ANL-75-2
NEA-CRP-L-118

ANL-75-2
NEA-CRP-L-118

INL Technical Library
072256

PROCEEDINGS OF THE NEACRP MEETING OF A MONTE CARLO STUDY GROUP

July 1-3, 1974

Argonne, Illinois, U.S.A.

PLEASE RETURN TO
MFC BRANCH LIBRARY

U of C-AUA-USERDA

ARGONNE NATIONAL LABORATORY, ARGONNE, ILLINOIS

Prepared for the U. S. ENERGY RESEARCH
AND DEVELOPMENT ADMINISTRATION
under Contract W-31-109-Eng-38

The facilities of Argonne National Laboratory are owned by the United States Government. Under the terms of a contract (W-31-109-Eng-38) between the U. S. Energy Research and Development Administration, Argonne Universities Association and The University of Chicago, the University employs the staff and operates the Laboratory in accordance with policies and programs formulated, approved and reviewed by the Association.

MEMBERS OF ARGONNE UNIVERSITIES ASSOCIATION

The University of Arizona
Carnegie-Mellon University
Case Western Reserve University
The University of Chicago
University of Cincinnati
Illinois Institute of Technology
University of Illinois
Indiana University
Iowa State University
The University of Iowa
Kansas State University
The University of Kansas
Loyola University
Marquette University
Michigan State University
The University of Michigan
University of Minnesota
University of Missouri
Northwestern University
University of Notre Dame
The Ohio State University
Ohio University
The Pennsylvania State University
Purdue University
Saint Louis University
Southern Illinois University
The University of Texas at Austin
Washington University
Wayne State University
The University of Wisconsin

NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Energy Research and Development Administration, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately-owned rights. Mention of commercial products, their manufacturers, or their suppliers in this publication does not imply or connote approval or disapproval of the product by Argonne National Laboratory or the U. S. Energy Research and Development Administration.

Distribution
Category: Special

ANL-75-2
NEA-CRP-L-118

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

PROCEEDINGS OF THE NEACRP MEETING OF A MONTE CARLO STUDY GROUP

July 1-3, 1974
Argonne, Illinois, U.S.A.

Sponsored by

The Nuclear Energy Agency Committee on Reactor Physics

Hosted by

Argonne National Laboratory

Chairman

Ely M. Gelbard

Compiled by

Cyrilla Hytry

Arrangements

Cyrilla Hytry
Miriam Holden
Joyce Steyskal

PREFACE

On July 1-3, 1974, a conference of Monte Carlo specialists was held at Argonne National Laboratory under the auspices of the Nuclear Energy Agency Committee on Reactor Physics (NEACRP). This was one of a continuing series of NEACRP working meetings convoked to expedite the exchange of views and information among specialists in many fields. All papers presented at the Meeting (and reproduced in these transactions) dealt with the application of Monte Carlo methods to neutron transport problems. Since the topic of the Meeting was already so narrowly defined, it seemed undesirable to attempt to organize each session around some still more sharply delimited subtopic. The reader will note, however, that the papers presented here do fall, generally, into two broad categories. Most of the papers read at the first two sessions were concerned with the implementation of Monte Carlo methods in certain specified codes. Papers in the last three sessions, for the most part, focussed on the methods themselves.

But the presentation of technical papers was not the sole or main purpose of this meeting. All information in these papers could, of course, have been disseminated in technical journals or company reports. A meeting, on the other hand, offers unique opportunities for discussion, evaluation, and speculation. For this reason the meeting schedule was arranged to allow ample time for discussion after each paper, and (during the closing session) for a period of general discussion after the formal presentation of papers had been completed. All the discussion (which was recorded and transcribed) has been reproduced here in some form. It has been extensively edited and, in fact, at certain points, is presented only in summary. Nevertheless we hope that the participants will find that the edited transcripts incorporated in these transactions are substantially correct.

We are greatly indebted to Dr. J. Royen and the other members of the staff of the NEACRP for the work they have done to make this meeting possible. We thank Mrs. Mirian Holden and Mrs. Joyce Stejskal of the Argonne Conference Planning and Arrangements Office for taking charge of the local meeting arrangements. Finally, it is a pleasure to acknowledge the assistance of Miss Cyrilla Hytry, who helped to organize these transactions, to deal with meeting correspondence, and to prepare the transcripts of comments and discussions.

Ely M. Gelbard
Chairman

CONTENTS

Opening Remarks - *C. E. Till (ANL)* 1

Session I

"Efficient Data Management Techniques Implemented in the Karlsruhe Monte Carlo Code KAMCCO," *G. Arnecke, H. Borgwaldt, V. Brandle, and M. Lalovic (KFZ-Karlsruhe, Germany)* 5

"The General Monte Carlo Code MONK," *J. G. Moore (UKAEA-Chulcheth, England)* 23

"Monte Carlo Code Development in Los Alamos," *L. L. Carter, E. D. Cashwell, C. J. Everett, C. A. Forest, R. G. Schrandt, W. M. Taylor, W. L. Thompson, and G. D. Turner (LASL, USA)* 53

Session II

"A Discrete Approach to Complex Planar Geometries," *E. Cupini and A. De Matteis (CNEN-Bologna, Italy)* 69

"The Organization of Cross-Section Data in the Monte Carlo Code SPARTAN," *R. C. Bending (CEGB-Berkeley, England)* 75

"The Geometry System Used in the General Monte Carlo Transport Code SPARTAN," *R. C. Bending and P. G. Easter* 95

"On Random Number Generation," *R. R. Coveyou (Union Carbide-Oak Ridge, USA)* 119

Session III

"The Bell Transform," *C. J. Everett and E. D. Cashwell (LASL, USA)* ... 133

"Coarse-Mesh Rebalancing Acceleration for Eigenvalue Problems," *T. Asaoka, Y. Nakahara, S. Miyasaka, K. Horikami, T. Suzuki, T. Nishida, Y. Taji, K. Tsuchihashi, H. Grotoh, Y. Seki, M. Nakagawa, and J. Hirota (JAERI-Ibaraki, Japan)* 149

"Monte Carlo Eigenfunction Strategies and Uncertainties," *R. C. Gast and N. R. Candelore (W-BAPL, USA)* 161

"Monte Carlo Work at Argonne National Laboratory," *E. M. Gelbard and R. E. Prael (ANL, USA)* 201

Session IV

"FOCUS — A Non-Multigroup Adjoint Monte Carlo Code with Improved Variance Reduction," *J. E. Hoogenboom (Delft Tech. U., Netherlands)* .. 243

"A Monte Carlo Technique for Local Perturbations in Multiplying Systems," *W. Bernnat (U. of Stuttgart, Germany)* 261

"Bounded Estimation of Flux-at-a-Point for One or More Detectors," *H. A. Steinberg (MAGI, USA)* 281

Session V

"Direction-Dependent Exponential Biassing," *R. C. Bending (CEGB-Berkeley, England)* 295

"Monte Carlo Methods for Shield Design Calculations," *M. J. Grimstone (UKEA-Winfrith, England)* 315

"MCT, A Monte Carlo Code for Time-Dependent Neutron Thermalization Problems," *E. Cupini and R. Simonini (CNEN-Bologna, Italy)* 339

Appendix

General Discussion on Monte Carlo Methods 353

Participants 362

OPENING REMARKS

C. E. Till, Director

Applied Physics Division
Argonne National Laboratory

On the behalf of the United States Atomic Energy Commission, and on behalf of Argonne National Laboratory as your host, I would like to welcome you all and to open this meeting of specialists on Monte Carlo methods. In particular, I would like to welcome our overseas guests — we are happy to see you here. It is the Laboratory's honor to act as host for what I am sure will be a most interesting meeting. Meetings of this kind, where information is exchanged between experts in the field, are possibly the most valuable of the types of technical meetings that we hold. Such meetings I think are important, not only because of their technical content, but also because they represent a continuation of the tradition of exchange and cooperation between our nations. I believe that it is important that we maintain such exchanges and indeed actively encourage them or inertia tends to act to reduce them, to the detriment of the programs of all of us.

I wish you all a most successful meeting. I know that the papers have been selected with care, and I am sure that your discussions will be lively and interesting. I would like to close by wishing you all a very pleasant stay in the Chicago area and here at the Laboratory.

SESSION I

EFFICIENT DATA MANAGEMENT TECHNIQUES IMPLEMENTED IN THE KARLSRUHE MONTE CARLO CODE KAMCCO

G. Arnecke, H. Borgwaldt, V. Brandl and M. Lalovic

Institut für Neutronenphysik und Reaktortechnik
Karnforschungszentrum Karlsruhe, Germany

* * *

Presented by H. Borgwaldt

EFFICIENT DATA MANAGEMENT TECHNIQUES IMPLEMENTED IN THE KARLSRUHE MONTE CARLO CODE KAMCCO

G. Arnecke, H. Borgwaldt, V. Brandl, M. Lalovic

Institut für Neutronenphysik und Reaktortechnik
Kernforschungszentrum Karlsruhe, Germany

ABSTRACT

The Karlsruhe Monte Carlo Code KAMCCO is a forward neutron transport code with an eigenfunction and a fixed source option, including time-dependence. A continuous energy model is combined with a detailed representation of neutron cross sections, based on linear interpolation, Breit-Wigner resonances and probability tables. All input is processed into densely packed, dynamically addressed parameter fields and networks of pointers (addresses). Estimation routines are decoupled from random walk and analyze a storage region with sample records. This technique leads to fast execution with moderate storage requirements and without any I/O-operations except in the input and output stages.

INTRODUCTION

The general purpose fast neutron Monte Carlo (MC) code KAMCCO has been developed, starting from previous work by U. Möller [1], as an instrument for (a) analyzing fast neutron experiments, including time-dependent pulsed neutron problems, (b) checking nuclear data available in the format of our KEDAK library [2], and (c) checking other techniques of fast reactor neutronic calculations. No stress was put on shielding and similar deep penetration problems. Following these aims, the natural approach chosen was to use a good model of the neutron physics involved, even at the cost of some increase in computation time, and to avoid many of the approximations usually accepted in reactor computation. Thus, KAMCCO fully includes the time-dependence of neutron transport and uses a continuous energy treatment instead of a multigroup model. In addition to a good physical model we required that the code should possess considerable freedom in problem specification (i.e. input and output options) and be easily adapted to improved physical models and/or unforeseen neutronic applications. All these requirements ask for modular programming and for refined data management techniques, as otherwise the required flexibility will involve an excessive overhead for clerical tasks. The programming language is FORTRAN IV, for a group of functional routines ASSEMBLER-routines may be substituted.

GENERAL FEATURES OF KAMCCO

In our code neutrons carry the energy E, the time t and a weight W as continuous variables. Russian roulette and non-analogue treatment of capture can be used to an extent specified by input. At each collision a special technique is used to reduce the variance normally connected with (n,2n) and fission processes. For this purpose the reaction type is selected from a suitably modified set of probabilities, so that the resulting neutron weight becomes independent of the isotope and reaction type involved. - KAMCCO has 2 main options:
The standard option 1 is mostly used for k_{eff} calculations, a source iteration scheme (i.e. generation cycle) is performed. A pool of fission neutron coordinates (cf. subpool A3 in diagram 1) is established and repetitively updated with coordinates of new fission neutrons. These coordinates are obtained as imaginary marks (positions) on the $(W\nu\Sigma_f)$-weighted cumulated neutron flight paths, with an exponential distribution of distances between these marks [3]. The regeneration process for this source pool can be made virtually free of bias, when the pool size is sufficiently large and some superimposed randomization is applied.

KAMCCO option 2 is a census time scheme, mostly used for the layout and analysis of neutron physics experiments. Neutrons with predetermined source coordinates are started at time t=0. These neutrons, and also all daughters from fission and (n,2n)-reactions, are pursued up to a certain time limit t_{max}. Additional variance reducing techniques operative for this option are a combination of truncated track length distributions at outer boundaries, ELP, and its improved modification MELP [4]. Weight limits for Russian roulette, ELP/MELP and splitting float with time t.

The geometry package allows to describe assemblies with all regions bounded by first and second order surfaces. Inner void regions are admitted, outer boundary conditions are total absorption (vacuum) or reflection on planes.

In order to retain the complete model of neutron physics, neutron collisions are processed on a per isotope basis. In elastic collisions first order anisotropy is introduced in the center-of-mass system. In this system the scattering angle is chosen uniformly from a cone, thereafter a transformation to the lab system is performed. Inelastic scattering at high energies is processed via the evaporation model, at lower energies via the excitation of discrete levels, that are represented in a way very similar to that of global cross sections (see next section). For fission a Cranberg spectrum is implemented, for the secondary distribution of (n,2n) neutrons an ad hoc spectrum has been introduced.

THE REPRESENTATION OF NUCLEAR DATA

Part of the KAMCCO input (cf.diagram 1) is a nuclear data set in a special format. Such data sets are prepared off-line, in principle once for each isotope, by running two codes, DASU and DISTRESS, and a third, DACONT, for merging, editing and listing their outputs. In a MC job the data set is handled as easily as any multigroup data set.

The code DASU [5] operates on the nuclear data file KEDAK and performs two tasks: (a) Point cross sections from the KEDAK file are linearly fitted on

optimally chosen energy grids; (b) other data, e.g. resonance parameters, are converted, so that during the MC run only a minimum of computer operations need be executed on such parameters. - DISTRESS is a special MC code that computes tabulated probability distributions for cross sections in the region of narrowly spaced resonances, for selected energies and 3 temperatures. DISTRESS uses a ladder method for generating, via the single level Breit-Wigner formalism, a large sample of cross sections, which is sorted and partitioned with respect to the size of the total cross section.

During the MC run, the neutron cross sections and other nuclear data are processed pointwise, on a per isotope basis, composite macroscopic cross sections by simple superposition. Main technique is table lookup with linear interpolation in energy. Maximum economy in computer storage is achieved by using tables with energy grids individually and optimally fixed for each isotope, reaction type and broad energy region.

For the representation of nuclear data 4 broad energy regions are distinguished. For a single isotope several regions or reaction types may be missing, such information and the boundaries of the broad energy regions are individually fixed for an isotope and included in a short block of control data.

In the lowest energy region elastic scattering is represented by interpolation coefficients A_i , B_i which allow a good representation

$$\sigma_s(E) = B_i + A_i E \tag{1}$$

with the energy E in an interval (E_i, E_{i-1}). The data actually tabulated in our data sets are sequences of triplets (E_i, A_i, B_i) with $E_{i+1} < E_i$. For absorption (and also for fission) a modified piecewise representation

$$\sigma_a(E) = (B_i' + A_i' E)/\sqrt{E}, \; E \varepsilon (E_i', E_{i-1}') \; , \tag{2}$$

is used for obvious reasons with tabulated (E_i', A_i', B_i').

For the next higher energy region, the resolved resonance region, our data sets contain converted resonance data. In this region, point cross sections are computed during the MC run, in the Breit-Wigner single level formalism, from a number of contributing resonances above and below the reference energy. This number is limited and fixed in advance. Doppler-broadening is optional, also this option must be fixed when the data set is generated.

In the unresolved resonance region, or more generally in the region of narrowly spaced, partially overlapping resonances, where we can trust the narrow resonance approximation, it is advantageous to use off-line generated probablility tables, instead of computing the contributions of many resonances in-line. In our scheme we can easily make use of isotope-based probability tables, in view of the superposition principle employed. In our data sets these tables are organized as a sequence of sub-tables for a number of reference energies. One such sub-table, for a specific reference energy, consists of a number, 3 to 10, of septets (decuplets) of data for non-fissile (fissile) isotopes. A septet (decuplet) of data starts with a value, between 0. and 1., of the cumulative probability distribution function followed by 6 (9) coefficients, which allow the necessary interpolation of

cross section values in energy and temperature.

In the highest energy region all cross sections, for elastic and inelastic scattering, absorption, fission and (n,2n) reactions, are represented by sequences of mesh points E_i and coefficients A_i, B_i used for linear interpolation as in Eq. (1). When necessary, small fixed groups of data, needed for generating the secondary spectra of fission or (n,n') neutrons etc., are included, contingent to the respective cross section data. For the sake of completeness, we mention a tabulation, in energy, of the minimum and maximum scattering cosines in the center-of-mass system and a representation of the excitation cross sections of discrete inelastic levels, along the lines of Eq. (1) with minor organizational refinements.

The data structure, explained here only in outline, is an ambitious scheme, in that we can represent data in as much detail as desired. Smooth data are given by a minimum of parameters. In addition the scheme is flexible, as we have been and will be able to introduce, locally, necessary improvements. On the other hand, the efficiency of such procedure, with multiple table lookup and interpolation, relies largely on good coding and the optimized layout of these data during the execution of a MC job.

OPTIONS FOR THE ESTIMATION PROCEDURE

In a MC code with a multigroup cross section library estimates of the neutron flux or the collision density obtained in regions of phase space are primary quantities, from which any reaction rate is derived. This is no longer true with our scheme of nuclear data. Effective cross sections, averaged over energy intervals, are not part of the input but may instead be obtained as an estimated quantity. This becomes obvious when we think of energy regions with resonances or threshold cross sections, where effective cross sections are based on a known weighting spectrum which does not enter our MC scheme working on a continous energy model.

The user is given a free and wide choice of options for the estimation procedure as part of his problem specification. All estimates are, of course, obtained with their standard errors.

The user may choose a collision or a track length estimator or an optimum combination. He may choose global estimates (multiplication, gain, loss, leakage, time to death/birth) and/or estimates of reaction rates integrated over regions in phase space. Here, the user specifies, for the purpose of editing only, an energy grid completely independent of any grid used in the nuclear data set.

Estimates may be specified in complete detail or as energy spectra in selected spatial regions or as distributions over all spatial regions. The random walk information can be analyzed with respect to regions, energy intervals, isotopes and reaction types (including flux). Furthermore, the desired estimates can be specified as directly estimated combinations of simpler quantities, e.g. by specifying: capture plus fission, in 3 identified isotopes, as a spectrum, summed over 5 identified regions. The code is used to a large extent for the analysis of time-dependent problems, i.e. pulsed neutron experiments in subcritical or non-multiplying assemblies [6]. Therefore, the extention of the estimation procedure to multiple estimates of the types just explained, for a series of time intervals, has been implemented.

Finally, we added another valuable feature that allows us to estimate directly the ratios of reaction rates, which are often measured in critical assemblies.

It is obvious that this flexibility in specifying estimation options introduces an amount of complexity, which can be handled only by programming that decouples almost completely the routines for the execution of the random walk process from the estimation routines.

OVERALL PROGRAM AND DATA STRUCTURE

Diagram 1 (see next page) demonstrates in a simplified form the data flow and organization of KAMCCO. The blocks R2 and E2 of subroutines execute the random walk and estimation, respectively. R1 and E1 are additional blocks of routines, the associated input processors. A rudimentary main program, which is not shown, calls these blocks in due sequence. The data pool B is loaded by block E1 with processed input data and address fields, which are required for the estimation routines.

Central part is the data pool A, at this time a COMMON block of 28 000 full words, containing all data needed by the random walk routines alone or shared with the estimation routines. It is subdivided into a 210-words subpool A1 with fixed-address data, largely for coordination, and 3 subpools, A2 to A4, that use dynamic addressing. Between these subpools no gaps are left. Subpools A2 and A3 are used only for the random walk. Subpool A2 contains, in densely packed form, all random walk input data, i.e. nuclear data and specifications for surfaces, boundaries, material compositions, and spatial regions. Program block R1 loads these data without gaps into subpool A2, converts and structures them into a complete network of parameters and pointers to their locations. During the random walk these pointers give fast access to any data needed.

Subpool A3 is the space allocated to the relevant data of source or split neutrons, its size is determined by input. Subpool A4 serves mainly for the communication of relevant information on the random walk process to the estimation routines. Such data associated with random walk paths and events comprise: identifiers of the spatial region and the energy interval, neutron weight, velocity (energy) , collision and escape times, mean free path, lengths of actual and virtual flight paths, and the associated macroscopic cross sections of all isotopes. This information is stored during the random walk computations and asynchronously analyzed by the estimation routines. The storing format is an 8-word record starting with a type identifier. Actually, types 1, 3, 5 signal a source event, a collision and a restart of a split neutron. The augmented type 2 communicates cross sections. Type 4 contains the flight path data and may be augmented for ELP/MELP. Types 8 and 9 are pure control records for switching between the program blocks.

This scheme decouples random walk and estimation almost completely without using external data carriers, that would involve costly, slow and, to some extent, unreliable I/O-operations.

Diagram 1

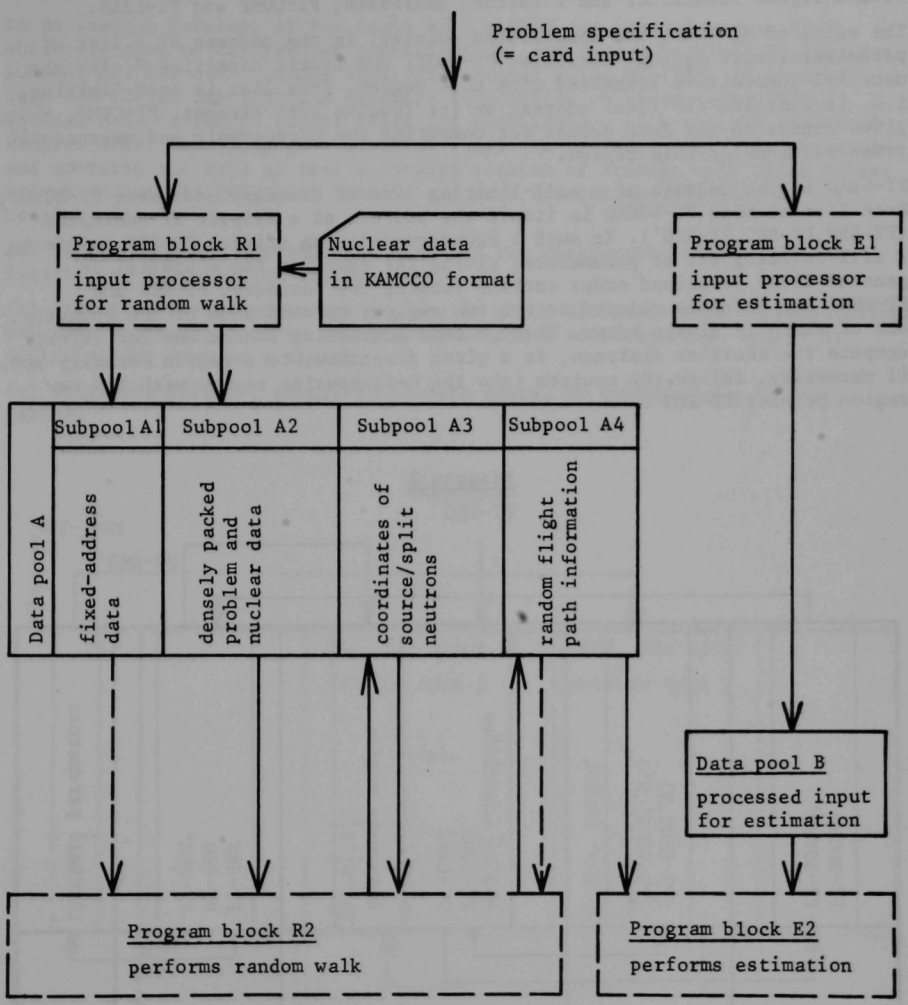
Problem specification
(= card input)
Program block R1
input processor
for random walk
Nuclear data
in KAMCCO format
Program block E1
input processor
for estimation
Data pool A
Subpool A1
fixed-address
data
Subpool A2
densely packed
problem and
nuclear data
Subpool A3
coordinates of
source/split
neutrons
Subpool A4
random flight
path information
Data pool B
processed input
for estimation
Program block R2
performs random walk
Program block E2
performs estimation

ADDRESSING SCHEME FOR THE GEOMETRY DATA

In diagram 2 we describe, in a simplified form, the addressing scheme employed in the random walk of a neutron. Basic pointer is PT-REG that contains the address of an 11-word block of data referring to the region in which the neutron is found. With fixed offset values of 0, 1, 2 to PT-REG we find a region identifier and 2 further addresses, PT-COMP and PT-LBAD.

The value of PT-COMP, the composition pointer, is the address of a list of parameters, viz. isotope addresses $PT\text{-}ISOT_i$ and atomic densities N_i for the material composition connected with this region. This list is self-limiting, i.e. it contains its final address as its leading list element. PT-COMP, thus, gives access to all data needed for computing the microscopic and macroscopic cross sections in this region.

PT-LBAD is the address of a self-limiting list of boundary addresses $PT\text{-}BOUND_i$. Each list element PT-BOUND is itself the address of a triplet of addresses (PT-SRF,PT-REG,PT-REG'). In such a boundary-defining triplet PT-SRF points to a self-defining set of parameters, giving the equation for a special or general first or second order surface forming the boundary. PT-REG and PT-REG' are pointers related to the two regions on each side of the boundary, one of which is always known. Through this addressing scheme one can quickly compute the shortest distance, in a given direction, to a region boundary and, if necessary, follow the neutron into the neighbouring region with the new region pointer PT-REG'.

