history is complete the records may be retraced to attribute appropriate contributions to each cell.

The contribution to the estimator for a given cell is given by:

$$\frac{\text{The score made by particle fragments that have had events in the cell}}{\text{The total weight of particle fragments emerging from events in the cell}}$$

The numerator is determined as follows.

> When a track segment registers a score in a nominated detector, the current sequence of events is backtracked to the source. The score is added into an accumulator for each cell in which an event has occurred on the route to the scoring segment. If several events occur in the same cell then the score is added in for each event.

The denominator is determined as follows.

> As each event occurs, the weight emerging is added into an accumulator for the cell appropriate to the parameters of particles emerging from the event.

This process accommodates all the effects of branching in the particle history and the weight changes associated with splitting/Russian roulette. By accumulating the squares of the contributions an estimate can be made of the variance of the estimated adjoint.

Contributions to the estimator are only obtained when a particle reaches a detector. In most practical calculations, such an event is improbable without the aid of an initial importance map. The following sections consider various ways in which this might be generated. The methods are approximate and, in general, are intended as a prelude to the above process. None of the methods is universally applicable but the range offered is such that one can usually be found that suits a particular problem.

3. Methods of generating importance functions.

3.1 Adjoint Diffusion solutions.

Diffusion theory is an approximation to transport theory; in adjoint mode it can be used to create an approximate importance function. This option has existed in MCBEND for many years and is more fully described in [2]. The problem space is overlaid by an orthogonal mesh (XYZ or RθZ) in which the adjoint, multigroup diffusion equation is solved by finite difference methods. Special diffusion constants are used and each group is represented by twin components - asymptotic and transient. The formulation overcomes some of the conventional limitations of diffusion methods and can be applied to neutron, gamma and coupled calculations. The results are an approximation to the problem adjoint, and are defined for each cell of the overlaid mesh and for each group of the diffusion calculation. During the Monte Carlo calculation, the progress of particles through the meshes and groups of the adjoint is monitored to control the application of S/R.

This method is simple, quick, convenient, robust and applicable to many practical cases. Its principal limitation is its inability to produce efficient importance maps in geometries that are dominated by voids and ducts.

3.2 Adjoint Monte Carlo

It is possible to execute a Monte Carlo calculation in adjoint mode to generate an importance for a forward calculation. However, the adjoint calculation also requires some form of acceleration. In this application, an importance can be evolved during sampling by raising the importance of regions that are being scored infrequently (or not at all) at the expense of those that are being well scored, this being known as the ADJUST option. Eventually, an estimate of the adjoint over the entire problem can be generated - though its statistical uncertainties in some regions may be poor. This preliminary step to the required forward calculation can be quite time consuming but will work where diffusion methods fail.

3.3 Sampling the adjoint.

Recalling that the adjoint is the expected score made by introducing a particle at a point, it could theoretically be produced by sprinkling the entire problem with test sources and tracking them until they die or score. Scoring samples will make a contribution to the adjoint estimate at their birth site. This is clearly impractical for a large case since the probability of scoring at a detector for a particle born in a remote region is extremely small - particularly in the absence of an importance map.

A refinement of this technique has been implemented in MCBEND. Initially, test samples are generated in a cell of the importance mesh that contains a nominated detector. These have a high probability of scoring and soon generate a reliable value for the adjoint in that particular cell. Test particles are then generated in the cells that are neighbours of the target and tracked. If they cross into the target cell their expected score is known from the value of the adjoint that was calculated. The history of these particles is immediately terminated. Particles that cross into unscored cells are allowed a limited number of events before being abandoned. When the immediate neighbours of the target cell have been completed, their adjoint values are estimated from the accumulated scores. Processing then moves to the next layer of cells surrounding those completed. Eventually, the entire importance map will be filled. In regions of the problem remote from the detector, test samples only have to be tracked as far as their scored neighbours to obtain a score rather than all the way to the detector. The method has been assigned the name ‘recursive’; its principles are illustrated in a one-dimensional mesh by Figure 2.

