Computer Physics Communications 185 (2014) 1628–1638

Contents lists available at ScienceDirect

Computer Physics Communications

journal homepage: www.elsevier.com/locate/cpc

Arbitrary perturbations in Monte Carlo neutral-particle transport

James Tickner *

CSIRO Centre for Process Science and Engineering, Building 67, Lucas Heights Science and Technology Centre, New Illawarra Road, Lucas Heights, NSW 2234, Australia

ARTICLE INFO

Article history:
Received 21 March 2013
Received in revised form
27 February 2014
Accepted 1 March 2014
Available online 11 March 2014

Keywords:
Monte Carlo
Perturbation methods
Neutral-particle transport
EGS

ABSTRACT

Monte Carlo techniques are widely used to model particle transport in complex models, as both the transport physics and level of geometric details can be simulated with arbitrary precision. The major draw-back of the Monte Carlo method is its computational cost. This is particularly true in design studies, where the effects of small changes in the model may be masked by statistical fluctuations unless prohibitively long simulation times are used. Perturbation methods have been developed to model the effects of small changes in material density, composition or reaction cross-sections. In this paper, I describe how this approach can be extended to allow nearly arbitrary perturbations in the transport problem specification to be made, including material properties, the model geometry and the radiation source description. The major problem, handling arbitrary variations in the model geometry, is overcome using a modified form of the Woodcock neutral-particle tracking algorithm. The approach has been implemented as an extension to the general-purpose Monte Carlo code EGSnrc. I discuss the details of this implementation, including how the specification of a perturbation simulation can be generated automatically from two or more unperturbed simulation models. I present an example of the application of the method to the modelling of a simple X-ray fluorescence instrument.

1. Introduction

Monte Carlo methods can be used to solve the radiation transport problem by performing a direct, microscopic simulation of individual radiation quanta. The energy, direction and other properties of a particle emitted by a radiation source are sampled from the appropriate distributions. The location and nature of the particle's interactions with the materials comprising the underlying system geometry are similarly sampled, the energy, type and nature of the particles arising from the interaction are determined, and the process repeated until the particle is captured, escapes or is otherwise terminated. Contributions to quantities of interest (*tallies*), such as flux or energy deposition, are recorded during the transport process. By simulating a large number of source particles N, average tally results can be estimated with a statistical precision that improves like $N^{-1/2}$.

The relatively slow convergence of the Monte Carlo method is a particular problem when estimating the effects produced by small variations in the problem specification. If two separate simulations are performed of slightly different problems, then in general the particle transport in the two runs quickly diverges, meaning that the statistical errors on the two simulations are uncorrelated. Very large number of events may then have to be simulated to estimate the effect of the variation with any confidence.

A partial solution to this problem is correlated-history sampling. Normally, the random sampling of source particle properties and subsequent interactions is carried out by drawing numbers sequentially from a pseudo-random number generator. Even if identical generators are used in separate simulations, as soon as one particle takes a different path (usually requiring a different number of random numbers to track its subsequent collisions), the pseudo-random number sequences used in the two simulations become desynchronised. From this point on, the events simulated in the two simulations are completely uncorrelated. The correlated-history approach insists that each simulated event starts at a known place in the random number sequence. A large number of random samples are allocated to each event, of which only a fraction are normally required. Unused samples are discarded at the end of each event and the random number generator is reset to the correct starting point for the next event to be simulated. For example, the MCNP code [1] allocates 152,917 random samples per event and uses a linear congruential generator that allows efficient skipping over the samples that are not required.

* Tel.: +61 8 8303 8430; fax: +61 2 9710 6732.
E-mail address: james.tickner@csiro.au.

http://dx.doi.org/10.1016/j.cpc.2014.03.003
0010-4655/

The correlated-history sampling approach improves the correlation between the errors on the separate runs, reducing the error on the difference in the tally quantity of interest. However, for individual events where the variation in the problem specification causes particles to follow different paths, particle transport in the two runs is uncorrelated as soon as the divergence occurs, so the correlation is not perfect. Further, by requiring two separate runs to be performed, the approach doubles the required simulation time.

For changes in material composition or density, or to estimate uncertainties due to using different cross-section evaluations, the perturbation approach is well known [1–8]. A single set of particle histories is generated and used to calculate tally quantities for both the original problem and for one or more variations, termed *perturbations*. This leads to very strong correlation in the statistical errors for the original and perturbed problems, and consequently small errors on the difference in tally results. For each perturbation, a perturbed weight parameter is introduced. Starting at unity, the value of this parameter is altered whenever a particle is transported or undergoes an interaction, by multiplying it by the ratio of the probability of the transport or interaction process occurring in the perturbed problem to the probability of it occurring in the original problem. Whenever a tally is made, the tally quantity for the perturbed problem is equal to the tally for the unperturbed problem multiplied by the perturbed weight.