Diagram 2

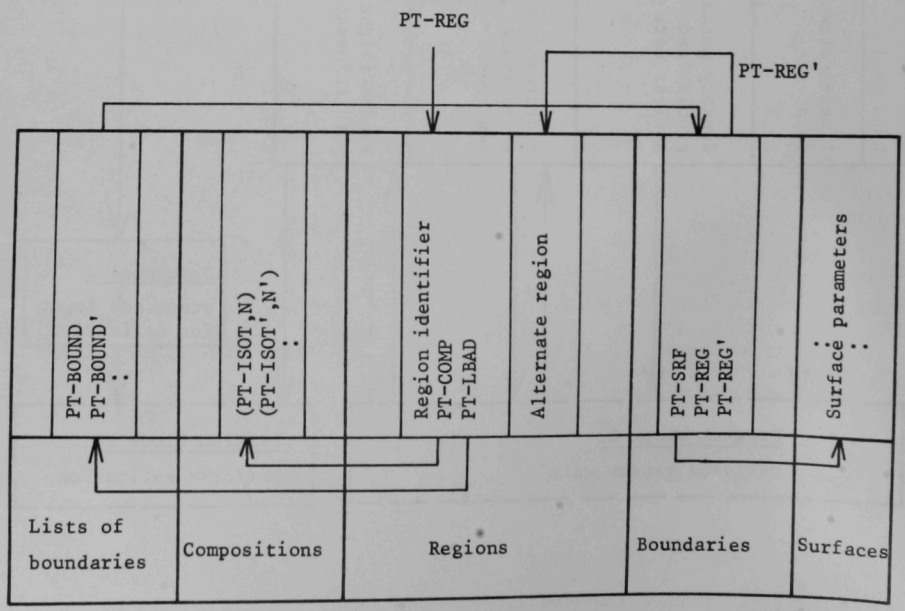

ADDRESSING SCHEME FOR THE NUCLEAR DATA

A step in the random walk process, that is very often executed, is the computation of neutron cross sections and similar nuclear data. In the preceding section we explained the addressing scheme

PT-REG → PT-COMP → PT-ISOT, from a region pointer via a composition pointer

to an isotope pointer. At the start of a KAMCCO run those isotopes $ISOT_i$, which are part of the problem specification, are loaded by block R1 into data subpool A2. For each isotope $ISOT_i$ the pointer value $PT\text{-}ISOT_i$ is inserted into all tables describing a material composition which contains $ISOT_i$. Nuclear data in KAMCCO format start for each isotope with a fixed amount of control data. When program block R1 loads the isotope $ISOT_i$, it structures and converts the data so that a 40-word section of storage with fixed offset values versus $PT\text{-}ISOT_i$ will contain all such control specifications. They consist of:

(a) simple tags and descriptors, e.g. for the distinction of fissile, non-fissile and fast-fission isotopes etc.,

(b) limits of the broad energy regions in which different data representations CRS_j are applicable, or in which specific cross sections, e.g. (n,2n) or fast-fission, must be computed,

(c) addresses $PT\text{-}CRST_j$ pointing, in most cases, to intermediate address tables.

Diagram 3

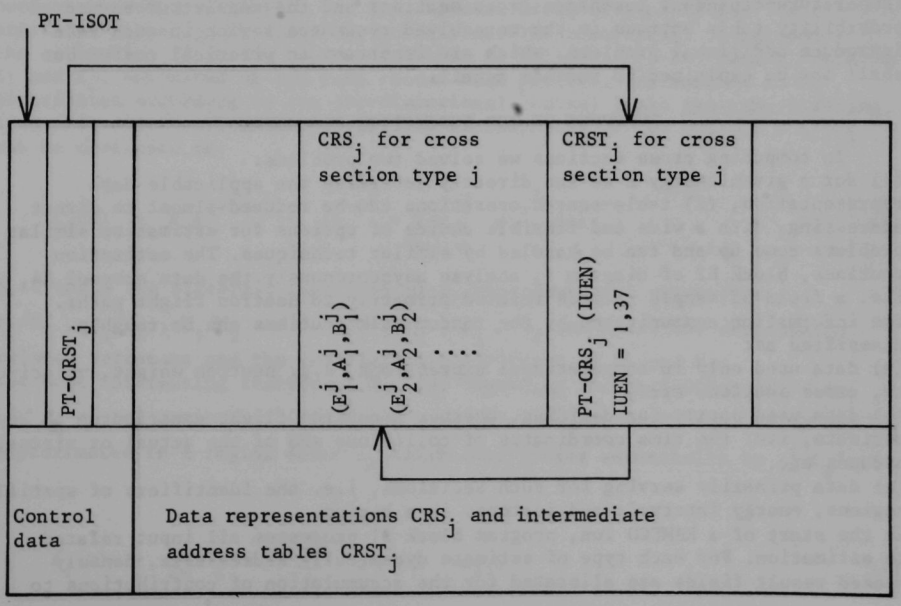

We recall, from Eq. (1), that the data for a representation CRS_j are, in the regular case, a sequence of triplets $(E_\ell^j, A_\ell^j, B_\ell^j)$ of energy grid points E_ℓ^j (in descending order) and interpolation coefficients A_ℓ^j, B_ℓ^j. When E is the neutron energy, table-lookup is performed to identify that triplet for which $E \varepsilon (E_\ell^j, E_{\ell-1}^j)$.

The intermediate address tables mentioned considerably reduce the time spent for table-lookup operations. Instead of starting at the top of the table intermediate addressing allows to start the search in the immediate vicinity of the unknown E_ℓ^j-entry. The intermediate address table for a representation CRS_j consists of a list of 37 addresses $PT\text{-}CRS_j$ (IUEN), IUEN=1,...,37. These tables are generated and inserted into the nuclear data by the program block R1 in loading the isotope. They are based on an internal, universal energy grid of 37 intervals between .01 eV and 15. MeV, with 4 intervals per decade. After each redefinition of neutron energy, e.g. by a collision, the interval index IUEN in this energy grid is determined. The intermediate address table, then, contains in its relative position IUEN that address $PT\text{-}CRS_j$(IUEN) which can be used to optimally initialize the table-search in the representation CRS_j, assuming the neutron has an interval index IUEN (cf. diagram 3).

This scheme is used, sometimes with minor modifications, whereever table-look-up operations must be optimized. Its benefit is achieved by some additional work in block R1, which is called only once and, thus, is comparable to a sequence of operations taken out from a loop.

The subroutines for computing cross sections also perform some additional book-keeping to save time. Computed microscopic cross sections are saved so that they can be re-used, when a neutron path crosses several regions. Temperature-dependent resonance cross sections and the consistent re-use of probability table entries in the unresolved resonance region in such case introduce additional problems, which are important in practical coding but shall not be explained in further detail.

FEATURES OF THE ESTIMATION PROCEDURE

In computing cross sections we solved two problems:

(1) for a given energy E we can directly determine the applicable data representation, (2) table-search operations can be reduced almost to direct addressing. With a wide and flexible choice of options for estimation similar problems come up and can be handled by similar techniques. The estimation routines, block E2 of diagram 1, analyze asynchronously the data subpool A4, i.e. a field of sample records related primarily to neutron flight paths.

The information communicated by the random walk routines can be roughly classified as:

(a) data used only in computational expressions, e.g. neutron weight, velocity, cross sections etc.,

(b) data used partly for deciding, whether a neutron flight contributes to any estimate, i.e. the time coordinates of collisions and of the actual or virtual escapes etc.

(c) data primarily serving for such decisions, i.e. the identifiers of spatial regions, energy intervals and isotopes in a region.

At the start of a KAMCCO run, program block E1 processes all input related to estimation. For each type of estimate dynamically addressable, densely packed result fields are allocated for the accumulation of contributions to

estimates, variances and covariances. Composite estimates are directly estimated and not only treated as an editing task. The addresses of these allocations are collected in address tables. All specifications for estimation are processed as a large number of extended tables which are either address tables or decision tables (containing the values 0,1). Both types are used, as one type requires more storage, the other more computing. These tabulations enable a fast decision regarding each estimate, as to whether a neutron flight with specified region and energy interval identifiers contributes to that estimate. In the positive case, all addressing of the data to be processed in the data subpool A4 or the result fields is performed through these pre-processed address tables.

Decisions performed on time coordinates cannot be accelerated to a large extent, in some cases one logical decision can be used to skip subsequent neutron flight paths. As mentioned above, one may obtain the same group of estimates for several time intervals. Originally, we proceeded without any sophistication by multiply scanning the subpool A4 and multiple swapping of data between the result fields and a save area. Recently, multiple time interval estimation has been fully incorporated into the addressing and decision tables. As in this case the sample records in A4 are scanned only once, a significant acceleration of the estimation procedure is obtained, when 2 or more time intervals are required.

A special technique has been developed for estimating ratios of more elementary estimates. Such ratios are needed for the analysis of experiments and for the determination of effective cross sections. MC estimates become problematic when denominator and numerator have significant statistical errors. Quite often, however, these quantities are positively correlated to some degree and their errors will cancel partly in the ratios formed. What we still need is an exact expression for the expected error of such ratios and a practical method to handle these quantities. The theoretical background of the estimation of ratios can be described as follows [7]: Two random variables X_1 and X_2, estimated in the same random walk process, are assumed to be distributed according to the two-dimensional central limit theorem. Starting from this distribution the probability distribution of the ratio $Z = X_1/X_2$ can be developed as:

$$f_o(z) = K\, e^{(w^2-c)} \cdot \left(e^{-w^2} + 2\, w \int_o^w e^{-x^2}\, dx\right) \qquad (3)$$

K, c and w are algebraic functions of z, dependent on 4 parameters, σ_1/M_1, σ_2/M_2, M_1/M_2 and ρ. M_1, M_2, σ_1, σ_2 and ρ are the expectation values, variances and the correlation coefficient of X_1 and X_2.
The most interesting feature of Eq. (3) regarding the estimation of Z is:
Let z_M be the maximum of $f_o(z)$, and $\sigma_1/M_1 << 1$, $\sigma_2/M_2 << 1$. Then $f_o(z)$ can be approximated in a region about z_M which contributes essentially to $\int f_o(z)dz$, by

$$f_1(z) = \frac{1}{2\pi}\,\frac{1}{\Sigma}\, e^{-\left(\frac{z-Q}{\Sigma}\right)^2}, \qquad \text{with} \qquad (4)$$

$$Q = M_1/M_2 \quad \text{and} \quad \Sigma = Q \sqrt{\left(\frac{\sigma_1}{M_1}\right)^2 - 2\rho\frac{\sigma_1\sigma_2}{M_1M_2} + \left(\frac{\sigma_2}{M_2}\right)^2} \tag{5}$$

In contrast to $f_o(z)$, the normal distribution $f_1(z)$ is analytically integrable and the statistical error of Z can therefore easily be estimated: At the end of each history one has to store additionally the mean value of the cross term $X_1 \cdot X_2$, which needs some more storage but a negligeable amount of computer time.

PERFORMANCE OF THE IMPLEMENTED SCHEME

Obviously, many different criteria may be taken as a basis of judgement on the performance of a proposed method. It is our belief, that in MC code development the following aspects are of special importance. A chosen technique of programme linking and data handling must not impose a rigid frame of its own. The resulting restrictions may create difficulties, when a user-supplied problem has an unexpected size, or with necessary new developments (e.g. incorporation of better physics models, improved estimation procedures or new eigenvalue strategies). Our code is run on an IBM/370-168 with 240-300 k bytes of storage, which is quite sufficient but could be expanded immediately. The implemented scheme has proved sufficiently flexible for code development and is virtually free of limitations of problem size.

A second requirement is, that partial tasks of a MC job must not require disproportionate shares of the total computer time. With KAMCCO this concerns mainly the techniques for estimation and handling cross sections. Estimation time is highly problem-dependent, in practice we found 10 and 30 per cent as minimum and maximum values, which we think acceptable.

The performance of cross section handling is best demonstrated by the following short selection of criticality calculations for reactors with different numbers of isotopes and differing neutron spectra.

Reactor	Multiplication k_{eff}	Error δk	Isotopes	Collisions
JEZEBEL	.9824	.0015	2	360,000
POPSY	.9691	.0063	4	490,000
SNEAK 3A1	.9686	.0077	8	310,000

The data set used was a standard set for MC code testing without adaptations to improve k_{eff}. Cross sections are computed after each collision, computations at low energies are more complicated. For each job about 6 minutes of CPU-time was used. As the number of collisions processed is only moderately affected by the numbers of isotopes and type of spectrum, we may conclude that our scheme of representing nuclear data is, indeed, practicable and efficient.

REFERENCES

1. U. MÖLLER, "MONTE, Monte-Carlo-Code zur Lösung der Transportgleichung," Kernforschungszentrum Karlsruhe, Private Communication (1966), U. MÖLLER, "Unterprogramme zur Lösung neutronenphysikalischer Probleme mit Hilfe der Monte-Carlo-Methode," Rep. KFK 298, Kernforschungszentrum Karlsruhe (1965).

2. B. HINKELMANN et al., "Status of the Karlsruhe Evaluated Nuclear Data File KEDAK at June 1970," Rep. KFK 1340 and EANDC(E) 136 "U", Kernforschungszentrum Karlsruhe (1971).

3. H. BORGWALDT and M. LALOVIC, "Untersuchung des Quelliterationsschemas für den Monte-Carlo-Code KAMCCO," in Rep. KFK 1273/4 and EUR 4974d, p. 122-2, Kernforschungszentrum Karlsruhe (1974).

4. H. BORGWALDT, "Comparison of 3 Methods to Control the Leakage of Particles in a Monte Carlo Game," in Proc. New Developments in Reactor Mathematics and Applications, Idaho Falls, USAEC (1971), p. 857, (also Rep.KFK 1391).

5. G. ARNECKE, H. BORGWALDT and M. LALOVIC, "DASU, Programm zur Erstellung von Wirkungsquerschnitten für den Monte-Carlo-Code KAMCCO," Kernforschungszentrum Karlsruhe, Private Communication (1971).

6. F. KAPPLER et al., "Measurements of Fast Neutron Spectra in Iron, Uranium and Sodium-Iron Assemblies," in Proc. Int. Symposium on Fast Reactor Physics, Tokyo, Oct. 1973, Vol. II, p. 883.

7. V. BRANDL, "Die Monte-Carlo Berechnung von Quotienten," Rep. KFK - , Kernforschungszentrum Karlsruhe (to be published).

DISCUSSION

Steinberg: It is not clear to me how your code handles problems where you have such a large number of isotopes that the data won't fit in the core.

Borgwaldt: I don't understand your query.

Steinberg: How do you handle a problem in which you have 20, 30, or 40 isotopes and the total volume of the data is such that the machine won't handle the entire energy range all at once for all the elements?

Borgwaldt: We are handling the complete energy range, and so far we have gotten along with the storage that we have. If the storage is not sufficient, then we can easily expand the code to use twice the amount, or three times the amount that we have now for storing data. On the other hand, there are two comments worth making. I think that it is not wise, in such a code, to include minor isotopes only for the sake of completeness: one can, in practical cases, reduce the number of isotopes to, let us say, 10. Further, with our scheme, we can describe some isotopes, in full detail, using only a few hundred full words because (in these cases) linearly fitting a curve works very nicely, works better than using a multigroup scheme. Other isotopes, for instance ^{238}U or plutonium isotopes, of course use more data. But this flexible scheme of energy grids allows us to describe data with a low number of storage positions in one case and a large number when it is necessary.

Kalos: I think that part of the problem here lies in the fact that you have the advantage over us in using the Karlsruhe data set, where we are at the mercies of ENDF/B, a library which includes a great deal more data which we feel constrained to present in the computer. For example, the full description of anisotropic elastic scattering itself requires a very large amount of data for us when we go from a few electron volts to 16 MeV. I believe that your data are much sparser.

Borgwaldt: No, I wouldn't say that. I think that the point is that if I perform a Monte Carlo calculation, I use a model; if I perform a computation in any multigroup scheme, I also use a model. We have at Karlsruhe, a 280-group data set for multigroup calculations. Usually we take a 26-group set. I take the position that if I create a nuclear data set for Monte Carlo calculations I will also compromise. I will not try to fit every minor deviation of the cross-section curve, but if the Monte Carlo data set has the quality of let us say a 40- or 60-group data set, then it will be sufficient for me. It is my personal point of view that I cannot put stronger demands on the data sets used for a Monte Carlo code than I would put on data used for a multigroup code.

Gelbard: Is it true that part of the reason that you don't have such serious storage problems is that you construct the resonance cross sections at execution time?

Borgwaldt: The resonance cross sections in the low-energy region (which range is seldom used since we are primarily interested in fast neutron calculations) are computed at execution time. Above the region of narrowly spaced resonances we use probability tables.

Gelbard: So that if you were doing a Doppler calculation you would in that case be concerned with the resonance structure, and then would the Doppler actually broaden at execution time?

Borgwaldt: Our probability tables also contain temperature dependence. Our code computes probability tables for different temperatures at specified energies and we can also use these tables for Doppler-broadened cross sections, insofar as the single-level Breit-Wigner formula, the ψ and χ function formalism, is valid.

Gelbard: But in the resolved range, below the unresolved resonance, you do use explicit resonance parameters?

Borgwaldt: Yes.

Gelbard: You do the Doppler broadening at run time?

Borgwaldt: Yes.

Gelbard: Is that a serious problem? When you Doppler broaden at run time, how much computer time is involved?

Borgwaldt: May I make another point here? Sometimes people make a sharp distinction between the unresolved resonance region and the resolved resonance region. I am prepared to use probability tables even in the resolved resonance region in situations where a neutron can jump over two or more resonances in a single collision. I don't stop using probability tables at the resolved resonance region; on the contrary, I try to use the probability tables as long as neutron physics, the collision physics, allows me to say that I am not concerned with what occurs in a single collision in a Monte Carlo but only with what occurs in the sample average. Does that answer your question? For instance, for ^{235}U our probability tables start with 60 or 100 eV and in ^{238}U our probability tables start with 1 keV and not with 2 or 3 keV, a figure that would be based on the physical knowledge of the resonance parameters.

Moore: It is interesting that you used your code to check nuclear data and other methods of calculation. We have been unable to interest reactor physics people in the United Kingdom in these aspects. I can only deduce from this that your variance reduction techniques are very good. The sort of problems that our physics people are interested in require a standard deviation in k_{eff} of the order of 0.1%, and so far we have been unable to achieve this. Can you say something about your success in this field?

Borgwaldt: Yes. Could I have Slide No. 8? We haven't been able to create enthusiasm among reactor physicists of the type you are thinking about, but at Karlsruhe we are doing certain experiments which can only be analyzed by Monte Carlo techniques. Here you see an iron cylinder which is pulsed; it is a pulsed nonmultiplying system. An analysis is done for the time-dependent leakage spectra and one wants to deduce inelastic cross sections from such an experiment. We have a three-dimensional problem with time dependence and all other codes break down.

Summary of Borgwaldt's remaining comments: Monte Carlo analysis of the pulsed cylinder suggested substantial modifications in the iron cross sections,

and particularly in the inelastic cross sections. Agreement between computational and experimental results was very much improved after introduction of the adjusted cross-section sets. Borgwaldt commented that Monte Carlo can do much more than it is called upon to do by the average user.

Moore: Can you give us some figures for the standard deviations you get on k_{eff} for critical systems?

Borgwaldt: In small critical systems like GODIVA and JEZEBEL, we get in 6 min on our 370/168, a deviation of 0.2-0.25%. On even smaller systems we obtain an accuracy of 0.1-0.2%. For typical larger systems, like POPSY-TOPSY and SNEAK assemblies, we obtained 0.5-0.7% in 6 min. But, I think we can improve on that by improved source iteration techniques.

Gelbard: I think probably that the laboratory that has the most experience with the analysis of criticals by Monte Carlo is Bettis; now I don't know if Bob Gast would be willing to discuss this experience, but I would like to ask Bob what kind of running times are generally required to get the kind of accuracy they need, and what kind of usage they have made of Monte Carlo.

Gast: Our experience has been rather gruesome and we have wound up spending many hundreds of hours on Monte Carlo runs. In my talk I would like to go into more details about what our experience has been, but I can't give you specific running times.

Borgwaldt: I might add that we are talking about fast-neutron calculations and I think that the conditions with fast criticals are more favorable than those for thermal criticals, especially if you have low-absorption regions in your thermal criticals. Our codes have broken down completely in a case where we had a D_2O reflector, I must admit; and I warn every user at our center when he tries to use our code on a thermal system. If you have a really thermal region with low absorption, then be careful.

Gelbard: I thought it might be worthwhile coming back to this question of philosophy. I think that there are two different points of view, at least, about Monte Carlo usage. The point of view that I am familiar with is that, since you are going to regard the Monte Carlo code as a standard, you cut as few corners as possible. This means that you not only have to represent the resonances in some detail, but also that you represent the very complicated cross-section structure of materials like iron and aluminum in great detail. There are situations where you are concerned about the accuracy of a multigroup scheme, and the function of the Monte Carlo code in such cases is to check the accuracy. For example, there was one situation that we had at Argonne where neutrons left the core into an aluminum reflector, and then returned from the reflector into the core streaming through the windows in the cross sections of the aluminum. These windows had a strong influence on the core spectrum, causing large peaks in this spectrum. It was for this reason that it was very difficult to define multigroup cross-section sets. For this sort of a Monte Carlo calculation you need as detailed a structure as you can get; so there seems to be many different levels at which Monte Carlo codes operate, and sometimes you are forced into a position where you want all the detail you can get.

Borgwaldt: I agree with you. What we feel is that by using a linear interpolation you are really far better off than by using normal multigroup sets. I agree with you that you may come into a position where it does not pay anymore to do linear fitting, but where you have to use the real point cross sections as they come from NEDF/B. But, just as in some situations you collapse neutron cross sections in a deterministic reactor computation, and in other cases you do not, you can also proceed in this fashion in Monte Carlo. You can create one data set which contains the iron or nickel data in great detail and so uses a large amount of storage, and you can reduce the representation of other data in which you are not so much interested, to some degree. We have not yet come to our upper limit. So far we have not had the problem that we have had to say to any user that we cannot represent his data.

Hoogenboom: I have a question about the anisotropic scattering in your code. As I understood it, you define a μ_{max} and μ_{min}. Is it correct that you choose isotropic scattering angles between the two boundaries?

Borgwaldt: Yes. In the center-of-mass system we find a scattering angle (by a rejection technique) which fits in with these μ_{max} and μ_{min} limits. The μ_{max} and μ_{min} are part of the data set and are fixed in such a fashion that the sum divided by two is the known average μ in this energy region. The scattering takes place in a cone. We feel that for our applications so far this approximation has been adequate, and also this scheme is expandable. We have put it into our code in a form so that if somebody really did come up with problems where anisotropy must be described in more detail, we can do so. So far nobody has come up with such a problem.

Hoogenboom: I don't understand. No matter what kind of anisotropy you have, the μ in the center-of-mass system always lies in the range from -1 to +1; you cannot exclude a certain range of μ values.

Kalos: Remember they are only trying to reproduce the average cosine of scattering in the center-of-mass system. That is the only datum about anisotropic scattering that they are trying to put in.

Moore: Oh, yes, I see.

Borgwaldt: We give only a rough picture of the scattering. The true distribution of the scattering angle in the center-of-mass system is replaced by a rectangular distribution; and, if you have strong forward scattering, then perhaps the approximation that the cosine of the scattering angle, μ lies between 0.7 and 1.0 is not so bad, particularly when detailed angular distribution data are unavailable.

Gelbard: One reason that I asked this question is that I noticed that the IAEA has started some sort of benchmark activity and there is some sort of benchmark activity in the American Nuclear Society. It seems to me that it might be worthwhile at some point to have detailed benchmark output from different Monte Carlo codes. So far the only data that people have reported are eigenvalues for JEZEBEL and some other reactors. One could get much more detailed output from different codes which are supposed to be using the same cross sections, so as to compare cross-section treatments in different ranges. We have found it rather difficult to decide when we have a good enough representation of resonances to meet our needs, and we have only been able to make

this decision by comparing one code with another. So it seems to me that it would be a very good idea to start some sort of cooperative benchmark activity.

Borgwaldt: I agree with you. Unfortunately, we have the same sort of trouble at our own center. Our data are based on KEDAK because we need to use KEDAK processing codes. The multigroup data sets which are mostly used for our computations contain some secondary corrections, because people say that the KEDAK set cannot be taken as a standard since we know for sure that in some areas the cross sections must be corrected. So we even have trouble comparing calculations performed with our own multigroup sets and our Monte Carlo data sets.