The method extends to three dimensional meshes. When the importance map is represented in multigroup form, a solution is first obtained for the lowest energy group followed by groups in the order of increasing energy. When tracking test particles in the higher groups, a down scatter into a scored group terminates the history with a known, expected score.

The adjoint estimate generated by this method can be very approximate and generally deteriorates with remoteness from the detector. The method is an alternative when diffusion methods are inapplicable and can form a useful starting point for subsequent improvement.

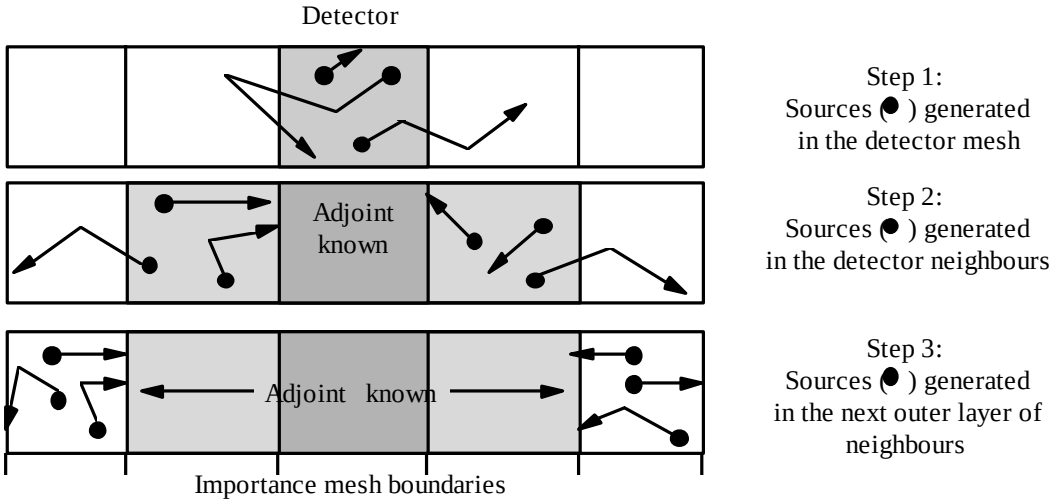

Fig.2. Steps in the recursive estimate of importance.

3.4 Density adjustments.

The principal reason why few particles can travel from source to detector without the aid of an importance function is the attenuating effect of the intervening materials. If the densities of all materials in the problem are sufficiently reduced then particles can reach the detector unaided. For example: if attenuation is approximately exponential then it can be reduced from, say, 10^8 to 10^2 by quartering all the densities. Using the technique described earlier, an adjoint for this non-physical problem can be scored. This adjoint can then be used to accelerate a variation in which the density reduction is less severe. Repetition of this process eventually produces an importance map that can be used in the real problem. In the diluted variation, the probability of scoring is governed by geometrical effects and the method is unsuitable for small detector volumes remote from the source.

4. Example Application.

The following simple test configuration has been used to illustrate the techniques described above.

The case is small enough to be solved in analogue mode but large enough to benefit from use of an importance map. It contains realistic materials and energy ranges. The objective is to calculate the neutron flux above 1.0 MeV in the detector from a fission source. The importance map is represented in eight energy groups between 10.0 and 0.6 MeV and contains 588 spatial meshes - a total of 4704 space/energy cells. The hatched area of iron is replaced by a void in a later model.

The following variations have been used.

1 Execution with no importance map.
2 Using an importance map from solution of the adjoint diffusion equation.
3 Using an importance map generated by the recursive method of §3.3 which took 6.0 minutes.
4 Using an importance map by density adjustments:
initial execution (5minutes) with no importance and densities scaled by a factor 0.25:
secondary execution (10 minutes) using the scored adjoint from the initial execution and densities scaled by a factor 0.5:
final execution using the scored adjoint from the secondary execution and real densities.