The perturbed weight multiplying factors arising during particle transport can be estimated by Taylor expansion [5], also known as the differential approach, or by the multiple-estimation method [6], also known as correlated sampling. For the rest of this paper, I use the correlated sampling approach, as the perturbation estimates are correct to all orders and the calculations are considerably simpler, particularly compared to the evaluation of second or higher order terms in the Taylor expansion approach. The main advantage of the differential technique is held to be that as derivatives are explicitly calculated, one can perform *a posteriori* sensitivity calculations [8]. I argue however, that the effects of multiple perturbations can be calculated in a single run at low computational cost using the correlated sampling approach and *a posteriori* results for perturbations of intermediate magnitude estimated using interpolation.

In comparison to the problem of perturbations to materials and cross-sections, methods for handling variations in problem geometry have received much less attention. Takahashi [9] describes a scheme that allows the effects of simple geometric variations to be estimated. Shuttleworth [10] discusses the application of this technique in the MCBEND Monte Carlo code. Sitaraman [11] introduces a method for calculating sensitivity to surface changes using first-order perturbation theory, and describes the incorporation of this method into a customised version of the MCNP code.

In this paper, I discuss how the perturbation approach can be extended to allow nearly arbitrary variations to be made to the materials, geometry and radiation sources in a simulation. The motivation for this work is to further the use of the Monte Carlo method for designing and optimising X-ray and neutron-based analysis instruments. The availability of what is essentially derivative information (a Hessian matrix) – in other words, changes in tally quantities with respect to parameterised changes in instrument design – serves two purposes. Firstly, it provides some direction on how a design might be modified to improve instrument accuracy, for example by making changes that reduce the statistical error of an analysis. Secondly, it provides information about the size of instrumental systematic, or non-statistical errors, arising, for example, due to an inaccurately positioned sample, or variation in the output of the radiation source.

The remainder of this paper is structured as follows.

Section 2 revisits the material perturbation technique and demonstrated how, in combination with the Woodcock tracking algorithm, it can be extended to calculate the effects of changing the size, shape or position of geometric regions. The more straightforward extension of the material perturbation technique to variations in the energy or angular distributions of particles emitted by the source is also described.

Section 3 discusses the handling of common types of Monte Carlo tallies in perturbation problems.

Section 4 describes the process by which the rather complex specification of simulations including one or more perturbations can be produced automatically from the descriptions of the original problem and its variations. An important consideration is what I term the 'super-set' problem, namely ensuring that all process that can occur at any point in any of the perturbed variations have non-zero probability of occurring at that point in the primary simulation.

Section 5 presents results for an example problem, namely a simple model of an X-ray fluorescence analyser incorporating an X-ray tube and silicon detector. The perturbation approach is used to estimate the errors resulting from movement of the sample or detector, a change in sample density and composition and a fluctuation in the operating voltage of the X-ray tube. Some comments on the efficiency of the perturbation approach are also presented.

Section 6 outlines the conclusions arising from this work.

2. The perturbation method

There is a close correspondence between the perturbation method and approaches for statistical variance reduction in Monte Carlo transport. Both rely on the introduction of a numerical factor termed 'weight' to allow particle properties to be sampled from non-physical or so-called non-analogue distributions.

Suppose that the properties $\mathbf{x}$ (for example, energy and direction) for a particle produced by a radiation source or in collision should be sampled from a physical distribution $p(\mathbf{x})$. We can sample the properties from an alternative distribution $q(\mathbf{x})$ provided that we multiply the particle's current weight by a factor

$$r = p(\mathbf{x})/q(\mathbf{x}). \tag{1}$$

Particle weights start at unity and may increase or decrease as the transport of the particle proceeds. Whenever a particle makes a tally contribution, that contribution must be multiplied by a statistical weight factor. For a simple, Boltzmann tally which does not depend on inter-particle coincidences, the statistical weight factor is given simply by the particle weight. For non-Boltzmann tallies where coincidences are important, the calculation of the statistical weight factor is more complicated. This is discussed in more detail in Section 4.

When playing variance reduction games, the user attempts to find non-physical distributions that result in a reduction in the statistical error achievable in a given amount of computer time. For example, these games may involve randomly terminating particles that are deemed unlikely to make a tally contribution, whilst increasing the weight of any survivors to ensure the process is unbiased.

The correspondence with the perturbation method is straightforward. Suppose that we sample all particle properties from the distributions corresponding to the normal, unperturbed problem. Tally outputs are recorded with a unit weight factor, and the tally results correspond to those expected for the unperturbed problem. In parallel however, we can calculate a perturbed weight factor, given by

$$r_p = p'(\mathbf{x})/p(\mathbf{x}) \tag{2}$$

where $p'(\mathbf{x})$ and $p(\mathbf{x})$ are the probabilities of a given interaction occurring in the perturbed and unperturbed problems respectively.