THE GENERAL MONTE CARLO CODE MONK

J. G. Moore

Safety and Reliability Directorate
United Kingdom Atomic Energy Authority
Culcheth, Warrington, Cheshire, England

THE GENERAL MONTE CARLO CODE MONK

J. G. Moore

Safety and Reliability Directorate
United Kingdom Atomic Energy Authority
Culcheth, Warrington, Cheshire, England

ABSTRACT

The Monte Carlo code MONK is a general program written to provide a high degree of flexibility to the user. MONK is distinguished by its detailed representation of nuclear data in point form i.e. the cross-section is tabulated at specific energies instead of the more usual group representation. The nuclear data are unadjusted in the point form but recently the code has been modified to accept adjusted group data as used in fast and thermal reactor applications.

The various geometrical handling capabilities and importance sampling techniques are described.

In addition to the nuclear data aspects, the following features are also described; geometrical handling routines, tracking cycles, neutron source and output facilities.

INTRODUCTION

The Monte Carlo code MONK was originally written to replace the Monte Carlo code GEM [1,2]. GEM was primarily written to help in the assessment of criticality in chemical and metallurgical plant processing fissile materials and also the storage and transport of these materials. MONK not only will perform these tasks but also can perform calculations of interest to the reactor physicist. In criticality assessment work, 'boundary' tracking is quite often used, especially for problems involved in the transport of fissile materials, since this option enables interactions between several similar items to be calculated as a subsidiary calculation. Fission to fission tracking has been incorporated as an optional choice in MONK to enable the effective multiplication constant, keff, to be calculated directly from the number of neutrons in successive generations.

MONK, in common with most Monte Carlo codes, is in principle capable of dealing with any geometrical arrangement. A description is given of the standard geometry types available which enable most systems to be reduced to

some idealised system. When complex geometries, not amenable to description by the basic types arise, the concept of 'hole routines' enables this to be done. In cases where an existing hole routine will not give adequate description it is a relatively simple matter to add further 'hole routine' modules, but because of the large choice available this is hardly ever necessary.

MONK is distinguished by its detailed representation of nuclear data. The total cross-section used in MONK is in the form of point data i.e., the cross-section is tabulated at specific energies instead of the more usual group representation. Linear interpolation is done between cross-section values enabling lethargy to be carried as a continuous variable in the calculation. A particular reaction is calculated by sampling the ratio of the partial cross-section to the total cross-section; this ratio is stored in a subsidiary group structure. The basic source of the nuclear data is the UK nuclear data library which covers the range from thermal energies up to 10 MeV. The nuclear data used in the point version is unadjusted, but recently the code has been modified to accept adjusted data as used in fast reactor calculations viz. FD5, and data used for thermal reactors viz. WIMS.

NUCLEAR DATA

The calculation of the distance that a neutron will travel in a material is found by sampling an exponential probability distribution; thus if $P(x)$ is the probability of a neutron travelling at least a distance x in an infinite homogeneous medium of macroscopic total cross-section Σ then:

$$P(x) = \exp(-\Sigma x) \qquad (1)$$

This equation is used to associate a given distance x with a random number Z by equating Z to $P(x)$. MONK uses equation (1) to determine the next collision point. Thus it is necessary to calculate Σ for each medium from the compositions and nuclear cross-sections.

Point Nuclear Data

In the version of MONK which uses point data the basic source of data is the UK nuclear data library [3,4]. In the data library cross-sections and energies are given in pairs such that linear interpolation on a log-log scale introduces acceptably small errors. MONK determines the interval in which the incident neutron energy lies and calculates the microscopic cross-section from equation (2):

$$\log \sigma = A \log E + B \qquad (2)$$

where σ is the microscopic total cross-section and E the incident neutron energy. The quantities A and B are determined by a data preparation program, POND [5] for each energy interval. Thus the total cross-section is as accurate as the main data library.

Since nuclides are often common to several materials, the microscopic cross-section is saved, so that if it is required in another material, at the same neutron energy, it is immediately available. This saves the time taken to identify the energy interval and to compute the value of the cross-section.

In order to do this it is also necessary to save the neutron energy to compare with the incident neutron energy for the next time that the total cross-section is needed for that nuclide.

Types of collision and their determination

When it has been determined that a collision has occurred with a particular type of nucleus it is necessary to determine the type of collision in order to determine the parameters of motion for any secondary neutrons.

MONK recognises six basic types of collision which are:

(i) elastic scatter
(ii) absorption
(iii) fission
(iv) (n,n') or inelastic scatter
(v) (n,2n) reaction
(vi) (n,3n) reaction

The data library gives cross-sections for these reactions at the same energy points and in the same manner as for the total cross-section. The probability of a particular reaction is identified by sampling the ratio of partial cross-section to total cross-section.

The data preparation program POND produces a subsidiary group structure for each nuclide. These are broad groups and the probability of each reaction is constant within the group. The choice of group boundaries is optimised by a group reduction technique for each nuclide but each will be one of the energy points given in the original data file. This means that when the interval in the total cross-section curve has been found the broad energy group interval is uniquely determined. The probabilities are sampled in the order in which the basic collisions have been defined.

It should be mentioned that although group constants are used for the partial cross-sections and point cross-sections for the total cross-section; the group constants used are those for the probability of the separate reactions. This implies that the partial cross-section varies as a constant fraction of the total cross-section for that group and is therefore not constant in a group in the usual way.

Group reduction process

The aim of the group reduction process is to find a good representation of the different probabilities using step functions. Given the number of steps to be used over a given range, POND will obtain the best set of boundaries and the value of the probability in each group.

This is achieved by discarding a single point at a time from the mesh used for the total cross-section curve. An average value of each probability is calculated for each interval and at every energy point the reaction with the largest weight change is found. The point with the smallest of these maximum changes is now discarded and a new group is formed from the groups on either side of the point with recomputed average probabilities. A means is provided whereby certain energy points may not be removed.

In this section the suffix k refers to quantities associated with energy points and the suffix j refers to reaction types. There are NT energy points and A actions.

Each cross-section at every point is converted to a probability using the sum of the partial cross-sections

$$p_{kj} = \frac{\sigma_{kj}}{\sigma_k} \quad \text{for } k = 1, \text{NT and } j = 1, \text{A} \tag{3}$$

where

$$\sigma_k = \sum_j \sigma_{kj} \tag{4}$$

When the data file is consistent, σ_k is equal to the total cross-section. The probabilities at each point are now replaced by a set of group probabilities in the (NT-1) intervals, as follows:

$$P_{k-1,j} = \frac{W_{k-1}\, p_{k-1,j} + W_k\, p_{k,j}}{W_{k-1} + W_k} \quad (\text{for } k = 2, \text{NT and } j = 1, \text{A}) \tag{5}$$

where W_k is an energy weighting function. $P_{k,j}$ is a quantity referring to the interval above the energy point E_k. The function W is constant, up to the peak of the fission spectrum, and above this energy follows the Rosen-Cranberg fission spectrum.

In addition to an energy weighting, each reaction type is assigned an importance w_j such that w_j is zero when neutrons are neither gained nor lost, unity when one neutron is either lost or gained, two if two neutrons are gained in an (n,3n) reaction and $\overline{\nu}$ for a fission.

The variation of probabilities from interval to interval, taking into account the reaction weighting, is incorporated in the quantities

$$q_{k,j} = w_j \left| P_{k,j} - P_{k-1,j} \right| \quad (\text{for } k = 2, \text{NT-1 and } j = 1, \text{A}) \tag{6}$$

At every point the maximum value of $q_{k,j}$ is used to obtain the function

$$Q_k = W_k \max_j (q_{k,j}) \qquad (\text{for } k = 2, \text{NT-1}) \tag{7}$$

where Q_k is the maximum weighted variation between the two intervals about the energy point E_k.

The procedure is to discard the energy point, E_k, with the smallest value of Q and form a single interval from the two intervals which meet at that point. This is done by preserving the weighted area under a graph of probability step function against lethargy. The probabilities for each interval are weighted by the mean of the weights at its end points.

The weighted area under two intervals of the graph S_k is given by equation (8)

$$2S_k = (W_{k-1} + W_k)P_{k-1,j}\,(U_k - U_{k-1}) + (W_k + W_{k+1})P_{k,j}\,(U_{k+1} - U_k) \tag{8}$$

A new value, $P^*_{k,j}$ is chosen for the interval between U_{k-1} and U_{k+1} which would preserve S_k if $P^*_{k,j}$. Thus,

$$P^*_{k,j} = \frac{(W_{k-1} + W_k)P_{k-1,j}(U_k - U_{k-1}) + (W_k + W_{k+1})P_{kj}(U_{k+1} - U_k)}{(W_{k-1} + W_k)(U_k - U_{k-1}) + (W_k + W_{k+1})(U_{k+1} - U_k)} \qquad (9)$$

A value of $P^*_{k,j}$ is computed for each reaction, the energy point at U_k is eliminated and the quantities q and Q are recomputed for the points at U_{k-1} and U_{k+1}. The process is repeated until the number of group intervals has been reduced to a specified number. It can be arranged that certain energy points may never be removed and this will always apply at natural threshold energies and any range boundaries.

Number of secondary neutrons

The number of neutrons emerging from a collision is automatically determined when the collision is elastic scatter, absorption, (n,n′), (n,2n) or an (n,3n) reaction. At a fission, however, the number of secondaries is a function of the incident neutron energy. This function is described in the data library by the mean number of neutrons per fission, $\bar{\nu}$, which is presented in the data library in a similar manner to cross-section, except that, generally, fewer points are used and interpolation is linear for $\bar{\nu}$ against energy. The number of secondary neutrons released in a fission, is an integral random variable whose expectation is equal to $\bar{\nu}$.

Anisotropic calculation of emergent direction cosines

If (l, m, n) and (l′, m′, n′) are the incident and emergent direction cosines of the neutron respectively, ϕ the angle of scatter (obtained from collision theory) in the laboratory system of co-ordinates and Ψ the azimuthal angle which is random in the range $(0,2\pi)$ then for anisotropic scattering in the laboratory system

$$\begin{Bmatrix} l' \\ m' \\ n' \end{Bmatrix} = \begin{Bmatrix} l \\ m \\ n \end{Bmatrix} \cos\phi + \begin{Bmatrix} -m \\ l \\ o \end{Bmatrix} \frac{\sin\phi \cos\Psi}{\sqrt{1-n^2}} + \begin{Bmatrix} -nl \\ -nm \\ 1-n^2 \end{Bmatrix} \frac{\sin\phi \sin\Psi}{\sqrt{1-n^2}} \qquad (10)$$

for $n^2 < 0.9999$. When $n^2 > 0.9999$ a degenerate formula is used:

$$\begin{Bmatrix} l' \\ m' \\ n' \end{Bmatrix} = \begin{Bmatrix} \sin\phi \cos\Psi \\ \sin\phi \sin\Psi \\ n^* \cos\phi \end{Bmatrix} \qquad (11)$$

where $n^* = \pm 1$ according as $n \simeq \pm 1$.

Isotropic calculation of emergent direction cosines

When the angular distribution is isotropic in the laboratory system, the direction cosines (l′, m′, n′) are independent of (l, m, n). In this case (l′, m′, n′) are computed from equation (11) with $n^* = 1$ and the angle ϕ is sampled from a cosine distribution so that

$$\cos\phi = 2z-1 \tag{12}$$

where z is a random number in the range (0,1). Because ϕ lies in the range $(0,\pi)$, $\sin\phi$ can be found uniquely.

Determination of azimuthal angle

In the preceding equations Ψ is a random angle in the range $(0,2\pi)$ and appears in the form as $\sin\Psi$ and $\cos\Psi$. In order to save computing time the following algorithm is used instead of direct calculation of Ψ followed by taking its cosine and sine.

If z_1, z_2 are random numbers, let $\xi = 2z_1 - 1$ and $\eta = z_2$, then ξ will be random in the range (-1, +1) and η random in the range (0,1).

$$\text{Let } \cos\Psi = \frac{\xi^2 - \eta^2}{\xi^2 + \eta^2}, \quad \sin\Psi = \frac{2\xi\eta}{\xi^2 + \eta^2} \tag{13}$$

where $10^{-4} \leq \xi^2 + \eta^2 \leq 1.0$, then it can be shown that Ψ occurs with equal probability.

Dependence of emergent energy upon scatter angle

At elastic scatter collisions and when a discrete energy level is excited in an (n,n′) reaction, the standard kinetic equations may be used. The following equations are used to calculate the scatter angle ϕ in the laboratory system from the scatter angle θ in the centre of mass system and the emergent energy E′ in terms of the incident energy E.

$$\frac{E'}{E} = \frac{1 + 2B\cos\theta + B^2}{(1+A)^2} \tag{14}$$

and

$$\cos\phi = \frac{1 + B\cos\theta}{\sqrt{1 + 2B\cos\theta + B^2}} \tag{15}$$

where

$$B^2 = A^2 + \frac{QA(A+1)}{E} \text{ and } B > 0 \tag{16}$$

where Q is the 'Q value' of the reaction, A is the ratio of the mass of the nucleus to that of a neutron.

Note that for elastic scatter, Q = 0 so that B = A and the equations simplify to a familiar form.

The condition $B^2 \geq 0$ leads to the threshold condition

$$E \geq -\frac{Q(A+1)}{A} \tag{17}$$

An approximate form may be used when A >> 1, because

$$B^2 \simeq A^2 \left(1 + \frac{Q}{E}\right) \tag{18}$$

giving that

$$E' \simeq E + Q \tag{19}$$

and $$\cos\phi \simeq \cos\theta \tag{20}$$

Determination of scatter angles

The angle through which a neutron is scattered at a collision is a function of the incident neutron energy and the type of collision nucleus. The angle of scatter is obtained by sampling a statistical angular distribution given in the nuclear data library as a frequency curve of probability against μ, the cosine of the scatter angle.

Distributions are sometimes given for particular energies with an implied interpolation, or over energy ranges. The data preparation program POND converts pointwise distributions to rangewise distributions.

Since the scatter angle data is given in the form of a probability distribution, to simulate this the standard statistical approach is to sample from a cumulative probability distribution. The probability of a neutron being scattered through an angle $\cos^{-1}\mu$ is $p d\mu$ so that in general it is necessary to solve the following equation for μ,

$$\int_{-1}^{\mu} p d\mu = z \int_{-1}^{+1} p d\mu \tag{21}$$

This is not a very convenient form for a Monte Carlo program so that unless an analytical solution is possible special techniques are adopted.

When the distribution is isotropic, the solution is simple, because $p(\mu)$ is constant and one obtains,

$$\mu = 2z-1 \tag{22}$$

When a distribution is anisotropic the program uses two techniques. The first technique is the discrete scatter angle approximation. POND produces 32 "equiprobable angles". At a scatter, the integral part of $(1 + 32z)$ is used to identify the appropriate angle. The angles μ_i are obtained as solutions of equation (21) for

$$z_i = \frac{2i-1}{64} \quad \text{for } i = 1, 32 \tag{23}$$

This technique is fast but needs a lot of computer storage. It is used for anisotropic scatters for the (n,n′), (n,2n), (n,3n) reactions which are relatively few.

The second technique evaluates μ as a cubic polynomial in z. 32 equiprobable angles are obtained as before and μ_i is now considered as a function of z_i. A cubic polynomial is fitted through the 32 points to obtain the

coefficients of

$$\mu = a + bz + cz^2 + dz^3 \tag{24}$$

This technique allows μ to be continuous and requires less machine storage. It is used for elastic scatter.

Calculation of secondary energies

When it is not possible to use the standard kinetic equations, the emergent neutron energy cannot be calculated from the scatter angle. The nuclear data library provides secondary energy laws in these instances as follows:

1. $$E' = E_d \tag{25}$$
   the neutron is emitted with a fixed energy, E_d

2. $$E' = k(E-E_d) \tag{26}$$
   the neutron is emitted with a fraction of the incident energy reduced by a discrete energy loss E_d

3. E', is obtained from a probability distribution which is independent of E

4. E', is obtained from a probability distribution dependent on E

5. E', is obtained from a probability distribution dependent on E, and $E'/\sqrt{E}$

6. E', is obtained from a probability distribution dependent on E, and E'/E

7. E', is obtained from an evaporation spectrum given by

$$q(x) = ax \exp(-x\sqrt{a}) \tag{27}$$

where $$x = E'/\sqrt{E} \tag{28}$$

Secondary energy laws hold over specified ranges and in any range one neutron may have any combination of these laws with an associated probability.

Laws 3, 4, 5 and 6 are represented in the data library as arguments of x which are E', E', $E'/\sqrt{E}$ and E'/E respectively against a probability function q(x). As with scatter angles this is not very convenient since one must solve the equation

$$\int_0^X q(x)dx = z\int_0^\infty q(x)dx \tag{29}$$

Except in the case of equation (27), equation (29) does not give an analytical solution. To overcome this, the program uses the discrete argument method and produces 64 values of X_i corresponding to

$$Z_i = \frac{2i-1}{128} \quad (\text{for } i = 1, 64) \tag{30}$$

and the integral part of (1 + 64z) is used to identify the value of X which is chosen.

Fission spectrum

The fission spectrum used by the program is common for all the fissile isotopes and incident energies. It is the Rosen-Cranberg fission spectrum.

$$f(E)dE = 0.4527 \exp(-E/0.965) \sinh(\sqrt{2.29E})\, dE \tag{31}$$

Although equation (31) does have an analytic solution involving two error functions, the discrete argument method is again used because of its smaller computation time.

Thermal data

Part of the input data to the program POND requires the specification of two energies, E_{min} and E_{max}; POND then processes the nuclear data library only between these energies. MONK uses the static nucleus model for collisions between E_{min} and E_{max} on the assumption that the kinetic energy of the nuclei of the medium between these limits is negligible compared with the kinetic energy of the neutron. The current values of E_{min} and E_{max} are 10^{-7} MeV and 10 MeV respectively. Neutrons below 10^{-7} MeV are assumed to neither gain nor lose energy in a collision and all scattering is isotropic in the laboratory system. Neutrons may also be absorbed or cause fission in the thermal region. Group constants are associated with this single energy group and are based on cross-sections averaged over a Maxwellian energy spectrum. Since inelastic scattering is not allowed, molecular binding effects are ignored.

The thermal constants are part of the input data to the processing program POND are therefore presented to MONK together with the data above 10^{-7} MeV. An option exists in the input specification of MONK of overwriting the thermal constants provided by POND by inserting the control word THDATA followed by the appropriate constants.

Resonance cross-section data

If resonances are present in the nuclear data library these are fully tabulated in the total cross-section. In the unresolved resonance region MONK merely takes the unresolved cross-section data without any further additions.

Accuracy of MONK calculations using point nuclear data

Calculations have been done with MONK on a wide range of critical systems. The calculation of k-effective for various hydrogen to fissile ratios has shown some definite trends. It appears that k-effective is under-predicted by about 5% in the under-moderated region, over-predicted by 5% in the region of optimum moderation. These accuracies are typical for unadjusted data and although this accuracy is probably good enough for criticality clearance work

it certainly would not be good enough for reactor physics applications.

The main use of MONK in the past has been in criticality clearance work where a versatile general code has been necessary in order to treat the diversity of problems found in chemical plant and transport problems. The use of point data covering a very wide range of problems, has, to date, been found preferable to the use of adjusted group data, although the accuracy is inferior. The advantages of point data are that it does not require a fundamental mode calculation to produce the data and secondly there is no need to calculate resonance self-shielding as a special exercise as this is automatically accounted for, if the resonance shape is present in the original data library.

GROUP NUCLEAR DATA

As mentioned above, point nuclear data is based on the nuclear data library which contains so called unadjusted nuclear data. A considerable expertise has arisen in the field of reactor physics known as 'nuclear data adjustment studies' [6]. In this, the aim is to adjust the library data or some condensed form of it in order to improve agreement between integral experiments and calculations done using either Diffusion Theory or Transport Theory. In the UK this has given rise to FD5, [7] the 37 group data for fast reactor calculations and the nuclear data associated with the computational scheme known generally as WIMS [8] [9]. It has become clear that if Monte Carlo codes are to be used in reactor physics work then in the first instance they must use nuclear data which is in common use in these fields. It was for this reason, and the fact that group data could give better representation in the unresolved resonance region and thermal region that MONK has been modified to accept FD5 and WIMS data. The version of MONK using group data is known as MONKG.

MONKG Using FD5 Data

In MONKG the input data is the same as for MONK except the nuclear data section. The nuclear data is provided by a fundamental mode calculation using MURAL, a collision probability code [10]. The FD5 37 group data was obtained from the 2000 group library FGL5 by condensation using the fundamental mode spectrum. At present the nuclear data is read from punched cards but may be written to magnetic tape by MONKG by inserting the word WRITE, subsequent runs can then read the nuclear data from tape. In contrast to the point data, the group data is in the form of macroscopic cross-sections. Previously a limit was imposed on the number of downscatter groups but this has been removed in the present version. The lowest energy bound for each group in turn is read in the input data followed by the fission spectrum for each group. In MONKG only isotropic scattering is possible, this is compensated in the usual way by using the transport cross-section Σ_{tr} where

$$\Sigma_{tr} = \Sigma_t - \bar{\mu}\Sigma_s \tag{32}$$

where Σ_t is the total cross-section and $\bar{\mu}$ the average cosine of the laboratory scattering angle.

The scatter within a group is not given explicitly and is therefore calculated from

$$\Sigma_{ssc}(I) = \Sigma_{tr} - \left[\Sigma_{cap}(I) + \Sigma_f(I) + \sum_{J+1}^{NS} \Sigma_{ds}(I, I+J) \right] \tag{33}$$

where

$\Sigma_{ssc}(I)$ = self scatter cross-section in group I

$\Sigma_{cap}(I)$ = capture cross-section in group I

$\Sigma_{ds}(I, I+J)$ = downscatter cross-section from group I to group (I+J)

NS = downscatter groups

The term in square brackets in equation (33) is referred to as the 'removal cross-section' Σ_{rem}.

At present no upscatters are allowed in this version of MONKG and none of the input cross-sections may be negative.

The number of neutrons per fission is calculated from $\nu\Sigma_f(I)/\Sigma_f(I)$ since ν is not given explicitly for each group.

MONKG Using WIMS Data

The data format produced by the WIMS scheme is different from that of FD5. In this case the number of upscatters need not be zero. The total cross-section is obtained by summing the partial cross-sections

$$\Sigma_t(I) = \Sigma_{cap}(I) + \Sigma_f(I) + \sum_{J=1}^{NG} \Sigma_{SC}(I, J) \tag{34}$$

where $\Sigma_{SC}(I, J)$ is the scatter cross-section from group I to J. The removal cross-section in this case is

$$\Sigma_r(I) = \Sigma_t(I) - \Sigma_s(I, I) \tag{35}$$

In the WIMS output $\Sigma_s(I, I)$ is adjusted so that

$$\Sigma_{tr}(I) = \Sigma_t(I) \tag{36}$$

which sometimes causes $\Sigma_s(I, I)$ to become negative, and $\Sigma_r(I) > \Sigma_t(I)$ so that the MONKG reduces $\Sigma_t(I)$ so that

$$\Sigma_r(I) = \Sigma_t(I) \tag{37}$$

Determination of reactions in MONKG

The scheme for the calculation of reactions in MONK using FD5 and WIMS data is as follows:

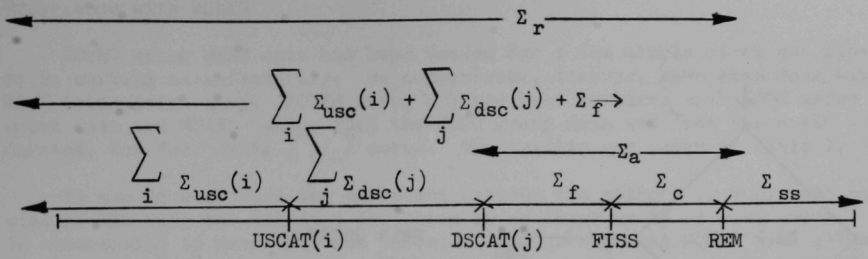

where Σ_{usc} = upscatter cross-section

Σ_{dsc} = downscatter cross-section

Σ_f = fission cross-section

Σ_a = absorbtion cross-section

Σ_{ss} = self scatter cross-section

Σ_c = capture cross-section

Σ_r = removal cross-section

i = number of upscatter groups

j = number of downscatter groups

and USCAT, DSCAT, FISS, REM represent the cumulative probabilities, obtained by dividing the appropriate cross-sections by the total cross-section.