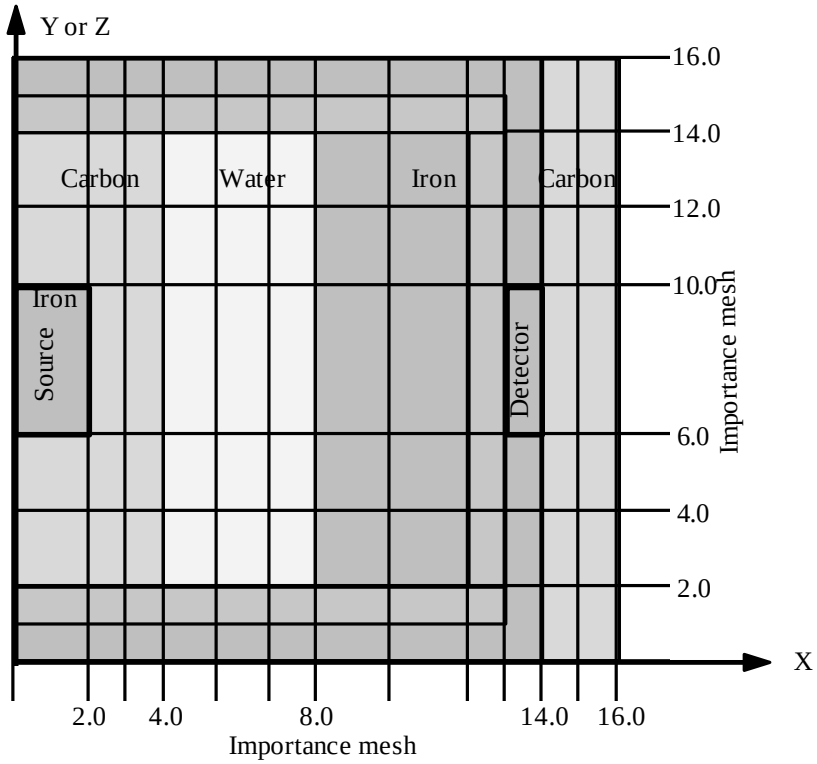

Fig.3. Example test geometry.

The efficiencies of the four variations are compared in the following table. The times quoted to reach a variance of 1% on the scored response include the times taken to generate the importance maps.

Table 1. Results using various methods of generating importance maps.

Variation	Description	Result	Time (minutes)	Relative efficiency
1	No importance	5.84E-3 ± 1.0%	368	1.0
2	Adjoint diffusion	5.76E-3 ± 1.0%	109	3.4
3	Recursive	5.83E-3 ± 1.0%	124	3.0
4	Density variation	5.83E-3 ± 1.0%	220	1.7

For such a small problem (the attenuation of the response between source and detector is only a factor 180) the gains in efficiency given by use of an importance function are not dramatic. The greatest gain is given by the diffusion solution - a typical conclusion when the geometry is suitable for the use of this method.

A variation on this test case was made by replacing the hatched area of Fig. 3 by a void region and adding extra importance mesh boundaries along the new surfaces created. In a diffusion solution, the presence of the void tends to flatten the gradients of the adjoint leading to a less efficient mesh. When the methods described above are applied to this case the results are as follows.

Table 2. Results in the presence of a void.

Variation	Description	Result	Time (minutes)	Relative efficiency
1	No importance	5.83E-3 ± 1.0%	368	1.0
2	Adjoint diffusion	5.88E-3 ± 1.0%	184	2.0
3	Recursive	5.91E-3 ± 1.0%	127	2.9
4	Density variation	5.92E-3 ± 1.0%	225	1.6

The scored responses show that the presence of the void has little effect on the attenuation through the problem but significantly reduces the efficiency of the importance map generated by an adjoint diffusion calculation. The effects on the efficiency of the recursive and density adjustment methods are negligible and may be attributed, in part, to the increase in the number of importance meshes.