A number RN between 0.0 and 1.0 is chosen randomly and the reaction determined by the following algorithm:

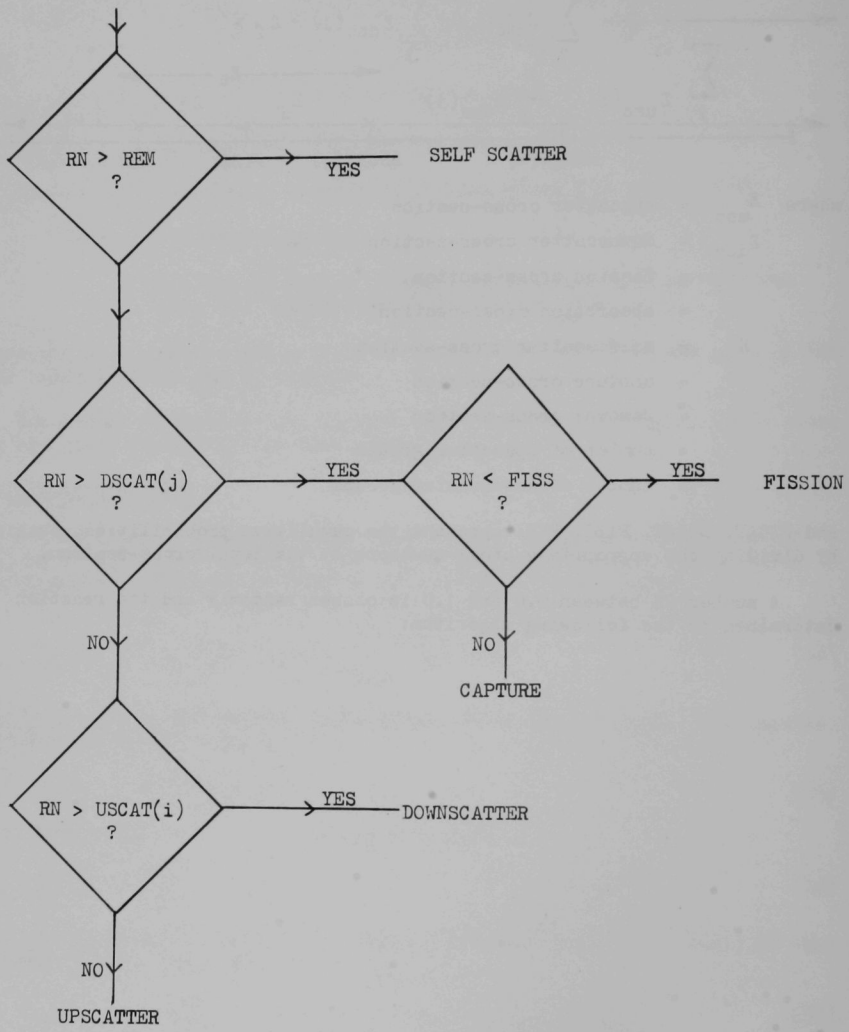

Sampling in this way is obviously the most efficient.

Experience with MONKG

MONKG using WIMS data has been tested for a few simple cases and appears to be working satisfactorily; no comparisons, however, have been done with MONK using point data. MONKG using FD5 has been compared with MONK using point data, and MURAL using FGL5 the 2000 group data set from which FD5 is derived, for four ZEBRA [11] cores. The results are shown in Table I.

It can be seen that the agreement between the methods is very good in view of the fact the standard deviation is as large as 2% in some cases. It is interesting to note that the calculations for MONK and MONKG were allowed for the same computing time and MONKG has about half the standard deviation of MONK. It was found that in fact MONKG was four times faster than MONK. The heavy penalty to be paid for the greater degree of sophistication in the nuclear data in MONK is obvious, with apparently no significant improvement in accuracy.

GEOMETRY

Basic Geometry

The geometry routines in MONK permit complicated geometrical configurations to be treated. Despite this the input remains simple and often it is possible to specify a system in a number of ways. Care has to be taken however in the choice of options, as the efficiency of calculation varies with the particular route chosen in the program.

The basic geometry of the system is described by a combination of "TYPES". TYPES may touch or be completely enveloped by another TYPE but may not intersect: in some cases a TYPE may have inner "REGION" boundaries which separate different materials or further geometrical detail. Each REGION may have one of the following basic shapes: sphere, cuboid, or a cylinder. The basic TYPES are as follows:

1. A NEST is a set of regions, each successive one entirely enclosing the previous region.

2. A CLUSTER consists of exclusive REGIONS not enclosing each other but contained in an outer surface.

3. An ARRAY is a three dimensional array of cuboidal TYPES which may be different in size and shape subject to the limitation that the entire array is specified everywhere as a series of TYPES.

4. SAME - This type will repeat a previously defined TYPE in a different geometrical position.

5. LIKE - This type will repeat a previously defined TYPE except that the contents of the REGIONS will be different.

Each type has its own origin and frame of co-ordinates, and in addition there is the facility for changing the origin of a REGION with respect to the TYPE origin. Provided the linkage between TYPES is specified in the data, MONK

produces a hierarchal system for the types with a single parent type. An error message is printed if this cannot be done. The program therefore automatically transfers between co-ordinate systems and hence the user need not specify the co-ordinates of the type origin with respect to the parent TYPE.

Hole Routines

In the basic geometrical TYPES described so far, the REGIONS may contain either another type or a single material. If this is insufficient it may be necessary to have more detailed representation of the geometry supplied by a 'HOLE routine'. HOLE routines will describe systems containing non-standard shapes e.g. hemispheres, cones, rings, intersecting cylinders or arrays of rods not suitable for ARRAY specification. A HOLE may in principle contain any number of objects of any shape, several such routines are available in MONK and others can easily be added as sub-routines.

In the normal type of tracking in a single material, the mean free path is calculated as the reciprocal of the macroscopic total cross-section and is used to calculate a random path. If this path, taken in the direction of motion of the neutron, crosses a region boundary, the neutron is moved along its line of motion just as far as the boundary and a new mean free path is taken in the next region. If the neutron does not reach a boundary, it is moved along the whole random path and is then assumed to have a collision in that material, which will give rise to some appropriate action.

When a neutron is in a region containing a HOLE, the mean free path is always taken to be that corresponding to the material in the region with the largest total cross-section, Σ, when a collision point is determined, a sub-routine called a HOLE routine is called upon to decide which material is present at that particular point, Σ_i. If the material at the point has cross-section Σ, then a collision occurs. If Σ_i is less than Σ then the program decides whether a true collision has occurred by comparing Σ_i/Σ and $1 - (\Sigma_i/\Sigma)$ with a random number, z. If $z > \Sigma_i/\Sigma$ the neutron continues on the same path with no change in energy.

In general provided that $\Sigma \geqslant \Sigma_i$ it can be shown that the correct distribution of collision density is obtained. The effect of this is to make a region containing a number of materials typified by a single mean free path for tracking purposes. Thus the tracking in HOLE routines is simpler since a calculation of boundary crossings has been replaced by the testing of inequalities.

The development which has taken place in this context has been the facility of replacing materials in HOLE routines by further HOLE routines and the process can be continued as often as required. In this way effectively all geometrical restrictions are removed from the program. A further refinement in HOLE tracking is associated with an attempt to remove one of the disadvantages. In a region containing a HOLE, the neutron tracking proceeds in steps of average length equal to the shortest mean free path appropriate to any material declared as present. If a material with a high cross-section is present in only a small volume of the region, the tracking steps are short over the whole region and consequently tracking is slow. In some cases this handicap can be overcome by isolating these materials and respecifying the

HOLE routines, when this is not feasible an option exists in MONK to include the material between two planes expressed as

$$Ax + By + Cz = D_1 \tag{38}$$

$$Ax + By + Cz = D_2 \tag{39}$$

in the co-ordinates appropriate to the region. Tracking then proceeds ignoring the materials occurring solely in the specified zone except when the track so calculated crosses that zone when reversion is made to the normal method.

IMPORTANCE SAMPLING

A 'FILTER' provides a simple form of importance sampling and the position and operation is optional to the user.

Filters

A FILTER can be used on any region boundary in a system and is typified by its power and the direction in which it is acting. Thus, when the power of a FILTER is N, only every N-th neutron moving outwards across the FILTER boundary is tracked; when a neutron returns to this boundary it will be taken N times. The N neutrons will have identical co-ordinates on entering the inner region but their tracks will diverge after one collision each.

NEUTRON SOURCES

A neutron source is used to begin the first stage or cycle of any calculation; in normal use the source neutrons are chosen at random from the fission spectrum. In other applications a source of specified spectrum and angular distribution may be used to start every stage or in some cases the first stage or any stage which becomes extinct.

Type of Source

The source may be defined in terms of geometrical shape, energy and direction.

Geometry

The source may be placed at any specified point, line, surface or distributed throughout a volume. In the case of a surface source, the surface may be that of any sphere, cylinder or cuboid or any particular face of a cylinder or cuboid.

Energy spectrum

The source energy distribution may be sampled from a fission spectrum or a set of energies defining equiprobable ranges.

Direction

The source neutrons can be isotropically distributed in direction, in a fixed direction or within given ranges of the direction cosines.

Magnetic tape and punched card options

An option exists in MONK for writing the neutron co-ordinates, energy, and direction cosines on to magnetic tape when neutrons cross a specified boundary; this tape can subsequently be used as a source of neutrons. In addition, at the end of a calculation, 'continuation cards' may be punched. These cards contain all the necessary information concerning neutron co-ordinates and boundary crossings to enable the previous run to be extended without any discontinuity. In this case a source of neutrons is not required to re-commence the calculation.

NEUTRON TRACKING

In MONK there are two types of tracking available and are known as "fission to fission" and "boundary" tracking.

In "fission to fission" tracking, the tracking cycle or stage is defined as that between one fission generation and the next. In "boundary tracking" the tracking cycle is defined by tracking neutrons which begin and end on a specified boundary.

Fission to Fission Tracking

Fission tracking enables an estimate of the multiplication constant k_{eff} to be made from the ratio of successive generations of neutrons.

$$k_i = \int_R N_{i+1}(x)dx / \int_R N_i(x)dx \qquad (40)$$

where x is a multidimensional vector in phase space, composed of geometrical and energy vectors. When sufficient cycles have been sampled the fundamental mode will be reached and

$$\lim_{i \to \infty} k_i = k_{eff} \qquad (41)$$

If k_{eff} is less or greater than unity the population either decreases or increases, in this situation MONK adjusts the neutron population at each stage so that a constant number of neutrons is tracked each time.

In addition to this, MONK also estimates k_{eff} from the probability of a fission occurring at every collision. Since every collision makes some contribution, one might expect this estimate to be better, statistically. Analysis of various calculations to date has shown that both estimates appear to be statistically independent and within one standard deviation of each other. It would seem that the best estimate of k_{eff} would be obtained by averaging the two estimates.

Boundary Tracking

The techniques of boundary tracking have been described elsewhere [1,2] so only a brief summary will be given here. The usual procedure is to enclose the fissile material by a boundary which is then used for the purposes of establishing the tracking cycle. The system is then divided into a "core" containing the fissile material and a "reflector" containing no fissile material.

If it is assumed that a settled distribution of neutrons has been achieved, then one neutron crossing the boundary into the core will on average produce M neutrons that return to the boundary, and if a neutron crossing the boundary into the reflector results, on average, in R neutrons returning, then the quantity MR is a measure of criticality. The system is supercritical, critical or subcritical according to whether MR is greater than, equal to or less than unity.

Neutron interaction between fissile units in an array

A problem commonly met in criticality work is the assessment of the criticality of an array of different interacting units. In principle a Monte Carlo code could estimate the criticality of such a system by tracking neutrons in the normal way. In large arrays the tracking time can be prohibitive. The alternative is to use a method referred to generally as the interaction parameter method [12].

Let F_i = number of neutrons leaving ith element of the array

P_{ij} = probability of a neutron leaving the jth element reaching the ith element without interacting with any other element of the array

M_{ij} = surface multiplication of the ith element of the array to neutrons reaching it directly from the jth element of the array

The quantity

$$Q_{ij} = M_{ij} P_{ij} \tag{42}$$

is known as the interaction parameter between elements i and j.

The steady state equations for the neutron fluxes are

$$F_i = \sum_{j=1} M_{ij} P_{ij} F_j \quad (i=1,n) \tag{43}$$

The criticality condition is then that the largest eigen value of the matrix equation

$$\lambda [F] = [Q] [F] \tag{44}$$

shall be equal to unity.

The Monte Carlo method is a suitable way of obtaining values of P_{ij} which when associated with estimates of M_{ij} can be used to assess the criticality of an array. A version of MONK known as BLACK MONK has been written which calculates the collision probabilities, P_{ij}, of neutrons arising from each source region in turn, having a first collision with each of the other units. The collision probability matrix is then multiplied by the surface multiplications and equation (44) is solved for λ.

In the case of an interacting array of identical units, MONK will automatically calculate the number of unit lattice cells which form a critical assembly from the results of a calculation relating to one unit cell. This is made possible since all that is required are the numbers of neutrons entering and leaving each face of the unit cell and these are stored during the boundary tracking. This method is known as the PQR method [12].

PROGRAM OUTPUT

The options have already been mentioned whereby information may be transferred to magnetic tape or punched card, normally for use in recommencing calculations. The printed output gives a complete edited output of all the input data, an edit of all the collisions at each stage and a final edit on completion of the calculations. The final edit contains information on all the different nuclear interactions which have occurred and region crossings for every type. The quantities k_{eff}, M and R are also printed and their standard deviations.

Neutron fluxes, defined as the total track length per unit volume, are calculated in every region of the system in the full energy spectrum but are finally condensed to sixteen energy groups. Reaction rates may also be obtained by specifying a nuclide with a negative concentration in the input data. The reaction rates are calculated as they occur in the calculation and are available in the detailed energy spectrum.

CHECKING FACILITIES

A comprehensive set of checks are built into MONK which are in operation during input and initialisation and a few which occur during the calculation. However the program may not necessarily fail when an erroneous specification is given. To overcome this difficulty the options in MONK known as SCAN and PERCY have been devised. These are particularly useful when the geometry is complex.

The Option SCAN

The option SCAN was written to check the geometrical arrangement of the materials and will print out material numbers to give a two dimensional picture of a plane section through the system. The plane can be at any angle to the co-ordinate axes. SCAN uses the tracking routines of MONK to follow paths through the system with a predetermined step length. The material number is printed according to a prespecified code of symbols. The tracking procedure scans a line at a time, at the end of each line the path is turned to take it to the next line. It is then turned again so that it is parallel to the original path. The scanning continues back and forth until the picture is complete.

The Option PERCY

In complicated geometries it is difficult to estimate the amount of a given material in region containing a hole routine. The option PERCY generates points at random evaluating the material number at each point. A prespecified number of points are sampled and the percentage of each material in the region calculated.

ACKNOWLEDGMENTS

The author would like to thank Mr. P. J. Hemmings who wrote the first version of MONK, Mr. J. J. Cruickshank and Mr. A. Brickstock who removed many of the errors in the code and finally Mr. E. R. Woodcock and other colleagues in the UKAEA for many helpful discussions.

REFERENCES

1. P. J. HEMMINGS, "The GEM Code", AHSB(S)R105 (1967).

2. T. C. LONGWORTH, "The GEM4 Code", AHSB(S)R146 (1968).

3. D. S. NORTON, "The UKAEA Nuclear Data Library", AEEW-M828 (1968).

4. A. L. POPE and J. S. STORY, "Minigal Output from UK Nuclear Data Library - NDLI (1973) Thermal Cross-sections, Resonance Integrals and Fission Spectrum Averages", AEEW-M1191.

5. P. J. HEMMINGS, "Use of the UK Nuclear Data Library in the Monte Carlo Code MONK", AHSB(S)R184 (1970).

6. C. G. CAMPBELL and J. L. ROWLANDS, "The Relationship of Microscopic and Integral Data", IAEA International Conference on Nuclear Data for Reactors IAEA-CN-26/116, Helsinki, Finland (June 1970).

7. J. L. ROWLANDS et al, "The Production and Performance of the Adjusted Cross-sections set FGL5", International Symposium on Physics of Fast Reactors, Tokyo (October 1973), Vol. III, p. 1133.

8. J. R. ASKEW, F. J. FAYERS and P. B. KEMSHELL, "A General Description of the Lattice Code WIMS", Journal of the British Nuclear Energy Society, Volume 5, Number 4, pp 564-585 (October 1966).

9. R. CHAWLA, "An Assessment of Methods and Data for Predicting Integral Properties for Thermal Reactor Physics Experiments", AEEW-R797.

10. J. D. MacDOUGALL, R. W. ROSS and J. L. ROWLANDS, "The Calculation of Neutron Spectra and Group Averaged Cross-sections Using the Computer Programs FRESCO and MURAL", AEEW-M843 (1968).

11. R. D. SMITH, A. R. BAKER, and J. L. ROWLANDS, "Theoretical and Experimental Work on the Physics of Fast Reactors", Journal of the British

Nuclear Energy Society, London Conference on Fast Breeder Reactors (May 1966).

12. A. F. THOMAS and F. ABBEY, "Calculational Methods for Interacting Arrays of Fissile Materials", Pergamon Press Ltd., International Series of Monographs in Nuclear Energy, Volume 108 (1973).

TABLE I

Comparison of k-effective using MURAL, MONK and MONKG for four ZEBRA assemblies

ASSEMBLY	FUEL	DATA SET	PROGRAM	k_{eff}	S.D.
ZEBRA 1	U	FGL5	MURAL	0.993	
		RD48	MONK	1.001	0.018
		FD5	MONKG	1.003	0.008
ZEBRA 2	U	FGL5	MURAL	0.992	
		RD48	MONK	0.979	0.021
		FD5	MONKG	0.996	0.008
ZEBRA 3	Pu	FGL5	MURAL	0.984	
		7ES184	MONK	0.984	0.019
		FD5	MONKG	0.987	0.008
ZEBRA 6A	Pu	FGL5	MURAL	0.976	
		7ES184	MONK	0.944	0.021
		FD5	MONKG	0.970	0.008

DISCUSSION

Kalos: Did I understand correctly that in black MONK you track from one region into another and then stop the history, regarding the next region as a purely black absorber?

Moore: Yes.

Kalos: How do you select the starting positions in the regions in which you start?

Moore: The points are taken at random on the surface of each array.

Kalos: Then that version of the code represents an approximation. The k that you get is then an approximation to the k_{eff} of the assembly. The equation you wrote down is correct only when the points are chosen in region I according to the equilibrium distribution of collisions in region I.

Moore: Black MONK is only used to calculate collision probabilities. It does not separately calculate the eigenvalue.

Kalos: That's true, but the approximation of the probabilities really should be calculated not from uniformly distributed collisions within the region but collisions distributed according to the eigenvector, the eigendistribution.

Moore: It is an approximate method, but one has to remember that it is used in criticality work and that ...

Kalos: Well, is there any reason why you couln't use the black MONK idea in the course of doing an ordinary MONK calculation? It is just a matter of tallying and then you can have your cake and eat it too.

Moore: I hadn't thought of that.

Gelbard: Now aren't you describing something very much like the Green's function technique where people accumulate the Green's function as they go along?

Kalos: Yes, I'm just pointing out that, to be correct, the Green's function must be accumulated in the eigenfunction.

Gelbard: In the main mode. Yes, that is an approach that has been used at a number of places.

Cashwell: Did I hear you say that you have an energy-dependent fission spectrum in the code?

Moore: Yes.

Cashwell: That means that you don't have to worry about prefission neutrons? Do you worry about that at all?

Moore: We sample from the formula I gave, which is the Rose and Cranberg fission spectrum.

Cashwell: Which is, I guess, what I call a Maxwellian?

Moore: Yes.

Whitesides: Why does the point version run so much longer than the multigroup? I really don't see what you are doing so differently in the point and multigroup versions so as to make the point version take so much longer to run.

Moore: Well, I guess it's because one does not have to do linear interpolation in the cross sections and there are fewer energy groups to look for in the calculation.

Whitesides: But linear interpolation shouldn't take that long to do.

Moore: There's also isotropic scattering in the multigroup version.

Whitesides: MONKG uses isotropic scattering?

Moore: Yes.

Whitesides: Did GEM do that?

Moore: No, it used anisotropic scattering. GEM was a point version.

Bending: In reference to a previous question: You said that the cross sections were evaluated by linear interpolation in the point version. Is this linear interpolation in energy or lethargy?

Moore: Lethargy.

Bending: May I ask a second unrelated question about the multigroup versions of MONK using WIMS data? I'm interested in knowing what action you take when your transport correction to the self-scatter term generates a negative self-scatter term. I believe there was something about this on one of your slides.

Moore: If any self-scatter term becomes negative, in that case we adjust the cross sections to keep it positive. We adjust the total cross section.

Bending: I am not able to give full details of the methods used from memory, but I know that another Monte Carlo code used in the UKAEA does treat the problem of negative self-scatter cross sections by introducing a special scattering law which, instead of being isotropic, includes a certain probability of anisotropic scatter and a certain probability of a direct forward scatter. I am not able to give accurate details of how that is done, but I know that it is done and gives very good results.

Borgwaldt: In the summary of your paper you referred to the "concept of filters" as a form of importance sampling. I'm not sure whether I missed this point. Could you elaborate?

Moore: I didn't mention it in the talk—it's in the paper. It is a very simple form of importance sampling and only applied to the fluxes. A filter can be used on any region boundary and is typified by its power and its direction so that if the power of the filter is n, every n-*th* neutron moving out is tracked, and the reverse happens when the neutrons come in the opposite direction.

Gelbard: I don't understand how this is a filter. Could you explain a little more?

Moore: If you have a filter you can have it operating in either direction. For instance, if you have n neutrons crossing a boundary, and the filter is operating in a way in which it reduces the flux, then only one neutron will be tracked across that boundary.

Kalos: Is this to be distinguished from Russian Roulette and splitting?

Moore: I don't really think that it is one or the other, but it is just a simpler version of it.

Gelbard: I might mention that I think the reason for some confusion is that there has been a paper on the use of filtering in Monte Carlo. There is a thesis on filtering in Monte Carlo as an application of techniques that are used in communication theory. Perhaps that is part of the reason for the confusion in this term.

Moore: The two types of filtering probably aren't related.

Gelbard: No, filtering, as you use the term, sounds like Russian Roulette.

Coveyou: You quoted a 7% error. Was that 7% in k or Δk?

Moore: A 5% error in k.

Coveyou: 5%?

Gelbard: That sounds like a rather big error. I think it would be much bigger than what you would get using ENDF/B unadjusted. That is my impression, based on information from people who have done critical calculations using ENDF/B.

Coveyou: Perhaps I misunderstood. Is this 5% as compared with the experimental values or is this a 5% statistical error?

Moore: These are on critical systems and one could get values of k ranging from 0.95 to 1.05. The standard deviation might be 1%. So it is unlikely that the discrepancies can be attributed to statistics. More likely they are systematic errors associated with the inadequacy of the data in certain energy regions. I have seen results of Monte Carlo calculations on other codes which gave similar accuracies for unadjusted data.

Taormina: Is your code able to treat shielding problems?

Moore: No.

Gelbard: I wonder if there are any other comments about the accuracy of critical eigenvalues using other data sets?

Borgwaldt: I should say that with our unadjusted data we normally lie, let us say 2 or 3% from the experimental values. That's normal. But the unadjusted data which we use are based on a rather old nuclear data library. Our KEDAK file is in the stage of being updated. We hope that in a few weeks we will have an updated set. But, 2 or 3% seems plausible for fast criticals.

Gelbard: Any other comments on this subject?

Hoogenboom: I have a remark on the formulas used with the scattering process. You select a direction in the center-of-mass system and then transform to the laboratory system. But there is a simple relation between the direction in the center-of-mass system before and after scattering, and that relation can be used whether the scattering is isotropic or anisotropic. The same relation can be used with only a small change if the scattering is elastic or unelastic. I think that it is much simpler than the transformation to the laboratory system.

Moore: I don't think there is much penalty in computing time in performing the transformation as we do.

Borgwaldt: I would like to comment on the concept of filters. Of course, the filtering in one direction is related to Russian Roulette. One know that Russian Roulette is free of bias but can introduce variance. This filtering technique, which is a systematic technique, seems to be of lower variance, but I question if one can guarantee that such a procedure is free of bias.

Moore: I should say that we use this filtering mainly in thick reflectors, so that we don't expend a lot of time in unimportant regions.

Kalos: Would you please explain again exactly what the filtering procedure is?

Moore: If we had a filter of the power n in a region, and if n neutrons impinge on that region in a given direction, then only one neutron will be tracked on the other side.

Kalos: On the average, or do you track every n-*th* neutron?

Moore: Just 1 over n of the number of neutrons.

Kalos: Is this just chosen at random?

Moore: Yes.

Kalos: Then it is Russian Roulette!

Whitesides: I think it is Russian Roulette because as you move into regions you deem to be less and less important, you track fewer and fewer neutrons, and as you move back you split into more and more neutrons and you do it at a boundary.

Kalos: Yes.

Whitesides: It is just a device for implementing Russian Roulette, and you split coming back in the other direction.

Gelbard: Is it correct that you do split when you come back out?

Moore: Yes.

Kalos: Then that technique is unbiased, and either introduces or does not introduce variance according as to how well the weights approximate the relative importance of the two regions, and according to other tactics as well.

Bending: As far as I understand it, the system used in MONK is precisely that which is normally termed splitting and Russian Roulette, and I think just by way of comment that the fact it is called filtering is historical accident based on the considerable age of the original GEM code. I think it is a rather unfortunate confusion, but I think had the code been written maybe five, six, or ten years later it would have been called splitting and Russian Roulette. Then there would be no confusion.

Kalos: Actually I think that the phrase "Russian Roulette" is the unfortunate phrase and I wish we had adopted the phrase "filtering" from the beginning. I would like to point out a fact that ought to be well known, but probably is not. In using splitting and Russian Roulette, one does two things: First, one assumes that the importance function depends upon the region and has a discontinuity at the surface of the region. Second, one accomplishes the biased sampling required for importance sampling by splitting and Russian Roulette at the surface. There are other ways of implementing exactly the same assumption about the importance function with lower variance in the final answer. These other techniques also lower the collision density in the regions deemed to be of lesser importance, with lower variance in the final answer. They lower the collision density in the regions deemed to be of lesser importance, but do so in a more systematic and continuous way. I will give you references later. One such scheme is implemented in the SAM codes, and is used systematically in place of splitting and Russian Roulette at the boundaries.

Coveyou: Actually in 05R for a long time we didn't do the splitting and Russian Roulette at the boundaries; we did it at the first collision point in the region.

Kalos: That's the worst scheme of all of these as far as ultimate variance of the Monte Carlo estimates is concerned.

Gelbard: I would like to refer back to the subject we were discussing before, mainly the utility of Monte Carlo in general. I noticed that there have been some Monte Carlo studies of the ZEBRA assemblies. These were safety assemblies in which there were small perturbations introduced to test the effects of accidents. In accident analysis you often have configurations in which the geometry is really pretty hideous. You have situations where you have voided the moderator and have vacuum streaming through large voids. It seems to me that no one really has any right to trust diffusion theory in this

sort of situation. Why can't we get practical people (and I imagine that excludes some of us here), why can't we get nuclear designers and safety analysts to use Monte Carlo in this sort of situation? It seems to me that one ought to be able to sell Monte Carlo capability in the analysis of a safety configuration. Now in ZEBRA I think there was a problem in that the perturbations were very small, which might be a difficulty. To some extent, maybe, one could get around this difficulty by having Monte Carlo in mind when an experiment of this sort is done. But, I really don't see how anyone can trust the standard diffusion theory calculation in this sort of case.

Moore: I'm familiar with the experiments you referred to and I have seen comparisons of diffusion theory and transport theory on these experiments. They usually refer to single subassembly experiments where the changes were very small and the agreement was remarkably good. What I suppose one is interested in are the ramps in reactivity because variations in these ramps lead to variation of explosion yield in fast reactor calculations. I have mentioned to the experimental people that when they design the next set of experiments they ought to look at larger configurations to see if they can get larger changes and then the Monte Carlo would come in on its own. But certainly the changes in reactivity that we're getting (typically about 0.1%) are way outside any Monte Carlo capabilities at present, I think.

Gelbard: Except possibly for Monte Carlo calculation using fancy perturbation techniques.

Moore: Are there any codes using such schemes?

Gelbard: Well, yes, I think a perturbation code is going to be discussed here later at this meeting. There are perturbation methods in the SAM codes, so that there are perturbations methods available for this sort of thing, but they have to be programmed separately.

Gast: I would like to make a few comments about Russian Roulette. In the early version of our program, we used the approach of using Russian Roulette at boundary crossings. However, as the geometry of the program built up, it became more and more difficult to track these additional boundary crossings. We, therefore, retreated to splitting and Russian Roulette at collision sites. As Kalos pointed out, we did find that there was a drop, to some degree, in the efficiency of the approach; however, it was not so severe that we did not retain it. We felt that the simplicity of the new scheme made up for the slight loss in efficiency.

Summary of Comments by Gast: Gast explained that, in the Bettis code, the details of implementation of the tracking method in complicated geometries made the use of Russian Roulette at boundaries very awkward.

Moore: So it's the extra geometry that caused the difficulty?

Gast: Yes, that's the reason we retreated, I guess I should say ...

Moore: Reverted?

Gast: Reverted — Coveyou indicated that this was an old approach in 05R. We feel it has worked out fairly well.

Gelbard: In your experience, what sort of splitting factors would you normally use on a boundary? Do you have any ideas as to what's reasonable and what's not?

Gast: Some people tend to use a splitting factor of about 5 and we may have a total of 10 splitting boundaries, roughly.

Kalos: Perhaps since both of you brought up some ancient history, I will too. In the very early days, people like Kahn and Von Neumann supposed that the optimum splitting was probably a factor of 2. One of the first things I did in Monte Carlo was to set up a model Monte Carlo problem which could be solved analytically with splitting. That is, I assumed I had a slab and there was straight-ahead penetration in the slab. The slab was divided into p sub-slabs and the splitting was adjusted so that the total ratio of importance function from slab to slab was constant. Then by means of generating functions, I calculated the variance for each case and found that indeed the optimum was about one-half, very close to one-half, with a rather broad middle. So, if you find yourself splitting by regions, a good rule of thumb is to put in a region when you judge that the importance changes by a factor of 2.

Gelbard: Do you think there is any geometry effect, and do you think that it would basically be the same sort of thing if you were nesting around a point?

Kalos: No, of course, the importance function must take into accound the geometrical effects of nesting around a point. If you are trying to get flux at a point by simply concentrating collisions in the neighborhood of a point (which I think is not a very good procedure), then you would get an infinite number of ..... .

Gelbard: I didn't mean literally a point. You might be dealing with a small cylinder, or a small three-dimensional region rather than a planar region, and then you're saying that this factor of 2 is not something that you could simply accept and generalize to such a problem. Then, it is not really clear theoretically what you have to do in this case.

Kalos: I think the rule of thumb is that the factor of 2 probably works pretty well there too.

Gast: Basically our experience with Russian Roulette has been somewhat disappointing in that a problem that is completely out of range in terms of machine time is not generally brought into range by Russian Roulette. We have had a few cases which might be called borderline where the technique has helped.

Gelbard: I wonder if anyone has ever tried to optimize splitting by using a technique that was described by Jerry Spanier in a paper presented at the ANS Topical Meeting in Idaho Falls. This is really one I can't claim to have looked at very thoroughly, and I wonder if anyone else has? This is a theoretical approach to optimizing parameters of Monte Carlo. Does anybody have any feeling about it?

Bending: I would like to go back, just briefly, to something which Gelbard mentioned and this is the application of Monte Carlo to the study, for example of small perturbations in safety problems. It seems to me, as has been pointed out, that the difficulty in using Monte Carlo to estimate perturbations is great and it seems to me that the role of Monte Carlo might be modified to avoid Monte Carlo perturbation calculations. Let us take, for example, a situation in which we would want to know the reactivity effect of introducing a void, and obviously we are a bit unhappy about using diffusion theory. Possibly one could use Monte Carlo directly to work out what would be an appropriate axial diffusion coefficient to use in this case, and then use diffusion theory. Perhaps one could argue that reactivity is the wrong thing to measure when one is looking at small configuration changes. Secondly, we have a little bit of experience in using Monte Carlo perturbation methods and we have found that there are significant difficulties in interpreting the statistics. There are obvious dangers when one introduces a small perturbation. The number of neutrons which contribute to the perturbation in the results is very small and the statistics on the perturbation are not so good, either as one might hope, or as good as the correlation type of studies on the results would suggest. We have found that one has to be careful that one does not underestimate the variance, and that the statistics don't lead one to underestimate the perturbation. I say this just by way of warning.

MONTE CARLO CODE DEVELOPMENT IN LOS ALAMOS

L. L. Carter, E. D. Cashwell, C. J. Everett, C. A. Forest,
R. G. Schrandt, W. M. Taylor, W. L. Thompson and G. D. Turner

Los Alamos Scientific Laboratory
Los Alamos, New Mexico 87544, U.S.A.

* * *

Presented by E. D. Cashwell

MONTE CARLO CODE DEVELOPMENT IN LOS ALAMOS

L. L. Carter, E. D. Cashwell, C. J. Everett, C. A. Forest,
R. G. Schrandt, W. M. Taylor, W. L. Thompson, and G. D. Turner

Los Alamos Scientific Laboratory
Los Alamos, New Mexico 87544, U.S.A.

ABSTRACT

The present status of Monte Carlo code development at Los Alamos Scientific Laboratory is discussed. A brief summary is given of several of the most important neutron, photon, and electron transport codes.

INTRODUCTION

Beginning with the ideas of von Neumann, Fermi, Ulam, and others after World War II when Monte Carlo emerged as a recognized discipline, Los Alamos has carried on a program of Monte Carlo code development. The value of this method in treating complicated particle transport problems was clearly demonstrated very early in the game, and Monte Carlo played an important role in weapons development and in many other programs at the Los Alamos Scientific Laboratory [1].

In the following remarks, some of the most important Monte Carlo codes at LASL will be described briefly. All of them are under constant development in order to meet the increased demands upon them, both as to the difficulty of the problems they are asked to solve and with respect to the task of handling the vast amount of cross-section data now available. With one exception, the neutron and photon codes considered here use a continuous energy treatment of the cross sections. Pointwise data are provided at discrete energies, with interpolation employed in between. The cross sections are read into the codes in considerable detail in an attempt to use the information with no significant approximations or distortions. This puts a considerable burden on the storage capacity of the computer, especially in view of the increasing size of nuclear data compilations such as the Evaluated Nuclear Data Files (ENDF).

The codes described below all have the same three-dimensional geometry package [2]. Each has the capability of handling an arbitrary configuration of materials in regions which are bounded by first- and second-degree surfaces, as well as some fourth-degree surfaces (elliptical tori) [3]. The language used in all of the following programs is FORTRAN IV.

MCN - A NEUTRON MONTE CARLO CODE

This [2] is our standard neutron transport code and it has formed the basis for a number of other codes, some of which will be mentioned here. A great deal of work has gone into making this program as general and versatile as possible, while at the same time keeping it simple to use. It is the latest of a series of Monte Carlo neutron codes at Los Alamos, the first of which was described by R. Johnston [4].

The pointwise cross sections used in MCN at present are from the nuclear data compilations of Howerton's group at LLL, from ENDF, or from the British (AWRE). The ENDF data is processed for use in MCN by D. R. Harris' group (Group T-2) at LASL. Pointwise data are provided us with the resonance treatment included, data may be thinned to prescribed tolerances, and in some cases interpolated data may be added in order to permit accurate linear interpolation in MCN. In particular, angular data may be excessive in some evaluations, and by thinning, storage requirements may be reduced considerably.

Included in the code are standard variance-reducing techniques, which are optional. These include particle splitting and Russian roulette, path-length stretching, and machinery for biasing the source with respect to both energy and angle of emission. Provision is made for forcing collisions in designated cells, obtaining flux estimates at point detectors, and for calculating reactions in small regions by using track-length estimators.

MCN employs an energy-dependent fission routine, with pre-fission neutrons emitted in an evaporation spectrum and the remaining neutrons in a fission spectrum. Both types of spectra depend upon the energy of the incoming neutron.

The code includes a neutron thermalization routine employing the free gas model. Below a thermal cut-in energy, the lighter atoms such as hydrogen and deuterium are assumed to be in thermal motion, with a Maxwellian distribution of velocities determined by the thermal temperature of the region. Each geometric cell of the problem has its own thermal temperature which may be given as a function of time--that is, these temperatures may be specified at discrete times t_i, with linear interpolation employed to yield the thermal temperature at time t. Scattering from the light nuclei includes the effect of the thermal motion. For nuclei not belonging to this select group of light atoms, and for energies in the thermal range, elastic scattering is assumed to occur isotropically in the laboratory system with no energy loss.

The standard output of MCN includes two-way currents as a function of energy, time, and angle with the normal, across any subset of bounding surfaces in the problem. Fluxes at any set of bounding surfaces are available as a function of time and energy. Similarly, the flux at designated points, and the average flux in a cell (track length per unit volume) are standard tallies. Reactions such as fissions or absorptions may be obtained in a subset of the geometric cells. In addition, particles may be flagged when they cross specified surfaces or enter designated cells, and the contributions of these flagged particles to certain of the tallies above are listed separately. All quantities printed out have their relative errors listed also--here the relative error is defined as the ratio of one standard deviation in the sample mean to the sample mean.

MCNA - A NEUTRON ADJOINT CODE

The code MCNA [5,6], a companion code to MCN, was set up to sample from the adjoint neutron transport equation. The identical cross sections and scattering models are used by both the forward and adjoint codes. This feature is often used to check for errors by comparing forward and adjoint results for standard test problems.

The principal use of MCNA at Los Alamos has been in computing F-numbers, i.e., in computing reaction rates in some portion of a system as a function of the neutron source distribution incident upon the system. The computation is typically carried out in a coupled mode to optimize the Monte Carlo. A short forward calculation is first made with a "typical" incident neutron source in order to compute the energy-dependent flux in each region of the system. These flux estimates are then used to bias the selection of the emerging energy at collision events in the adjoint calculation by numerically approximating the near-optimal adjoint density function

$$f(E;\, E')_k = \frac{\phi(E)_k \; \Sigma(E \rightarrow E')_k}{\int \phi(E'')_k \; \Sigma(E'' \rightarrow E')_k \; dE''} \; ,$$

where $\phi(E)_k$ is the energy-dependent flux averaged over the cell volume denoted by k and $\Sigma(E \rightarrow E')_k$ is the cross section for an energy transfer from E to E′ at a collision. The "particle" weight is adjusted at each collision in the calculation to obtain unbiased results.

MCG - A GAMMA-RAY TRANSPORT CODE

We shall give a brief description of MCG [7], a Monte Carlo program for transporting gamma rays which has been in existence for many years. This code has a number of features in common with MCN, such as the geometry package, the variance-reduction techniques, and some of the output. Photon interaction cross sections from the LLL library [8] are used.

Since this program is designed to transport high energy photons, it considers only the following reactions:

(1) Compton Scattering - Sampling from the Klein-Nishina scattering law is accomplished very efficiently through the use of simple but accurate fits to the inverse of the distribution function [9, 10].

(2) Pair Production - When an electron-positron pair is produced by a gamma ray, the difference between the energy of the incoming gamma ray and the resulting annihilation radiation is deposited locally at the site of collision. The two 0.511 MeV photons emitted are assumed to be emitted isotropically at the site of collision and are transported further.

(3) Photoelectric Effect - This reaction is treated as absorption, resulting in the removal of the photon with its energy deposited locally.

The current and flux tallies described in MCN are also standard in MCG. Information about the energy deposited in each cell of the problem is automatically listed. The energy deposited due to Compton collisions, pair production, and photoelectric absorption is listed separately for each type of reaction, along with the total energy deposited per cell.

MCP - A GENERAL MONTE CARLO PHOTON CODE

The gamma code MCG is not designed to deal with low energy photons, say in the range 1-50 keV for some materials. Fluorescence and coherent scattering are ignored, and incoherent scattering is assumed to be governed by the Klien-Nishina cross section for free electrons.

The Monte Carlo code MCP [7] corrects some of the deficiencies in MCG. Designed to handle photons of energies 1 keV to 100 MeV, MCP provides for fluorescent emission, and modifies the Thomson and Klein-Nishina differential cross sections by scattering factors which take binding effects into account.

With the exception of a more sophisticated collision routine, the code MCP has much in common with MCG, containing the same geometry routine, variance-reducing techniques, and output. The distinctive features of the collision routine are summarized:

(1) The cross sections in the compilation of Storm and Israel [11] are used, with log-log interpolation, on the energy range .001-100 MeV.

(2) Incoherent scattering is governed by the function

$$K(\alpha,\mu)\ I(Z,\nu)\ d\mu \quad ,$$

where $K(\alpha,\mu)$ is the Klein-Nishina differential cross section and $I(Z,\nu)$ is a form factor correcting for bound effects. Here α is the incoming photon energy, Z is the atomic number of the scattering atom, μ is the cosine of the scattering angle, and ν is proportional to the momentum transfer to the electron. Sampling the incoherent scattering density is effected by a rejection technique, using the rapid sampling scheme [10] for the Klein-Nishina distribution described above for the code MCG.

(3) The coherent scattering density is proportional to the function

$$C^2(Z,\nu)\ T(\mu)\ d\mu \quad ,$$

where $C(Z,\nu)$ is a form factor modifying the energy independent Thomson cross section, $T(\mu) \sim (1 + \mu^2)$. A rejection scheme is also used to sample the coherent scattering.

Although the collision routine is much more complex in this code than it is in MCG, the running times are not greatly different for all but the simplest geometries.

MCNG - A COMBINED NEUTRON-GAMMA RAY TRANSPORT CODE

This Monte Carlo code [12] has the capability of transporting neutrons with the code MCN, producing gammas from neutron reactions by means of stored gamma production cross sections, and transporting gamma rays with the code MCG. It is designed to replace the separate codes MCN and MCG, and it may be run in three modes: (1) mode 0 - MCN only; (2) mode 1 - MCN and MCG (combined mode); and (3) mode 2 - MCG only.

When run in the combined mode 1, the code provides for the study of total heating or energy deposition in the medium by neutrons, photons, and products of neutron reactions. In a specific neutron reaction, the conservation of energy requires that

$$k + Q = H + N' + \Gamma \quad ,$$

where k is the kinetic energy of the neutron, Q is the Q-value of the reaction and depends only on the rest energies of the material particles involved, H is the total kinetic energy of the charged particles or nuclear products resulting, N' stands for the sum of the kinetic energies of all outcoming neutrons, and Γ' is the total energy of the gammas emitted. Partly because of the complexity of gamma production data and the existence of cross sections for materials consisting of natural mixtures of isotopes for which the component cross-section data are unknown, the code MCNG computes the heating per neutron collision,

$$\bar{H} = k + \bar{Q} - \bar{N}' - \bar{\Gamma} \quad .$$

$\bar{H}$ is the expected nuclear product heating per neutron collision and $\bar{Q}$, $\bar{N}'$, $\bar{\Gamma}$ are the corresponding expectations per collision of the rest energy change, of the neutron energy emitted, and of the gamma energy produced. An additional reason for obtaining the charged particle heating per collision is to reduce statistical fluctuations inherent in sampling the individual reactions. The expected heating per collision, $\bar{H}$, is listed on the same energy grid that the total reaction cross sections use.

Liberated photons may produce further heating from pair production, Compton scattering, and photoelectric absorption. This photon energy deposition is also given as the expected heating per photon collision. Similarly, the product neutrons will give rise to additional heating from further collisions. In transporting neutrons by MCN, the gammas produced from collisions with the medium are stored on tape, and they form a source which is processed by MCG at the end of the neutron cycle. The code is designed to cycle between these two codes until the required statistical accuracy is attained. The

heating tallies give the energy deposition per starting neutron. In addition to collision tallies, provision is made in the code to obtain heating in small regions by track length tallies.

Currently the gamma production cross sections are used in multigroup form (30 neutron groups and 12 gamma groups), even though our transport codes use a continuous energy dependence on cross sections. This has been a matter of convenience, since some multigroup data existed when the code was set up. Our present plans call for the gamma production cross sections to be given on the same energy points our neutron cross sections use, with the gamma output spectra to be given in much finer detail.

MCK - A MONTE CARLO CRITICALITY CODE

One of the codes which is a modification of MCN--with essentially the same input, collision routine, biasing techniques, and general geometry package--is the code MCK. This program is an iterative scheme, similar in many respects to the code KENO [13], for finding the multiplication in a static configuration containing fissile materials. From a given fission source, the code transports the neutrons, creating a new fission source. The procedure is iterative, using the latest distribution as the source for a new cycle until convergence is attained. The multiplication constant k_{eff} is estimated from the ratio of the neutron population in the $(n + 1)^{st}$ generation to that in the nth generation.

The source is self-corrective in the sense that if the user decides to use a source of N particles, then the code generates a new source of approximately N neutrons. Of course, the total weight started from the source at the beginning of each cycle remains constant. Since the code is adapted from MCN, it contains the thermalization routine as well as the energy-dependent fission treatment.

Simple to use and versatile because of its general geometry, this program has been applied to a wide range of problems, e.g., in the design of very complex reactors [14] as well as in safety calculations.

MCMG - MULTIGROUP MONTE CARLO

The continuous energy Monte Carlo codes have been, and will continue to be, the backbone of our Monte Carlo effort. However, there are some applications for which our present codes are somewhat limited. Consequently we have developed a multigroup code with three general problem areas in mind: (1) the treatment of adjoint problems not amenable to solution with MCNA, e.g., coupled neutron and neutron-induced gamma-ray transport problems; (2) calculations by both multigroup Monte Carlo and multigroup discrete ordinates methods, using identical cross section sets; (3) the simulation of particle transport with relatively simple cross-section input.

The programming of the multigroup code MCMG is based on our combined neutron-gamma code MCNG. Most of the input data, the geometry package, the biasing schemes, the majority of the output formats, and many of the control

cards are identical to their counterparts in MCNG. The code permits the multigroup cross sections to be input in a variety of standard formats. It will operate in either the forward or adjoint mode.

Various methods have been suggested in the literature for treating the angular scattering distribution after a group-to-group scattering event. Three options are provided in MCMG. The first uses the first N Legendre moments to construct a step function with N - 1 non-zero steps which conserves (in some cases, approximately) these N moments. The area under each non-zero step is chosen to be 1/(N - 1), so that the μ_i-endpoints of the non-zero steps are determined from the N moments. The second option is the method used in the MORSE code which allows the scattering to occur at discrete angles. The third option samples the scattering angle from a density function that is proportional to the absolute value of a truncated Legendre series. The last option can lead to negative weights and, from a Monte Carlo standpoint, is not usually a desirable mode. It is included as an option since it is a close analog to the discrete ordinates treatment.

MCGE - A COUPLED ELECTRON-PHOTON TRANSPORT CODE

One of the newer members of our family of Monte Carlo codes is an electron-photon transport code containing the general three-dimensional geometry package. This program goes to considerable lengths to include the physics of electron-photon interactions and the transport of both types of radiation. In the primary-photon mode, the effect of secondary electrons can be studied definitively. Cases in which secondary electrons are important include (1) response function calculations of detector systems if the range of the typical secondary electron is comparable to the dimensions of the system; (2) gamma-ray shielding calculations if bremsstrahlung is important--in a 6 m.f.p. lead shield irradiated by 10 MeV photons, bremsstrahlung increases the Goldstein-Wilkins exposure buildup factor by 60% and the dose albedo factor by 140%; and (3) gamma-ray heating problems if local energy deposition is a poor assumption (because the energy is transported by secondary electrons and bremsstrahlung). In the primary-electron mode, the code is useful in the study of electron transport and the production of bremsstrahlung. Applications include verification of experimental results from high-energy electron accelerators, laser fusion calculations, and energy deposition in tissue and bone for medical studies.

MCGE [15] treats the physics of electron and photon transport, including the generation of bremsstrahlung radiation, in considerable detail. The electron portion of the code is based on the procedures used by Berger and Seltzer in ETRAN [16] but includes several modifications and refinements not found in ETRAN--most apparent is the generalized geometry, mixed-media capability of MCGE.

Photon histories are followed with standard collision-by-collision Monte Carlo methods, accounting for Compton scattering, pair production, and photoelectric absorption. Photoelectric absorption is considered explicitly in order to include the effect of the photoelectron. Pair production in MCGE refers only to the absorption of a photon and the subsequent production of an electron-positron pair. Annihilation is a completely separate event, occurring only at the end of the positron's path (path and range are not equated)

and not at the point where the electron-positron pair was produced. If the positron escapes from the system, the annihilation quanta are not produced.

Histories of electrons (and positrons) are followed in MCGE by using multiple-scattering theories to describe the electron's behavior over a small increment of its path, not collision-by-collision since tens of thousands of collisions may be associated with each history. The electron energy loss (including straggling effects) over an increment of its path is determined by sampling the appropriate theories to account for ionization, bremsstrahlung, and the angular deflection. The production of bremsstrahlung in each increment is determined from corrected cross sections derived by using the Born approximation.

All generations of electrons and photons are accounted for in MCGE: delta rays, bremsstrahlung, characteristic x-rays, photoelectrons, Compton electrons, the electron-positron pairs, annihilation quanta, fluorescence radiation, and Auger electrons. The code requires about 300,000 decimal words of computer storage, and the problem of 10-MeV photons normally incident on a 6 m.f.p. lead slab requires about two hours of CDC-7600 computer time for 40,000 primary photon histories.