5. Merging Adjoint estimators with an Initial Importance.

At a certain point in the calculation the results for the adjoint estimator in various space/energy cells of the importance map may be classified into three categories:

- those that are reasonably well scored i.e. those that lie on the principal route from source to detector:
- those that are poorly scored i.e. those peripheral to the principal route:
- those that are unscored i.e. those remote from the principal route.

A method is required of merging the partially completed adjoint solution with the initial estimate of the importance map to take advantage of the information learned. Trials have shown that simple substitution of the well scored adjoint estimates for the initial estimates of importance in corresponding cells is unsatisfactory. The transition between the retained, initial importance values and the substituted ones can lead to large step changes. The splitting/Russian roulette process then becomes inefficient or, in extreme cases, unstable.

A better solution is to scale the initial importances outside the envelope of well scored adjoint estimates so that there is no change in magnitude at the transition. The gradients outside the scored envelope are preserved.

These alternative techniques are illustrated, in one dimension, in the following sketches.

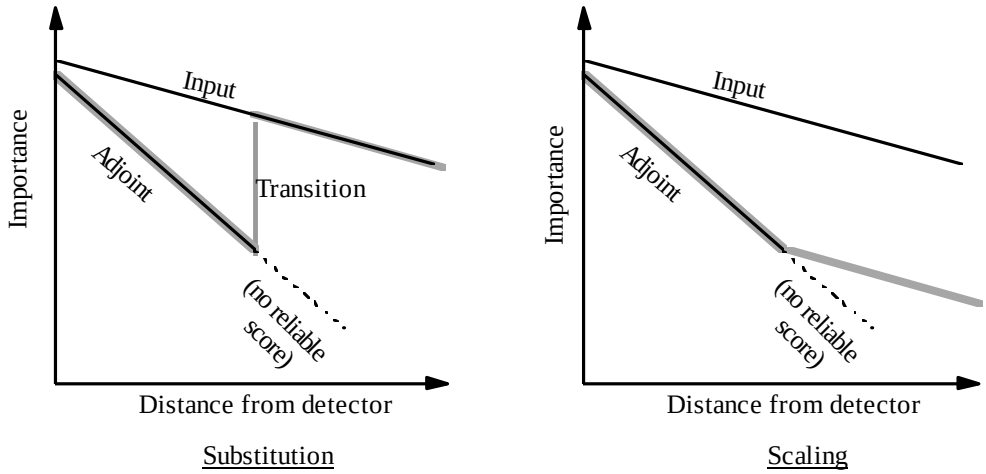

Fig.4. Methods of Merging Importance Maps.

6. Frequency of Importance Map revision.

Methods have been outlined for the production of initial importance maps and their subsequent merging during the course of a calculation with estimates of the scored adjoint. MCBEND allows the frequency of merging to be controlled by specifying the number of started samples, the sampling time or the number of samples scoring at a nominated detector. The merging may be carried out once at a specified point in the calculation or at regular intervals. An optimum strategy has yet to be evaluated but the following points have emerged from preliminary studies.

When the importance map is revised, the weight of individual particles reaching the detector will change though the expected results should, of course remain constant. Consider the extreme case where an initial, flat importance is replaced part way through the calculation by a realistic adjoint. Before replacement, few particles will reach the detector but their weight will be that with which they were born in the source. After replacement, many more particles will reach the detector but their weights will be reduced by the S/R process to remove bias. In the calculation as a whole, there will be a large spread in the weight of scoring samples which is contrary to the objective of variance reduction. In less extreme cases, there will still be a tendency for a change in the importance map to cause a spread of scoring weights. If this does not occur then it is probable that the revision was minor and unnecessary.

If a single revision is made too early, the adjoint scores will not be very reliable and the envelope of cells that are sufficiently well scored will be small. The calculation risks introducing a spread of particle weights for a small change in the efficiency of overall importance map.

If a single change is made too late then the remainder of the calculation will be too short to benefit from any improvement.