MCGB - A GAMMA CODE WITH BREMSSTRAHLUNG

The conventional gamma code MCG described earlier has been modified to include the effects of secondary electrons (i.e., bremsstrahlung) for situations where the dimensions of the system are much greater than the typical range of a secondary electron--e.g., in gamma-ray shielding calculations. MCGB includes the effects of bremsstrahlung, without going to the trouble and expense of actually following secondary electron histories, by using thick-target bremsstrahlung theory [17]. Only 50,000 decimal words of storage are needed (much of this is taken up by the geometry package--a one-dimensional version would need only a few thousand words) and the 40,000 histories of 10 MeV gamma rays in lead require only 10 minutes of CDC-7600 time. MCGB is basically a simple Monte Carlo code. Compton scattering, pair production, and photoelectric absorption are accounted for using the conventions of MCG (photoelectric absorption is not "weighted-out", however, since the photoelectron is needed for bremsstrahlung production). The energy and direction of the Compton recoil electrons are easily found from energy and momentum relations. The photoelectron is assumed to be produced straight ahead (a very good assumption if the energy of the electron is greater than 1 MeV) with energy equal to the energy of the absorbed photon, neglecting the binding energy. The energy distribution of the members in the electron-positron pair is assumed to be equally split, with each member having kinetic energy equal to $(k - 2\ m_o c^2)/2$, k being the energy of the incident photon. The emission angle of the pair is assumed to be $m_o c^2/T$ radians, where T is the kinetic energy of the electron or positron.

REFERENCES

1. E. D. Cashwell and C. J. Everett, *A Practical Manual on the Monte Carlo Method for Random Walk Problems*, Pergamon Press, New York, (1959).

2. E. D. Cashwell, J. R. Neergaard, W. M. Taylor, and G. D. Turner, "MCN: A Neutron Monte Carlo Code," Los Alamos Scientific Laboratory report LA-4751 (1972).

3. E. D. Cashwell and C. J. Everett, "Intersection of a Ray with a Surface of Third or Fourth Degree," Los Alamos Scientific Laboratory report LA-4299 (1969).

4. R. R. Johnston, "A General Monte Carlo Neutronics Code," Los Alamos Scientific Laboratory report LA-2856-MS (1963).

5. L. L. Carter and N. J. McCormick, "Coupled Sampling with the Monte Carlo Method in Neutron Transport Calculations," *Nucl. Sci. Eng.*, 39, 296-310 (1970).

6. L. L. Carter, "MCNA: A Computer Program to Solve the Adjoint Neutron Transport Equation by Coupled Sampling with the Monte Carlo Method," Los Alamos Scientific Laboratory report LA-4488 (1971).

7. E. D. Cashwell, J. R. Neergaard, C. J. Everett, R. G. Schrandt, W. M. Taylor, and G. D. Turner, "Monte Carlo Photon Codes: MCG and MCP," Los Alamos Scientific Laboratory report LA-5157-MS (1973).

8. E. F. Plechaty and J. R. Terrall, "Photon Cross Sections 1 keV to 100 MeV," Lawrence Livermore Laboratory report UCRL-50400, Vol. VI (1968).

9. C. J. Everett and E. D. Cashwell, "Approximation for the Inverse of the Klein-Nishina Probability Distribution Function," Los Alamos Scientific Laboratory report LA-4448 (1970).

10. C. J. Everett, E. D. Cashwell, and G. D. Turner, "A New Method of Sampling the Klein-Nishina Probability Distribution for All Incident Photon Energies Above 1 keV," Los Alamos Scientific Laboratory report LA-4663 (1971).

11. E. Storm and H. I. Israel, "Photon Cross Sections from 1 keV to 100 MeV for Elements Z = 1 to Z = 100," *Nuclear Data Tables* 7, 565-681 (1970).

12. Report in preparation.

13. G. E. Whitesides and N. F. Cross, "KENO - A Multigroup Monte Carlo Criticality Program," Union Carbide Corporation, Nuclear Division, Report CTC-5, Oak Ridge, Tennessee (1969).

14. C. W. Watson and E. D. Cashwell, "Monte Carlo Development and Applications in the Los Alamos Nuclear Rocket Program," *Proceedings of the Fourth International Conference on Reactor Shielding, Vol. 1*, Paris (Oct. 1972).

15. W. L. Thompson, "Gamma-Ray and Electron Transport by Monte Carlo," Doctoral Dissertation, University of Virginia, Charlottesville, Va. (Feb. 1974).

16. M. J. Berger and S. M. Seltzer, " Electron and Photon Transport Programs," National Bureau of Standards Report 9837 (June 1968).

17. C. R. Emigh, "Thick Target Bremsstrahlung Theory," Los Alamos Scientific Laboratory report LA-4097-MS (Mar. 1970).

DISCUSSION

Borgwaldt: I understand that your code MCNA is a continuous energy adjoint code. Could you give a few details? How does this code work? Of course, we all know about the problem of fitting an adjoint code into a continuous energy model.

Cashwell: Well, I am not absolutely sure about what it is that you want to go into here. The sampling is done essentially as is MCN in a reverse procedure, but the code uses a continuous model at all stages. The outcoming energy is sampled continuously and so is the scattering angle.

Kalos: There are several papers published on this subject. I had two in *Nuclear Science and Engineering* about 1969: "On the Integration of the Adjoint Gamma Ray Transport Equation," and "On the Monte Carlo Integration of the Adjoint Neutron Transport Equation". These papers explained in nauseating detail how one does the continuous sampling of an adjoint neutron problem.

Borgwaldt: Excuse me, may I reply to this directly? I was asking what type of strategy Carter is really using? That was my point. Does he use any special tricks?

Kalos: No tricks! If I remember, Carter does not calculate the kind of adjoint cross sections that are defined in our paper. He instead looks at the kernel and samples it in a more or less direct way using weights that derive from the normalization of the adjoint kernel.

Cashwell: Carter has a paper on this subject, besides the two papers to which Kalos referred. Carter and McCormick, I guess, had the first paper; and in his Los Alamos report, Carter goes into great detail on how all of this is carried out. I can give you a reference to it, if you would like, but he simply samples in every stage from a continuous distribution.

Gelbard: Is that a Los Alamos report to which you are referring?

Cashwell: Yes, there is a paper in *Nuclear Science and Engineering* by Carter and McCormick, and then the paper to which I am referring is a Los Alamos report which I think has more detail.

Gelbard: I have some brief questions about the resonance treatment. Your resonance cross sections come, I believe you said, thinned, from Don Harris' group?

Cashwell: That's right.

Gelbard: Does the MCK code then use a separate energy structure for each nuclide or does it combine them all into one?

Cashwell: No, we make no attempt to combine nuclides. The only code I have talked about that combines the energy mesh is the MCG Code — the gamma code. We can do this in MCG fairly easily because it treats high-energy photons, but all the neutron codes use a separate mesh for each nuclide.

Gelbard: Do you have any idea, off-hand, how many points are needed to treat ^{238}U? The reason that I ask is that it turns out to be a lot more difficult in our experience, to get an accurate cross-section representation, than we originally had expected.

Cashwell: Well, I really don't know but I would guess it's perhaps 15,000-20,000.

Steinberg: In the beginning you just barely mentioned the fact that your codes do include point and small volume detectors of some sort. I would like to know whether this is true; and could you describe whatever treatment that you have to keep the variance within reason?

Cashwell: Well, we don't have the Steinberg-Kalos treatment of the point detector, so that is a weakness as far as point detection is concerned. We have our difficulties, in some cases. There are a lot of detector problems where the detector is sitting some distance away from the collision point. If it is not, we may have some trouble. We have looked into the Steinberg-Kalos methods, and we are in the process of trying to implement some such methods, but have not done so yet. In some cases we do get average flux in a small region, but, as you know, that is not very feasible in some applications.

Bernnat: How can you treat light moderators in the thermalization process, light moderators in which the free gas model is not valid?

Cashwell: No, we have no provision for anything beyond the free gas model. Keep in mind, now, that this model is a little more applicable to some of our work. You're thinking perhaps of a more exact treatment using $S(\alpha,\beta)$ or something of that sort. No, we don't have that. Elliot Whitesides, didn't you look into the $S(\alpha,\beta)$ scheme?

Whitesides: Yes, we have an $S(\alpha,\beta)$ sampling scheme.

Cashwell: And how expensive is it, and how does it work?

Whitesides: The whole code is reasonably inexpensive. As a matter of fact, that's why I was a little bit concerned about the computing time required by the British code, because we find that treating everything in very much detail and using point data really costs us little more than the use of a multigroup form.

Cashwell: I agree with you about running time. Where we have compared our point code with our multigroup code, we find that, if the geometry is at all complicated, the time is essentially spent in the geometry and you can get very little benefit, I think, from using a multigroup code in those cases.

Kalos: Elliot, someone at KAPL or RPI published a paper on $S(\alpha,\beta)$ which involved, essentially, a preprocessing of everything, so that the sampling was very, very short. Is yours the same as that or similar?

Whitesides: Yes, we simply took this work and extended it somewhat for actual application in our code.

Kalos: Cady was one of the people involved.

Whitesides: I have forgotten the names now. Lester Petri of our group has done this work. I believe he gave a paper at Miami in which he reported some results on using this technique.

Kalos: How did it work?

Whitesides: It worked pretty well, although with the ENDF data, of course, we don't get very good results for normal systems. But we attribute that more to the data than to the method of treating cross sections, because we simply don't get very good results for thermal systems in any fashion.

Gelbard: I think it is worth mentioning that the alternative scheme, the older scheme for treating thermal neutrons, is simply a transfer matrix method. This scheme has been used in a number of places and is also a perfectly feasible one. But, Elliot, you say that you do not get very good results with thermal systems. You mean water-moderated systems, or some other kind of systems?

Whitesides: For PWR and BWR-type fuel elements and their environments. We have been looking mainly at shipping casket environments, and have been doing some criticals at Oak Ridge with some 5% metal rods that are about the diameter of a pencil. We don't get very good results for those experiments.

Gelbard: Did you mention a tamped system in your talk?

Cashwell: Yes.

Gelbard: What is a tamped system?

Cashwell: I mean by the term "tamper" a reflector of some sort.

Gelbard: You also mentioned something about a step function that preserved moments? Are you talking about a histogram treatment?

Cashwell: Yes, that's right. A histogram treatment which preserves the first n Legendre moments. Carter is still working on the method. He can only preserve the moments approximately, in some cases, but the method does, in the cases he's tried, give a step function which mimmicks the scattering distribution quite well. Of course, it has the advantage that everything is positive.

SESSION II

A DISCRETE APPROACH TO COMPLEX PLANAR GEOMETRIES

E. Cupini and A. De Matteis

Comitato Nazionale Energia Nucleare
Centro di Calcolo
Bologna, Italy

* * *

Presented by R. Simonini

A DISCRETE APPROACH TO COMPLEX PLANAR GEOMETRIES

E. Cupini and A. De Matteis

Comitato Nazionale Energia Nucleare
Centro di Calcolo
Bologna, Italy

ABSTRACT

Planar regions in Monte Carlo transport problems have been represented by a finite set of points with a corresponding region index for each. The simulation of particle free-flight reduces then to the simple operations necessary for scanning appropriate grid points to determine whether a region other than the starting one is encountered. When the complexity of the geometry is restricted to only some regions of the assembly examined, a mixed discrete-continuous philosophy may be adopted. By this approach, the lattice of a thermal reactor has been treated, discretizing only the central regions of the cell containing the fuel rods. Excellent agreement with experimental results has been obtained in the computation of cell parameters in the energy range from fission to thermalization through the ^{238}U resonance region.

INTRODUCTION

A discrete approach to planar complex geometries has been described in Ref. 1 and special algorithms for geometrical configurations which may be described by an overlay of convex polygons and circles have been given in Ref. 2.

Planar regions are represented in a computer memory by means of a set of points belonging to a square grid. A region index is associated to each point and the simulation of a particle free-flight between two grid points reduces to the simple operations necessary for scanning appropriate grid points to determine whether a region other than the starting one is encountered.

The geometrical part in the tracing of particle histories becomes in this approach independent of the complexity of the configuration examined, avoiding, too, the usual onerous tests to ensure correct correspondence between the particle position and the region index. On the other hand, it demans a large storage capacity, although only a few bits will be sufficient to encode the main information, i.e. the index of the region to which a point belongs.

When the complexity of the geometry is restircted to some regions only, a mixed discrete-continuous philosophy may be adopted by converting integer coordi-

nates into real ones (or vice versa) at a certain boundary (this conversion requires some adjustment to ensure that the new coordinates belong to the new region entered). In this way both memory economy and speed are obtained in large simple regions.

The question of the degree of approximation of this approach is partially answered by the excellent agreement with the experimental results obtained in the computation of cell parameters, in the energy range from fission to thermalization through the ^{238}U resonance region.

SCANNING ALGORITHM AND BOUNDARY CONDITIONS

The first operation to perform is the transformation of a geometrical configuration into a discrete map, trying to preserve the area of the regions and their geometrical shape as much as possible. The configuration is embedded in a rectangular background giving rise to a "boundary region" surrounding the configuration, so that when the image has been digitized in some way, the whole figure will be described by a rectangular matrix of elements. Those corresponding to the boundary region may contain information about the boundary conditions, while all others represent a square element of area h^2, h being the mesh size of the grid. Although any elementary-given figure (such as a circle or a rectangle) may be represented with the approximation $h^2/2$, the final represented area of the regions may have a larger approximation, due to the operation of overlaying elementary figures one upon the other. The small difference between the true and the represented areas may be taken into account in some final results (typical: the reaction rates) through a correction assuming linear dependence of the effects on the represented areas.

Once a discrete map has been prepared, given the starting point, the direction of motion and the free-flight of a particle, a scanning algorithm will determine the grid points visited in succession by the particle, which jumps from one grid point to another keeping as close as possible (in a sense later specified) to the ideal path in the continuum, until a new region is encountered. Let the mesh size h be the unit of length, x_0, y_0 the integer coordinates of the starting point, ℓ the free-flight, θ' the angle between the line of flight and the x-axis. The end-point coordinates of the free-flight are

$$x_f' = x_0 + \ell \cos \theta' , \qquad y_f' = y_0 + \ell \sin \theta' .$$

Let x_f and y_f be the integers resulting from the rounding off of x_f' and x_f'. If $(x_0,y_0) = (x_f,x_f)$, the path is terminated. Otherwise the direction of motion is corrected from θ' to θ with $\mathrm{tg}\ \theta = (y_f - y_0)/(x_f - x_0)$. In the case $0 \leq \theta \leq \pi/4$, the scanning procedure may be essentially described as follows. Denoting by [q] the integer part of q, the point visited at the n-*th* step is

$$x_n = x_{n-1} + 1 = x_0 + n$$

$$y_n = y_0 + [n\ \mathrm{tg}\ \theta + 0.5] ,$$

i.e. the y_n coordinate is the grid coordinate nearest to the ideal one,

$$y_0 + (x_n - x_0)\ \mathrm{tg}\ \theta\ .$$

Two types of direct transition from one grid point to another are thus possible: along one of the coordinate directions and along the diagonal of a grid cell. Therefore (for $0 \leq \theta \leq \pi/4$), for each x_i only one point is scanned (it must be noted that the direct transition along the diagonal of a grid cell could fail to detect the crossing of a new thin region only if this region is very poorly represented in the digitized image). Figure 1 shows an example of a path simulation.

Periodicity for infinite lattices of rectangular or hexagonal cells is taken into account by reflection on the cell boundary during path simulation. In the algorithm given in Ref. 2, this operation is very simply accomplished by moving the particle from one boundary to another according to information stored in the grid points corresponding to the boundary region, whenever this region is reached.

SOME RESULTS

The geometrical approach described was tested against experimental results concerning the lattice of a D_2O-moderated, H_2O-cooled reactor, with hexagonal cells at a pitch of 29 cm, at room temperature [3]. The fuel cluster was composed of 19 rodlets (diameter 2 cm) of natural UO_2, arranged in three rings of 1, 6, and 12 elements, respectively, contained in an aluminum pressure tube (internal diameter 10.6 cm) surrounded by an air annulus and a Zircaloy calandria tube (external diameter 12.11 cm). In this lattice the measurement of k_∞ was obtained by the null reactivity method (PCTR) and the neutron densities by activation of copper foils, with and without the coolant.

In the Monte Carlo computation only the part inside the pressure tube has been discretized with about 10^5 points, while the other regions were treated as a continuum; moreover the Zircaloy sheaths of the fuel elements (too thin to be realistically represented by a discrete image) were taken into account by a dilution in the coolant. In the energy range 16-MeV- keV, a 12-group approximation for cross sections (including anisotropic scattering and inelastic scattering) was adopted. In the energy region 50 keV-5 eV, 64 groups were considered, except for ^{238}U whose resonance cross sections were computed at each energy, with a random generation of resonance parameters for the unresolved resonances. Finally, below 5 eV, 256 energy points were used to describe cross sections with, moreover, two $33 \times 33 \times 10$ scattering matrices $\sigma\ (E' \rightarrow E, \cos\theta)$ for H_2O and D_2O in the thermal region below 1 eV.

In Table I, experimental and Monte Carlo results (based on 8000 histories) are quoted. The numerical results obtained, together with other positive comparison made for similar fuel assemblies of Canadian type, seem to encourage further experimentation with discrete geometry simulation.

REFERENCES

1. A. DE MATTEIS, "Digitizing Planar Regions for Monte Carlo Transport Problems," CNEN Report RT/FIMA(69)6.

2. A. DE MATTEIS and W. GUZIŃSKI, "Algorithms for Discrete Geometry in Monte Carlo Transport Problems," CNEN Report RT/FIMA(72)7.

3. A. GARAGNANI et al, CNEN Report CEC(70)12 (Bologna).

TABLE I

Average Neutron Densities in Various Regions and k_∞

	$\rho_{H_2O} = 1$		$\rho_{H_2O} = 0$	
	Experiment	MC	Experiment	MC
I ring	0.056	0.056	0.095	0.098
II ring	0.068	0.069	0.104	0.109
III ring	0.108	0.110	0.136	0.143
Coolant	0.124	0.124	-----	-----
Pressure tube	0.190	0.192	0.184	0.183
Calandria tube	0.199	0.195	0.192	0.187
Moderator	0.254	0.254	0.289	0.280
k_∞	1.054	1.057 ± 0.005	1.15	1.15 ± 0.005

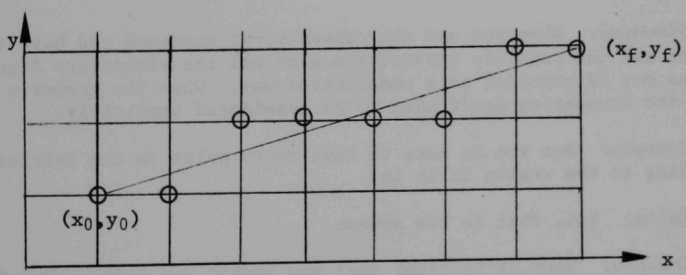

Fig. 1. A path simulation with tg θ = 2/7. The grid points scanned are circled.

DISCUSSION

Kalos: There is one thing that I have not understood here. I can see the advantages of discretizing the geometrical description, but I don't see the advantage of discretizing the position, or insisting that the particle be on one of the lattice points of your discretized geometry. Why not let the position be a continuous variable, and not approximate θ or θ'? The arithmetic is almost exactly the same. Secondly, would you repeat the ratio of the computing times for the geometrical tracking?

Simonini: The time to treat seven centered annular regions discretized to about 107 points was about 1.5 times that required by the usual approach, and the time was found to be roughly proportional to the square root of the number of grid points used.

Gelbard: You seem to be saying that it takes longer to treat the discretized geometry than to use the usual approach.

Kalos: Suppose you took a geometry that you were interested in, and discretized it to a degree that you thought was pretty good. Do you think it would run faster or slower?

Simonini: I don't know; it would depend on the geometry. If you add more elementary regions to the geometry, the tracking time using the discretized approach remains the same.

Gelbard: Are you also saying that one substantial advantage is in programming simplicity, not in the running time but in the fact that you can use one program to treat many different types of geometries?

Simonini: Yes. In any case when the geometry becomes very complicated I believe that my approach would be much faster than the usual one.

Kalos: Why?

Simonini: When you use the conventional approach you have to take into account all the possible intersections of all the elementary figures and these figures may be combined in a complicated way. When the geometry is discretized the boundaries don't have to be considered explicitly.

Coveyou: But you do have to have every point in the lattice labeled according to the region it is in.

Kalos: Yes, that is his scheme.

Coveyou: Then it is clear that you cannot use this method for a very large number of lattice points.

Simonini: No.

Kalos: And it would be very cumbersome for three-dimensional geometries.

THE ORGANIZATION OF CROSS-SECTION DATA IN THE MONTE CARLO CODE SPARTAN

R. C. Bending

Berkeley Nuclear Laboratories
U.K. Central Electricity Generating Board
Berkeley, Gloucestershire, England

THE ORGANIZATION OF CROSS-SECTION DATA IN THE MONTE CARLO CODE SPARTAN

R. C. Bending

Berkeley Nuclear Laboratories
U.K. Central Electricity Generating Board
Berkeley, Gloucestershire, England

ABSTRACT

The Monte Carlo code SPARTAN is a general-purpose code intended for neutron or gamma transport calculations. The code is designed to accept physics data from a number of external libraries (which may be used singly or in combination) and to use this data with as little alteration as possible.

Data obtained from one or several libraries is placed in an interface file on magnetic tape or disk, using a general hierarchical structure which allows particular data items to be assessed in a straightforward way. The interface file, with or without additional data from cards, is regarded as a data source for the main Monte Carlo calculation.

Within the Monte Carlo calculations itself, generality in the evaluation of cross sections and secondary particle parameters, as well as economy of coding, is achieved by the use of unified data structures for function evaluation and random sampling.

This paper includes a summary of the functional forms, sampling distributions, and particle interaction laws which are available at present, and indicates some of the mathematical methods used.

INTRODUCTION

One of the features of Monte Carlo techniques for particle transport calculations is the ability to use a very detailed physics data. Unlike most numerical methods, in which material cross sections are reduced to a multigroup form with fairly simple scattering laws, the cross sections and angle or energy distributions used in a Monte Carlo code may, in principle, be specified in great detail and in almost any functional form. For this reason, Monte Carlo techniques are useful for the evaluation of physics data against exprimental results, and for checking the validity of approximate physics models. At the other extreme, where multigroup data is well-established or

where it is required to check a numerical technique using consistent data, it is useful for a Monte Carlo code to accept multigroup physics data.

In the Monte Carlo code SPARTAN, which is being developed at Berkeley Nuclear Laboratories (United Kingdom Central Electricity Generating Board), this wide range of requirements is being met by the provision of a comprehensive physics data system designed to accept data from a number of different sources. Data from different source libraries is combined on an interface in a standard, but very general, hierarchical structure which reflects the data facilities required by the Monte Carlo calculation. All or part of this interface may be transferred into the computer core in order to carry out such a calculation. A unified approach to function evaluation and random sampling allows great generality as well as economy in coding. The system is designed to permit extension to other data libraries and formats as required.

In this paper a brief description is given of the data formats and structures used by the system and of some of the random sampling techniques employed. The structure of the physics data interface, used to store data on magnetic tape or disk, is also described, and a summary of the present capabilities and future development of the system is given.

FACILITIES REQUIRED BY A MONTE CARLO CALCULATION

During the course of particle tracking, a number of occasions arise when physics data is required. Whenever a particle enters a region, it is necessary to evaluate the total cross-section (and possibly the partial cross-section) for each of the nuclides which make up the material concerned. On the basis of this information, it is decided whether the particle has a collision in that region.

If the particle has a collision, further information is required. Partial cross-sections for all possible actions in the material of interest are used to select which type of event occurs. Further data is required to find the number of secondary particles produced. For each secondary, the number of available scattering laws is determined and a law chosen. Detailed data for this law is then employed to sample for a scattering angle and secondary energy.

It will be evident from the above paragraphs that there is a hierarchical structure implicit in the physics data used by a Monte Carlo calculation. This hierarchy, which is reflected by the way in which the data is structure in SPARTAN, is summarised in figure 1.

STANDARD DATA FORMAT FOR FUNCTION EVALUATION

At several points in the calculation it may be necessary to evaluate a function. These include:

a) Cross-section as a function of particle energy

b) Number of secondaries as a function of energy

c) Scattering angle as a function of energy

d) Secondary energy as a function of scattering angle.

In SPARTAN, a unified approach to function evaluation is adopted and a single routine is used for all the above situations. The functional forms which are included permit a wide range of specifications, and may be extended in a straightforward way as required. Consider, for example, a cross-section which is inversely proportional to particle velocity. Approximate representation of such a cross-section could be obtained using a multigroup approach or linear interpolation between point values. Alternatively, the cross-section could be described exactly by specifying log-log interpolation between the values at two points, or by describing the function as a general power series in energy.

The permitted functional forms and interpolation methods are summarised in Table I.

STANDARD DATA FORMAT FOR RANDOM SAMPLING

Similarly, there are several occasions on which random sampling is required. Of these, the selection of scattering angle and/or secondary energy are of great importance, and a wide variety of functional forms for angle and energy distributions may be found in different physics data libraries. The philosophy of SPARTAN is that, rather than risk the introduction of systematic errors by reducing these functions to a standard representation, this variety should be retained, and facilities provided within the code for sampling from frequency functions in many forms.