If regular changes are made throughout the calculation then the weight of scoring samples will also keep changing with a deleterious effect on the final statistics.

At present, only two tentative recommendations can be made:

- if the importance map is reasonably good then leave it alone:
- if the importance map is very poor then revise it at regular intervals.

The first recommendation is supported by executions of test case (figure 1) using an adjoint diffusion solution for the initial importance. The case was executed for 100 minutes using seven different strategies for merging the estimated adjoint with the initial importance: no merge; merge after 5 or 10 or 20 or 40 or 80 minutes; merge every 10 minutes. The detector results for the first six cases all achieved a variance of 1.2%; the final variation achieved 1.3%. The nuances of the effect are lost because the variances are only printed to the nearest 0.1% but the general ineffectiveness of the revisions is illustrated.

The second recommendation is illustrated by the test case (figure 1) using a flat initial importance. The case was executed for 200 minutes using six different strategies for merging the estimated adjoint with the initial importance: no merge; merge after 20 or 40 or 80 or 160 minutes; merge every 50 minutes. The detector results for the cases achieved a variance of 1.4%, 1.7%, 1.4%, 1.4%, 1.4% 1.2% respectively.

7. Practical Applications.

The methods described above have been applied to a number of full-scale problems. One of these involved the calculation of pressure vessel damage rates below the core of a MAGNOX reactor. The characteristics of this problem are a complex steel support structure with many gas filled ducts. It is generally unsuited to the use of diffusion methods so the original assessment used an importance map generated by a lengthy, adjoint Monte Carlo calculation. The importance map was represented in RθZ geometry using 17 x 5 x 13 (=1105) spatial meshes and 15 energy groups - a total of 16,575 cells.

This case proved to be suitable for application of the density adjustment method using steps in which the densities were scaled by 0.125, 0.25, 0.5 and 1.0. The three preliminary steps were each executed for 2.5 hours; the final, physical step for 14 hours. The entire process was launched as a single execution

using the looping and parameterisation facilities of MCBEND to control the sampling times and density factors in each pass.

The variances of the results were compared with those obtained in a 14 hour execution of the original case with a Monte Carlo Adjoint. The comparison showed that the importance map produced by the new method was, as to be expected, less efficient - by about a factor 2.5. The case was also executed using the recursive method with a similar conclusion.

The approximate importance maps are not intended to compete in efficiency of execution with a proper adjoint produced by Monte Carlo or transport methods. Their merits lie in their ease of application in routine calculations when user effort has a higher premium than computing time.

8. Summary.

A range of methods has been described for the automated production of importance maps for acceleration of Monte Carlo calculations.

An algorithm has been defined whereby the adjoint solution of a problem can be estimated during the normal, forward execution.

Mechanisms have been described for the revision of an initial importance by the estimated adjoint at intervals during a calculation.

The methods have all been implemented in the Monte Carlo code MCBEND to provide a variety of options with a convenient user image.

Further development is needed to refine the strategies to be used when applying the new methods.

Acknowledgements.

The funding for these developments has been provided by the HSE/IMC Nuclear Safety Research Program

References.

1 S.J.Chucas et al. “ Advances in the Monte Carlo Code MCBEND” , Proceedings of 1996 Topical Meeting: Radiation Protection and Shielding. No. Falmouth, Massachusetts, USA, 1996

2 P.C Miller et al. “The Use of an Inbuilt Importance Generator for Acceleration of the Monte Carlo Code MCBEND” , Proceedings of the International Conference on the Physics of Fast Reactors. Marseilles, France, 1990

Manuscript information file

1 34
2 Self-adjusting Importances for the Acceleration of MCBEND
3 Author
4 Shuttleworth, Edmund
5 AEA Technology, Winfrith, Dorset, DT2 8DH, UK
6 +44 1305 251888
7 +44 1305 202746
8 ted.shuttleworth@aeat.com
9 End Author
10 Acceleration
11 Monte Carlo
12 Self-learning