The functional forms at present permitted are shown in Table II. In most cases the sampling methods are trivial; those methods in which the random sampling technique is of some interest are described in Appendix A.

SECONDARY LAW SPECIFICATION

Once a particle has suffered an event, it is necessary to calculate the number, energy and direction of any secondaries produced. This may be done in a number of ways. For example, energy may be sampled randomly and angle calculated from it, or both may be sampled, and so on. In SPARTAN, this variation in the logic of secondary law specifications is retained, and several law types are available. These are listed in Table III. Provision is included for either energy **or** angle to be sampled, or for both to be sampled independently. Bivariate sampling, particularly useful for thermal neutron studies, is not yet available, but the data structures would allow it to be included without difficulty.

In many cases, the energy range of interest is divided into a number of 'law energy ranges' and a separate secondary law specification is given for each range. At any particular energy, a choice must be made between the law specification given for the upper end of the relevant energy range and that given for the lower end. In SPARTAN, there are three alternatives:

(a) For all energies in the range, the specification given for the lower end of the range is used.

(b) The law specification 'varies linearly' over the energy range. This means that if E_2 and E_1 are the upper and lower energy limites and E is the particle energy, then the specification given for the lower end of the range is used with probability $(E_2-E)/(E_2-E_1)$.

(c) The law specification 'varies logarithmically' over the energy range. The specification given for the lower end of the range is used with probability $\log (E_2/E)/\log (E_2/E_1)$.

THE SPARTAN PHYSICS INTERFACE

Data from various source libraries is combined to form an interface file which has a standard, though very general, structure. This file normally resides on magnetic tape or disk, and is used as the physics data source for the Monte Carlo calculation itself. It was noted above that the physics data has an implicit hierarchical structure, and this is reflected in the physics interface. Each record in the interface refers to information at the level below it in the hierarchy, and the use of a standard record structure enables particular parts of the interface to be located by a routine which 'knows' nothing of the meaning of the data. The standard record structure is shown in Table IV. The records are stored in the 'natural' order for a hierarchical file, which is best illustrated by the example shown in figure 2. The hierarchy used in the physics interface file is shown in figure 3, and the various types of record are described below.

Interface Level 1

This is the primary level of the interface file, containing the number of data source libraries and the interface title.

Source Level 2

There is one record for each data source library, containing the number of species used from that library, and information on energy group structures used in that source.

Species Level 3

This record gives the number of possible actions for the species concerned, a species code and the atomic mass.

Action Level 4

This record gives the number of data energy ranges and the action classification, which indicates whether the action is scatter, capture or fission.

Data Energy Range Level 5

The volume of data in some physics libraries is such that it is not possible to retain all relevant data in core at the same time. This is particularly true of neutron resonance data, as found in the FGL [1] and GENEX [2] libraries; when such a library is used it is necessary to bring data into core from time to time during particle tracking. In SPARTAN, this may be accomplished by dividing the energy range of interest into several data energy ranges. When the interface file is read in preparation for a Monte Carlo calculation, only the data for the first data energy range is retained in core, while the remainder is placed in temporary storage to be read as required. The introduction of data energy ranges is made at the action level in recognition of the fact that these large quantities of data are often encountered only for particular actions in particular nuclides, such as resonance capture in uranium-238. The data energy range record specifies the number of cross-section and law energy ranges and the energy limits of the data energy range. The facilities for using several data energy ranges have not yet been used in SPARTAN.

Cross-section Energy Range Level 6

The energy range of interest may be divided into a number of sections for the description of nuclide cross-sections. This is usually done to enable a different functional form to be used at different energies, and is particularly useful when a cross-section vanishes above or below a particular threshold energy. The cross-section energy range record contains the lower energy bound of the range and the cross-section, defined in the standard format for function evaluation. No reference is made to records at a lower level in the hierarchy.

Law Energy Range Level 6

The law energy range record gives the number of distinguishable secondaries, the lower energy bound of the range, a description of the average number of secondaries, and a parameter indicating the way in which the law data is selected for energies within the range (see section on secondary law specification).

Secondary Level 7

The secondary record give the number of laws which may apply for the secondary concerned.

Law Level 8

This record gives the probability, delay and energy release for this law, the co-ordinate system in which the angular distribution is given, the law type and law data (see section on secondary law specification).

For many problems, all the physics data required will be available in the interface file, and the user will need only to provide instructions to the program to locate the relevant data. Facilities are provided in SPARTAN, however, to combine data from the interface with data on cards (or indeed to use card data alone) in order to set up the files used in the main Monte Carlo calculation. Reference to the interface may be made at the species, action, data energy range or cross-section energy range level. Appendix B gives an example of the card data used to set up an interface file and to construct from it, with additional card data, a file for use in the Monte Carlo calculation.

PRESENT AND FUTURE DEVELOPMENT

At the time of writing, the SPARTAN physics system is able to use data from the FD5 [1], a multigroup neutron data library suitable for fast reactor calculations and GAMBLE [3], (a gamma data library) data sets. An earlier version of the system uses the thermal reactor neutron library associated with the WIMS-E program [4], and the DICE neutron library, which is intended primarily for shielding calculations [5].

Future development of the code will provide access to the FGL libraries [1] and possibly to ENDF/B data. It is anticipated that the range of data libraries which can be used, and the facilities provided within the physics system, will be extended from time to time, and the system has been designed to allow this.

CONCLUSIONS

A description has been given of the physics data system used by the Monte Carlo code SPARTAN. The system is designed to allow data to be obtained from number of different source libraries, and used in the Monte Carlo calculation with as little alteration as possible. The generality is achieved in an economic way through the use of standard systems for function evaluation and random sampling, and a formalised hierarchical structure for the relevant data files. Facilities exist at present to read both neutron and gamma data, and access to additional data libraries will be provided in the future.

ACKNOWLEDGEMENT

This paper is published by permission of the United Kingdom Central Electricity Generating Board.

REFERENCES

1. J.ROWLANDS, C.J.DEAN, J.D.McDOUGALL and R.W.SMITH, "The Production and Performance of the Adjusted Cross-Section Set FGL5", in Proc.Intl. Symp. on the Physics of Fast Reactors, Tokyo (1973) Paper A30.

2. R.J.BRISSENDEN and C.DURSTON, "A User's Guide to GENEX, SDR and related computer codes", UKAEA Rep. AEEW-R622, Atomic Energy Establishment, Winfrith (1968).

3. A.FODERARO, "Subroutine HEITLER", UKAEA Rep. AERE-M1956, Atomic Energy Research Establishment, Harwell (1967).

4. J.R.ASKEW, F.J.FAYERS and P.B.KEMSHELL, "A General Description of the Lattice Code WIMS", J.Brit. Nucl. Energy. Soc., 5, 4, p.564(1966).

5. J.B.PARKER, Ed., "Dice Mk V. The Preparation of Nuclear Data into a Form Suitable for Monte Carlo Calculations Using an Electronic Computer", UKAEA Rep. AWRE O-27/66, Atomic Weapons Research Establishment (1966).

APPENDIX A

SPECIAL RANDOM SAMPLING METHODS

Maxwellian, evaporation and Watt spectra

The Maxwellian, evaporation and Watt spectra are given by the respective formulas

$$m\ (E)\ dE = k_m E^{\frac{1}{2}} \exp\ (-a_m E) dE$$

$$v\ (E)\ dE = k_v E \exp\ (-\ a_v E) dE$$

$$w\ (E)\ dE = k_w \sinh\ (\sqrt{bE}) \exp\ (-\ a_w E) dE$$

where k_m, k_v, k_w, a_m, a_v, a_w and b are independent of E.

Noting that the Maxwellian frequency function is χ^2 frequency function with three degrees of freedom, we may sample for E by forming the sum of the squares of three random normal deviates and applying an appropriate scale factor.

Similarly, the evaporation spectrum is a χ^2 frequency function with four degrees of freedom.

Sampling from the Watt spectrum is more difficult, and we employ a rejection method. Consider the function

$$f(E) = k_w \sqrt{bE} \exp\ (-\ qE)$$

In order to construct a rejection test, we need to find a value of q such that for all positive E,

$$f(E) \geq w(E)$$

It is found that the value

$$q = a_w - b/6$$

has the required property, for with this value

$$f(E) - w(E) = k_w \exp(-a_w E)\left((bE)^{\frac{1}{2}} \exp(bE/6) - \sinh(bE)^{\frac{1}{2}}\right)$$

$$= k_w \exp(-a_w E)\left(\left(\frac{1}{2!6^2}\right) - \frac{1}{5!}\right)(bE)^{5/2} +$$

$$\left(\frac{1}{3!6^3} - \frac{1}{7!}\right)(bE)^{7/2} + \ldots)$$

$$\geq 0$$

The sampling procedure for the Watt spectrum is to select a value of E from the Maxwellian spectrum obtained by normalising f(E), f*(E), and accept it with probability

$$w(E)/f(E)$$

A typical Watt spectrum, and the corresponding bounding function f(E), are shown in figure 4.

Legendre Polynominal Series

Let f(u) be the frequency function of the cosine of the angle of scatter, and let f(u) be described by means of a series of Legendre polynominals,

$$f(u)du = \sum_{i=o}^{n} a_i \, Pi(u) \, du$$

In order to construct a rejection test, we seek a frequency function of similar shape to f(u) which can, by means of a constant factor, be converted to an upper bound for f(u).

Using the results

$$|Pi(u)| \leq \sqrt{2/(\pi i(1-u^2))}$$

$$Po(u) = 1$$

$$P_1(u) = u$$

We may write

$$f(u) \leq \frac{1}{\sqrt{1-u^2}} \sqrt{\frac{2}{\pi}} \sum_{i=2}^{n} |a_i|/\sqrt{i} + C_1 (1+u) + C_2 + C_3 (1-u)$$

$$= h(u), \text{ say,}$$

where

$$C_1 = \max (a_1, o)$$

$$C_2 = \begin{cases} a_o - a_1 & a_o \geq a_1 > o \\ a_o + a_1 & -a_o \leq a_1 < o \\ o & \text{otherwise} \end{cases}$$

$$C_3 = \max (-a_1, o)$$

writing $b = \sqrt{\frac{2}{\pi}} \sum_{i=2}^{n} |a_i| / \sqrt{i}$

the frequency function proportional to h(u) is

$$h^*(u) = \frac{b/\sqrt{1-u^2} + C_1 (1+u) + C_2 + C_3 (1-u)}{b\pi + 2C_1 + 2C_2 + 2C_3}$$

We regard h*(u) as the sum of four frequency functions,

$$\frac{1}{\pi\sqrt{1-u^2}} \quad \frac{1+u}{2} \quad \frac{1}{2} \quad \frac{1-u}{2}$$

with appropriate coefficients. In order to sample from h*(u), we select one of these frequency functions (the probability of selection being equal to the coefficient) and sample from it using elementary methods.

Having obtained a sample from h*(u), we accept it with probability

$$f(u)/h(u)$$

The frequency function f(u) and bounding function h(u) are shown for an notional angular distribution in figure 5.

APPENDIX B

EXAMPLE OF CARD INPUT USED TO READ PHYSICS DATA

(Explanatory comments are in brackets)

```
CREATE PHYSICS

DATA NEW UNIT 10                    (Interface is to be created on
                                     FORTRAN unit 10)

*  TEST GAMMA DATA                  (Interface title)

FROM UNIT 15 GAMBLE SPECIES 4       (GAMBLE library is on FORTRAN unit 15;
                                     four species are to be read)

6  7  92  1000                      (Species numbers in the library)

ENERGY 1.0E7 1.0E4 1 SOURCE 1 1     (Two energy structures; one from the
                                     interface, one given explicitly)

2  4  1.0E7 1.0E6 1.0E5 3.0E4 1.0E4

SPECIES

CARBON SOURCE 1  1                  (Carbon, uranium and pure Compton
NITROGEN 14.0  ACTIONS  3            scatter data are obtained direct
                                     from the interface.  For nitrogen
                                     the interface is entered at the
1  SOURCE  1  2  1                   action level, permitting an additional
2  SOURCE  1  2  2                   action, a fixed absorption specified
3  ABSORPTION DATA 1  1  0           in the card data, to be included)
   ENERGY 1.0E4 XSEC FIXED 0.3E-4
URANIUM SOURCE 1  3

COMPTON SOURCE 1  4

MATERIALS

1 CONTAINS 3 CARBON 0.01 NITROGEN 0.003 COMPTON 0.06

2 CONTAINS 2 URANIUM 1.0E-4 COMPTON 9.2E-3

(when this data has been read, the files used in the main calculation
will have been set up)
```

TABLE I

STANDARD FORMATS FOR FUNCTION EVALUATION

FORMAT NUMBER	FUNCTIONAL REPRESENTATION	
1	f(x) is independent of x	
2	f(x) is tabulated at equal intervals	Interpolation between x values may be:
3	f(x) is tabulated at equal intervals in log x	f(x) constant (value given at lower bound) f(x) linear in x
4	f(x) is tabulated at intervals defined in a separate array	f(x) linear in log x log f(x) linear in x
5	f(x) is tabulated as (x,f(x)) pairs	log f(x) linear in log x
6-10	Not yet allocated	
11	f(x) is a polynominal in x	
12	f(x) is a polynominal in log x	
13-20	Not yet allocated	
21	f(x) is a series of real powers in x	
Over 21	Not yet allocated	

TABLE II

STANDARD FORMATS FOR RANDOM SAMPLING

FORMAT NUMBER	SAMPLING DISTRIBUTION
1	x has a fixed value
2	x is uniformly random on (-1, 1)
3	x is uniformly random between given limits
4	x is given by a histogram with equal intervals
5	x is given by a set of equi-probable intervals
6	x is given by a general histogram whose intervals are defined in a separate array
	The frequency function of x within each interval may be constant or proportional to 1/x
7-10	Not yet allocated
11	x is given by a Legendre polynominal series
12	x is given by a Maxwellian spectrum $f(x) \alpha \sqrt{x} \exp(-ax)$
13	x is given by an evaporation spectrum $f(x) \alpha x \exp(-ax)$
14	x is given by a Watt spectrum $f(x) \alpha \sinh \sqrt{bx} \exp(-ax)$

TABLE III

LAW TYPES

LAW TYPE NUMBER	DESCRIPTION
1	Angle sampled from standard disribution; secondary energy computed from angle by conservation of momentum.
2	Angle sampled from standard distribution; secondary energy a standard function of angle.
3	Angle sampled from standard distribution; secondary energy sampled independently from standard distribution
4	Angle sampled from standard distribution; secondary energy a standard function of primary energy.
5-10	Not yet allocated
11	Energy sampled from standard distribution; angle computed from energy by conservation of momentum.
12	Energy sampled from standard distribution; angle a standard function of secondary energy.
13-20	Not yet allocated.
21-30	Allocated for bivariate distributions.
31	Angle sampled from standard distribution; secondary energy computed from angle by Compton scattering law.
Over 31	Not yet allocated.

TABLE IV

STANDARD RECORD STRUCTURE ON PHYSICS INTERFACE

ITEM	MEANING
1	Number of branches from this record to the next level in the hierarchy.
2	Length of data block (in words).
3	Data block (containing all other physics information).

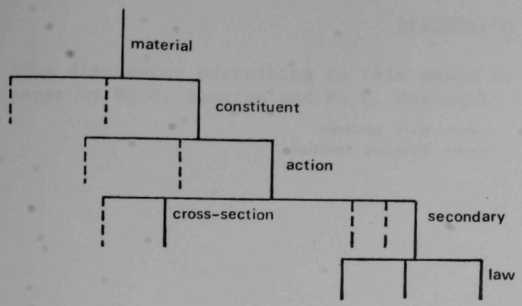

Fig. 1

Hierarchy of physics data used in Monte Carlo calculation.

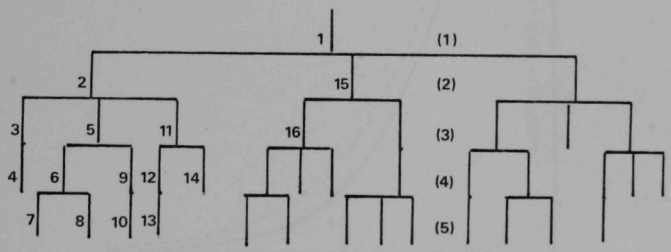

Fig. 2. Example of hierarchical structure.

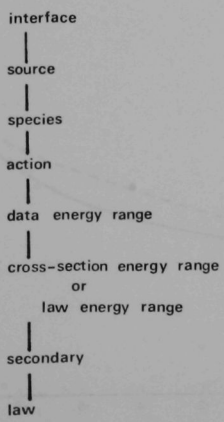

Fig. 3

Hierarchy used in the physics interface file.

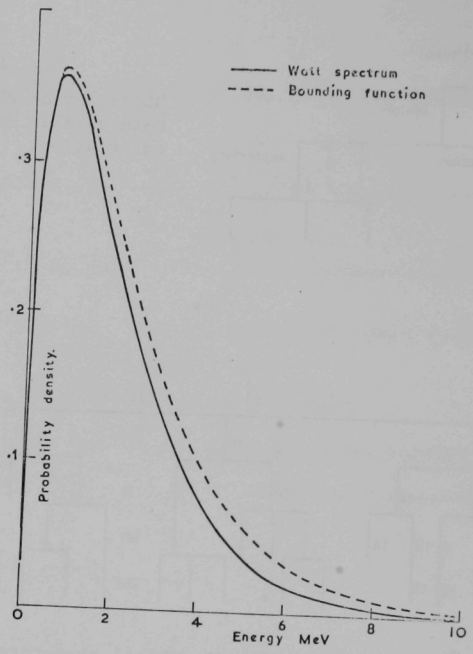

Fig. 4. Sampling from a Watt spectrum.

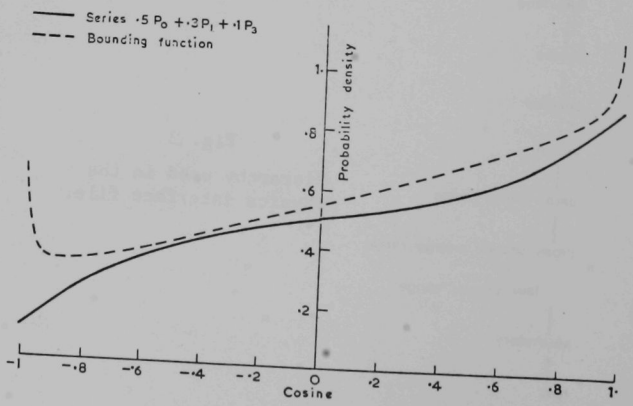

Fig. 5. Sampling from a Legendre polynomial series.

DISCUSSION

(The discussion pertaining to this paper is included with that for the next paper by R. C. Bending and P. G. Easter.)

THE GEOMETRY SYSTEM USED IN THE GENERAL MONTE CARLO TRANSPORT CODE SPARTAN

R. C. Bending and P. G. Easter

Berkeley Nuclear Laboratories
U.K. Central Electricity Generating Board
Berkeley, Gloucestershire, England

THE GEOMETRY SYSTEM USED IN THE GENERAL MONTE CARLO TRANSPORT CODE SPARTAN

R. C. Bending and P. G. Easter

Berkeley Nuclear Laboratories
U.K. Central Electricity Generating Board
Berkeley, Gloucestershire, England

The paper describes the geometry routines used in the general-purpose, three-dimensional particle transport code SPARTAN. The code is designed to deal with the very complex geometries encountered in lattice cell and fuel handling calculations, health physics, and shielding problems. Regions of the system being studied may be represented by simple shapes (Spheres, cylinders, and so on) or by multinomial surfaces of any order, and many simple shapes may be combined to make up a complex layout.

The geometry routines are designed to allow the program to carry out a number of tasks (such as sampling for a random point or tracking a path through several regions) in any order, so that the use of the routines is not restricted to a particular tracking or scoring method. Routines for reading, checking, and printing the data are included.

INTRODUCTION

SPARTAN (System for PARticle Tracking and ANalysis) is a general-purpose three-dimensional Monte Carlo code capable of tracking neutrons or photons. This paper describes the SPARTAN geometry routines, which are to a great extent independent of the rest of the code, and which allow very complex geometrical layouts to be described.

One of the features of Monte Carlo techniques, as opposed to deterministic methods, is their applicability to geometries of arbitrary complexity. A general purpose Monte Carlo code must therefore contain routines capable of accepting a description of a complex geometry and of tracking particle trajectories through it. These tasks are carried out in SPARTAN by a set of routines which are largely modular in construction, and could therefore be of use in other Monte Carlo codes. The way in which the routines are called allows the user considerable freedom in his choice of tracking or scoring techniques.

In most nuclear reactor systems, complex geometrical layouts are usually made up of many simple shapes, such as cylindrical fuel pins. For this reason much of the SPARTAN geometry system is intended to deal with simple shapes, such as spheres and prisms of various cross-sections. Objects having one of these shapes (known as geometric units) may be subdivided to form, for example, a set of concentric spheres. Geometric units may contain other geometric units in any positions and with any orientation. This facility allows complex layouts, such as a lattice of fuel pins with control elements, or even a whole reactor, to be described economically.

Despite the generality of the system provided by the standard shapes, circumstances do arise in which the layout cannot conveniently be described by this means. For such situations, the SPARTAN geometry system allows the surfaces of physical regions to be described by any number of multinomial surfaces of (in principle) any order. This makes it possible to describe, for example, a human phantom, or a venturi region in a cylindrical pipe.

The facilities which are provided by the geometry routines (apart from slight restrictions where multinomial surfaces are used) include the selection of uniformly random points in a region or on a boundary, the location of intersections of a particle trajectory with a set of regions(such as concentric annuli in a cylinder), evaluations of volumes and areas, calculation of a vector normal to a boundary, determination of the location of a given point with respect to a co-ordinate mesh (used when scores are required in great geometric detail) and reflection and translation operations for use at the boundary of the geometry layout itself. The system also includes routines which read geometry data in a comprehensible form, print it, and check it for internal errors.

REPRESENTATION OF THE GEOMETRY LAYOUT

The geometry layout is described by specifying the surfaces of every physical region. Regions are usually grouped into geometric units (see below). Geometric units may themselves be grouped to form structures (e.g. ring), and may be contained within one another.

A standard geometric unit is a group of nested surfaces having a common shape, orientation and origin, and differing in only one size parameter. The shape of the unit is one of the standard shapes described below. Examples of standard geometric units are a set of concentric spheres or of coaxial cylinders. The regions which make up the geometric unit are distinguished by their nesting level within the unit; the outermost region is at level 1, the next at level 2, and so on.

The innermost region of a standard geometric unit may itself contain other geometric units in any position and orientation. The entire geometry layout may be defined in terms of nested units in this way. The outer boundary of the layout (which may be fictitious) must be the outer surface of a standard geometric unit.

If some of the units contained within the innermost region of any geometric unit form a particular structure, either a right-angled lattice or a ring, they may be described as such in the data input, giving savings of both storage and data. Facilities are also provided to generate a random

distribution of contained units (all units have the same orientation, and their origins have a uniformly random distribution within the innermost region of the containing units, with a minimum separation between origins imposed on the distribution).

A non-standard geometric unit is one which is created by a group of multinomial surfaces, of which further details are given below. A non-standard unit is treated within the program in much the same way as a standard unit, except that it must be contained within a standard unit and cannot contain any further units. The way in which standard and non-standard geometric units are used to describe a complex layout is illustrated by Figures 1 to 4. Figure 1 shows a square reactor lattice cell with a central cylindrical channel containing a cluster of pins. Three geometric units (the outer block, the cylindrical channel and a fuel pin) are sufficient to describe the layout.

Figure 2 shows a model of a group of buildings, which might be used to assess the effect of radiation from airborne contamination. The use may be noted of intersecting regions (to describe the hemispherical dome) and the orientation of standard units (for the horizontal cylinders and the building set at an angle).

An example using multinomial surfaces is shown in Figure 3, which represents a simple human phantom as might be used for health physics studies. The head, trunk and legs of the phantom are described by second order equations. In order to truncate the surfaces at appropriate planes, each surface could be enclosed by a (physically non-significant) cylinder, and the layout completed by specifying a standard unit to provide an outer boundary.

Figure 4 shows a venturi region of a cylindrical pipe. In this case a single high-order multinomial surface is required. The card data required by SPARTAN to describe each of these geometry layouts is given in Appendix A.

Several standard shapes for geometric units are available. Table I gives the shape names and definitions, the size parameters used to describe them, the size parameters which may be varied to produce a nested geometric unit and the positions of the local origin and co-ordinate axes for each shape. Only one size parameter may be varied within one geometric unit. The standard shapes and permitted nesting are shown in Figure 5.

Coding is provided in the shape routines (see following section) for various tasks which may be required for Monte Carlo calculations. Details of the tasks are given in Tables II to V.

Any of the standard shapes may be used for the boundary of the whole layout. Apart from some restrictions noted in Table II, the boundary conditions may be specified separately at each boundary face as either periodic transmission, mirror reflection or diffuse reflection.

Circumstances arise in which it is not possible (or not convenient) to describe the layout solely in terms of the standard shapes listed in Table II. In this case a more general representation is available. A geometric unit consisting of several disjoint regions may be created by a number of surfaces expressed as multinomials of any order in the Cartesian co-ordinates (x,y,z).

For data input purposes, general surfaces are described in a form directly resembling the polynomial concerned (see examples in Appendix A). Each disjoint region created by the group of surfaces is identified to the computer by specifying a point which the region contains.

The specification of the geometry layout may allow two or more geometric units to intersect. This is permitted in SPARTAN, but certain conventions are used to determine in which of several interesecting regions a point lies. These conventions are applied in the tracking routines of SPARTAN, outside the geometry system, and their operation is beyond the scope of this report. The conventions are as follows:

(1) Any geometric unit is truncated by the surface of the innermost region of the geometric unit which contains it.

(2) If two or more units which are contained within a single unit intersect one another, then any point common to the two units is assumed to be in the unit which was specified earlier in the input data.

(3) If a unit intersects another unit other than one of its contained units, then it may not contain any units. (In other words, a unit which is truncated must be empty).

Intersecting regions are particularly useful when a plane of symmetry cuts some of the regions of a layout, or where compound shapes (such as a cylindrical channel with slotted walls) are required.

ROUTINES IN THE GEOMETRY SYSTEM

Input of geometry data is performed using the routines GEOIN, GUNIT (which is called from GEOIN), and POLYIN (which is called from GUNIT). This group of routines reads data either from cards, in a free-format code-word-controlled form, or in binary form from a disk or magnetic tape. For examples of the data input, see Appendix A.

The routine GEOUT either gives a printed summary of the data or dumps the data file onto an external device for subsequent reading by GEOIN.

Detailed geometry calculations are carried out by a set of shape routines, one for each standard shape (SPHERE, CYLIND, TRIANG, CUBOID, HEXAGO), and one for groups of multinomial surfaces (SURFAC). This division enables certain routines to be omitted from a load module which does not require them, leading to a saving in core space. For each shape, a set of tasks is available. Table II lists the tasks, noting which tasks are available for each shape routine. Most of the calculations are straight forward; Appendix B gives details of any special techniques used.

The geometry system contains a number of other routines which are used to obtain or check data, carry out rotations or boundary transformations, evaluate multinomial expressions and carry out random sampling tasks. Details of non-trivial methods are given in Appendix B.

CONCLUSIONS

The geometry system of the SPARTAN Monte Carlo code allows a wide range of geometric layouts to be described, including several commonly-encountered shapes and more general shapes for unusual geometries. No restrictions are placed on the position or orientation of geometric objects, and, subject to certain conventions, geometric shapes may intersect one another. The routines which make up the system are in most respects independent of the rest of the code, and may readily be adapted for use outside SPARTAN.

ACKNOWLEDGMENT

This paper is published by permission of the United Kingdom Central Electricity Generating Board.

APPENDIX A: EXAMPLES OF DATA INPUT

Channel and pin cluster (Figure 1)

```
PIN CYLINDER RADIUS NEST 3  1.50  1.30  1.25 HEIGHT 1000.0 EMPTY
CHAN CYLINDER RADIUS 20.0 HEIGHT  1000.0 CONTAINS 3
     RING 6.0  1.570796  4 PIN
     RING 10.0  1.047198 6 PIN
     RING 16.0  0.523599  12 PIN
CELL CUBOID SIDES  48.0  48.0  1000.0 CONTAINS 1
     AT  24.0  24.0  0.0 CHAN
```

Group of buildings (Figure 2)

```
TANK CYLINDER RADIUS 2.0 HEIGHT  6.0 EMPTY
SILO CYLINDER RADIUS 8.0 HEIGHT  16.0 EMPTY
DOME SPHERE RADIUS  8.0 EMPTY
BLOCK CUBOID SIDES 40  10  10  EMPTY
OUT CUBOID SIDES 200  200  50  CONTAINS 6
     AT 91  95  2 ANGLE  0.0  -1.5708  0.0  TANK
     AT 96  95  2 ANGLE  0.0  -1.5708  0.0  TANK
     AT 112  90  0  SILO
     AT 112  90  16 DOME
     AT  50  110  0 BLOCK
     AT 100  120  0 ANGLE  -0.524  0.0  0.0  BLOCK
```

Human Phantom (Figure 3)

```
HEAD SURFACE  1  100 X (2) + 49Y(2)  =  4900
     REGIONS  2  CONTAINS  0   0   0
                 CONTAINS  100  100  100
CYLH CYLINDER RADIUS 10  HEIGHT 24  CONTAINS 1 HEAD
TRUNK SURFACE 1  1 X (2) + 4Y(2)  =  400
     REGIONS  2  CONTAINS 1  1  1
                 CONTAINS 10 10 10
CYLT CYLINDER RADIUS 20 HEIGHT 70 CONTAINS 1 TRUNK
     LEGS SURFACE 1 25 X (2) + 100Y(2) - 1Z(2) - 40Z = 400
     REGIONS 2 CONTAINS -1  -1  -1
               CONTAINS 100  100  100
```

```
CYLL CYLINDER RADIUS 20 HEIGHT  80 CONTAINS 1 LEGS
OUT SPHERE RADIUS 100  CONTAINS 3
     AT  0  0  70  CYLH
                   CYLT
     AT  0  0  -80 CYLL
```

The head, trunk and legs are specified as second order multinomial surfaces; each is surrounded by a fictitious cylinder which truncates it vertically. The whole phantom is surrounded by a sphere which provides the boundary of the layout.

Cylindrical venturi (Figure 4)

```
VENT SURFACE 1  1.0X(2) + 1.0Y(2)-4.0E-8Z(8) + 1.6E-5Z(6)-4.0E-4Z(4)
                -0.24Z(2) = 9
REGIONS 2 CONTAINS 0  0  0
          CONTAINS 4  0  0
CYL CYLINDER RADIUS 5.0  HEIGHT 20.0  CONTAINS 1
     AT  0.0  0.0  10.0  VENT
PIPE CYLINDER RADIUS 5.0  HEIGHT 100.0  CONTAINS 1
     AT  0.0  0.0  50.0  CYL
```

APPENDIX B: SPECIAL TECHNIQUES USED BY THE GEOMETRY ROUTINES

Sampling for a random point in a spherical shell

Let the inner and outer radii of the shell be r_1 and r_2. Then the frequency function of the radius of a uniformly random point is:

$$f(r)dr = 3r^2dr/(r_2^3 - r_1^3)$$

This curve is bounded above (see Figure 6) by

$$g(r) = 3r_2r/(r_2^3 - r_1^3),$$

which is proportional to the frequency function

$$g^*(r) = 2r/(r_2^2 - r_1^2)$$

In order to select r from f(r), we select r from $g^*(r)$ using RANRAD (see below), and accept it with probability

$$p = f(r)/_{g(r)}$$

i.e. $$p = r/r_2$$

The efficiency of this rejection test (number of trials/acceptable value) is

$$2(r_1^2 + r_1r_2 + r_2^2)/(3r_2(r_1 + r_2)),$$

which has a minimum value, when $r_1 = 0$, of 2/3.
The angular co-ordinates of the point are found using RANG3D (see below)

Sampling for a random point in a triangle

Triangular shapes in the SPARTAN geometry system are always orientated with respect to local co-ordinate axes so that the base of the triangle lies along the local x axis (see Figure 7). A pair of uniformly random points are chosen on the base AB, and from each of these is constructed a line parallel to the side to which it is nearest (i.e. the right-hand point is nearest to side BC, the left-hand to side AC). The intersection of these lines is then uniformly random over the area of the triangle.

Sampling for a random point in the area between two hexagons

Figure 8 shows two concentric hexagons, and we need to select a uniformly random point from the area between them. The method adopted is to construct the rectangle shown along the uppermost side, and select a point in the rectangle. The point is accepted provided it does not lie in the shaded portions shown. This rejection test has an efficiency

$$(r_1 + r_2)/2r_2$$

where r_1 and r_2 are the inner and outer radii, and gives a minimum efficiency, when $r_1 = 0$, of 1/2.

A uniformly random integer between 0 and 5 is selected, and is used to rotate the sample point, if required, to another segment of the hexagon.

Calculation of the intersection of a given trajectory with a group of surfaces.

The method used in SURFAC to calculate the distance along a trajectory to a particular surface depends on the order of the multinomial expression used to describe the surface. For order one or two, elementary algebraic methods are used. For higher order, a Newton-Raphson iterative technique is used to find all the roots of the appropriate equation. As each real root or pair of complex roots is found, the order of the polynomial is reduced appropriately, so that exactly the correct number of roots is found. Only real, positive roots are retained as meaningful intersections.

Calculation of the normal to a given multinomial surface

Use is made of the fact that the surface $\phi(x,y,z) = 0$ has a normal given by the vector

$$\Delta\phi(x,y,z).$$

Rotational transformation (routine GEROT)

The routine carries out the transformation of a point (given with respect to a right-handed set of Cartesian axes) due to a rotation of the axes through specified Euler angles.

The total rotation is brought about by three successive rotations through angles, α, β, γ (see Figure 9) in that order. The rotations are -

(i) about the original z-axis
(ii) about the new x-axis
(iii) about the new z-axis,
and the axes rotate in a positive sense for positive Euler angles.

The matrix describing the transformation is

$$\begin{pmatrix} \cos\alpha\cos\gamma - \sin\alpha\cos\beta\sin\gamma & \sin\alpha\cos\gamma + \cos\alpha\cos\beta\sin\gamma & \sin\beta\sin\gamma \\ -\cos\alpha\sin\gamma - \sin\alpha\cos\beta\cos\gamma & -\sin\alpha\sin\gamma + \cos\alpha\cos\beta\cos\gamma & \sin\beta\cos\gamma \\ \sin\alpha\sin\beta & -\sin\beta\cos\alpha & \cos\beta \end{pmatrix}$$

which premultiplies the co-ordinate vector of the point.

Mirror reflection (routine MIRROR)

The routine is supplied with the direction cosines of the normal to a surface and the direction cosines of the flight of an incident particle. The direction cosines of the particle flight after mirror reflections are calculated using the relations

$$u_r = u_i - u_n\alpha$$

$$v_r = v_i - v_n\alpha$$

$$w_r = w_i - w_n\alpha$$

where $\alpha = 2(u_i u_n + v_i v_n + w_i w_n)$,

(u_i, v_i, w_i) are the incident direction cosines,

(u_n, v_n, w_n) are the normal direction cosines,

(u_r, v_r, w_r) are the reflected direction cosines. (See figure 10)

Sampling for angles (routines RANG and RANG3D)

The routine RANG samples for the cosine and sine of an angle uniformly random on $(0,2\pi)$. The von Neumann method is used. Pairs of uniformly random numbers ξ_1 between -1 and 1 and ξ_2 between 0 and 1 are selected until a pair is obtained for which

$$\xi_1^2 + \xi_2^2 \leq 1$$

Then $$u = \cos\theta = \frac{\xi_1^2 - \xi_2^2}{\xi_1^2 + \xi_2^2}$$

and $$v = \sin\theta = \frac{2\xi_1\xi_2}{\xi_1^2 + \xi_2^2}$$

The speed of this method is assisted by the fact that no square root operation is required.

The routine RANG3D samples for the direction cosines of a uniformly random direction is three dimensions. The routine RANG is used to find direction cosines relative to the x and y axes, and the third direction cosine is taken to be uniformly random on (-1,1).

Sampling for a uniformly random point in an annulus (routine RANRAD)

Two numbers ξ_1 and ξ_2, uniformly random on (0,1), are selected, and the radius is set using

$$r = \max \{((r_2 - r_1)\, \xi_1 + r_1),\ ((r_2 + r_1)\, \xi_2 - r_1)\}$$

where r_1 and r_2 are the inner and outer radii.

This formula is derived as follows. The frequency function of r is

$$f(r)dr = 2r/(r_2^{\,2} - r_1^{\,2})\, dr$$

with distribution function

$$F(r) = \frac{r^2 - r_1^{\,2}}{r_2^{\,2} - r_1^{\,2}}$$

Now,

$$\text{prob } [\max\{((r_2 - r_1)\, \xi_1 + r_1),\ ((r_2 + r_1)\, \xi_2 - r_1)\} \le r]$$

$$= \text{prob } [\ (r_2 - r_1)\, \xi_1 + r_1 \le r \text{ and } (r_2 + r_1)\, \xi_2 - r_1 \le r]$$

$$= \text{prob } [\xi_1 \le \frac{r - r_1}{r_2 - r_1}] \times \text{prob } [\xi_2 \le \frac{r + r_1}{r_2 + r_1}]$$

$$= \frac{r - r_1}{r_2 - r_1} \qquad \frac{r + r_1}{r_2 + r_1}$$

= F(r) as required.

Diffuse reflection (routine WHITE)

The routine WHITE generates the direction cosines of a particle flight path after diffuse reflection at a surface with given normal direction. By diffuse reflection, we mean that the cosine α of the angle between the flight path and the normal has a frequency function proportional to α.

The direction cosines relative to axes tangential to the surface, β and γ are found using RANG, while α is found as the larger of two uniformly random number on (0, 1) (see Figure 11). If the normal to the surface is (u_n, v_n, w_n), then the reflected direction cosines (u_r, v_v, w_r) are given by

$$u_r = \sqrt{\frac{1 - \alpha^2}{1 - w_n^2}} \{u_n w_n \gamma + v_n \beta\} + u_n \alpha$$

$$v_r = \sqrt{\frac{1 - \alpha^2}{1 - w_n^2}} \{v_n w_n \gamma - u_n \beta\} + v_n \alpha$$

$$w_r = - \sqrt{(1 - \alpha^2)(1 - w_n^2)} \gamma + w_n \alpha$$

TABLE I

Standard shapes

Shape code	Shape name	Description	Size parameters	Nesting	Local origin and co-ordinate axes
1	Sphere	Sphere	Radius r	r	Centre of sphere. No restriction on axes.
2	Cylinder	Right circular cylinder	Radius r Height	r or h	Centre of base of cylinder. Positive z axis along axis of cylinder.
3	Triangle	Triangular prism	Sides a,b,c Height h	h	One corner of base. Positive z axis parallel to axis of prism. Positive x axis along edge a of base. Triangle lies towards (y > 0) side of x axis.
4	Cuboid	General cuboid	Sides a,b,c	a or b or c	One corner of base. Positive x, y and z axes along edges, a, b, c of cuboid.
5	Hexagon	Hexagonal prism	Radius r, Height h	r or h	As cylinder. The x axis cuts a vertex of the base.

TABLE II

Tasks performed by shape routines

Task	Description	Notes
1	Find a uniformly random point in a given region of a given geometric unit	Not available in SURFAC
2	Find whether a given point lies in a given geometric unit	Not available in SURFAC
3	Given the direction cosines and starting position of a line, find all the positive distances along the line to intersections with the regions in a given geometric unit or group of surfaces	
4	Find the total surface area of a given geometric unit or group of surfaces	In SURFAC, an error is generated unless the appropriate values have been included as input data
5	Find the total surface area of a given geometric unit or group of surfaces	
6	Given a point close to the surface of a region, find the inward pointing normal at that point	In SURFAC, the direction of the normal is defined by the vector grad ϕ, where ϕ is the multinomial expression
7	Find a uniformly random point on all or part of the boundary of a given region	Not available in SURFAC
8	Given a co-ordinate mesh by which a region is subdivided, and a point in the region, find in which subdivision the point lies	See table III
9	Given a point close to the surface of a geometric unit, find the corresponding boundary type and adjust the point and direction cosines accordingly	Not available in SURFAC see tables IV and V. Periodic boundary conditions cannot be used for vertical faces in TRIANG

TABLE III

Subdivision of regions for scoring

Shape	Co-ordinate system in which subdivision are made, in conventional order.
Sphere	r, cosθ, ϕ
Cylinder	r, θ, z
Triangle	*, *, h (*indicates no subdivision permitted)
Cuboid	x, y, z
Hexagon	*, θ, z (No. of θ divisions must be 1 or 6)

TABLE IV

Boundary conditions used in shape routines task 9

Shape	Use
1	Periodic boundary condition. Direction cosines unchanged, position changed to appropriate point.
2	Specular reflection boundary condition. Position unchanged, direction cosines undergo mirror reflection.
3	Diffuse reflection boundary condition. Position unchanged, direction cosines random with inward-facing cosine distribution.

TABLE V

Conventional ordering of boundary faces

Shape	Order (bracketed items included for task 6 and 7 but not task 9)
Sphere	(Inner radius), outer radius.
Cylinder	(Inner radius, outer radius, base, top.
Triangle	Base of triangle, second side, third side, base of prism, top of prism.
Cuboid	$x = 0$, $x = a$, $y = 0$, $y = b$, $z = 0$, $z = c$.
Hexagon	(6 faces of inner surface), 6 faces of outer surface, base, top. (Note: vertical surface count anticlockwise from x axis).
Surface	Order defined by the order in which the surfaces are specified.

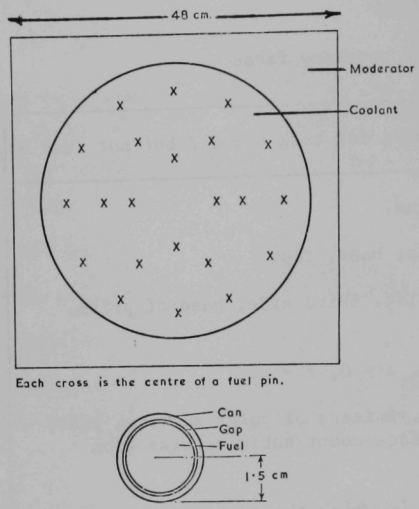

Fig. 1

Pin cluster in cylindrical channel.

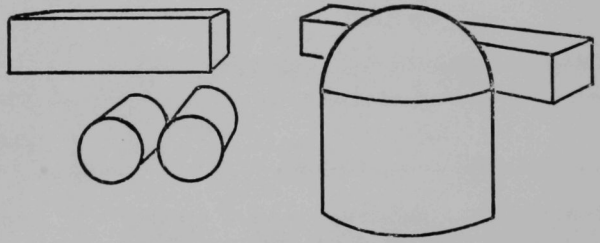

Fig. 2

Group of buildings.

Fig. 3

Human phantom.

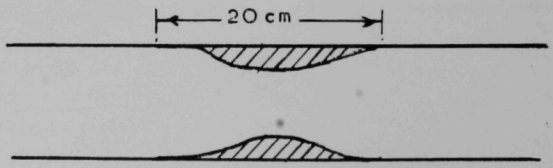

Fig. 4. Cylindrical venturi.

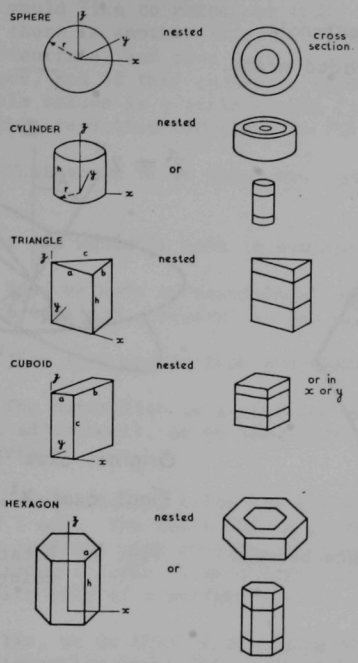

Fig. 5
Standard shapes.

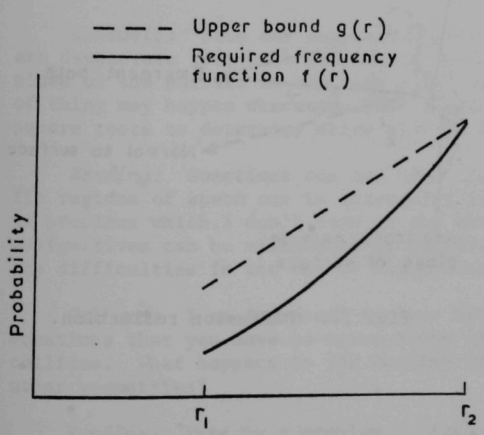

Fig. 6. Sampling in a spherical shell.

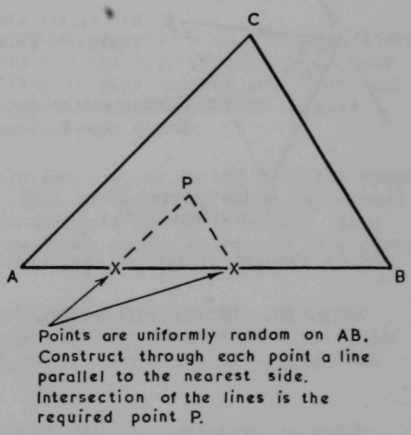

Fig. 7. Sampling for a random point in a triangle.

1. Select a point in the rectangle.
2. Reject points in the shaded areas.
3. Rotate if required.

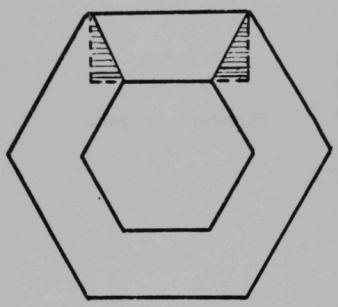

Fig. 8. Sampling from the space between two hexagons.

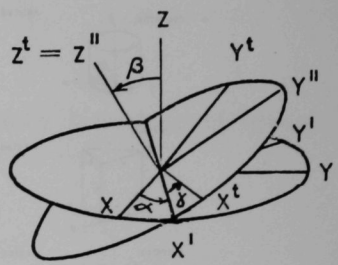

Original axes X, Y, Z.
Final axes X^t, Y^t, Z^t.

Fig. 9. Rotation of axes using Euler angles.

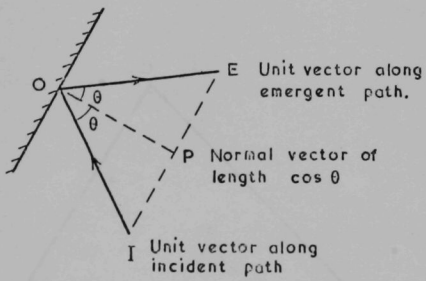

Results obtained by setting $\overrightarrow{IP} = \overrightarrow{PE}$

Fig. 10. Mirror reflection.

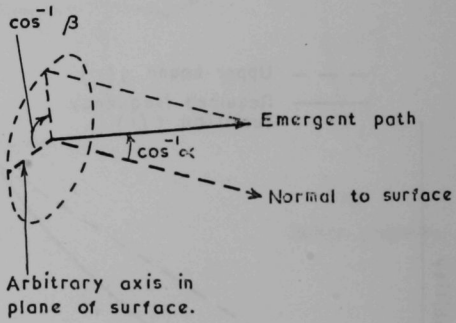

Fig. 11. Diffusion reflection.

DISCUSSION

Kalos: I would like to point out that, if you want to sample from the Watt spectrum, there is another rejection technique which can be used. If you sample two exponentials, and then calculate a quadratic expression in one of the sample values, and if this quadratic exceeds the second sample value, then one of the sample values is a valid sample from the Watt spectrum. The efficiency of this rejection technique is 70%.

Bending: That seems to be about the same as the efficiency of the technique I used.

Kalos: But you wouldn't have to compute a hyperbolic sine....

Bending: Yes, we have an exponential and a hyperbolic sine to compute against your two logs and a square, so you will probably win. Thank you.

Cashwell: You also sample from the Maxwellian. How do you do that?

Bending: The Maxwellian is a χ-square with three degrees of freedom (with a trivial adjustment), so we sample by taking the sum of squares from three normal deviants.

Cashwell: Incidentally, Kalos has a very nice way to do that also. One other point, if I may. You don't have a great deal of difficulty when you go to higher-order surfaces with geometry routines? Considering the difficulty we have had treating a torus, I am a little surprised. Do you find it necessary to know which side of a surface you are on?

Bending: Yes, we do this by arranging the surfaces into groups. We divide our configuration space into subspaces with perhaps half a dozen surfaces in a particular subspace. Then we define regions within each subspace by looking at the sign of each multinomial expression at any given point.

Cashwell: From our experience with toroidal geometries we find that there are degenerate tori with properties such that you can find yourselves on both sides of the surface at the same time, and that is what bothers me. That sort of thing may happen where you have fourth-degree surfaces and have to take square roots to determine which side of a surface you are on.

Bending: Sometimes one can have trouble identifying to the computer specific regions of space one is interested in. This is something which has caused us problems which I don't know of any simple and general way to avoid. Some ambiguitives can be eliminated by introducing extra planar surfaces. But there are difficulties in the use of higher-order surfaces.

Gelbard: I am impressed by just the problem of solving the polynomial equations that you have to solve every time you want to find where a particle collides. What happens to the running time, as compared to running time for other geometries?

Bending: This is a problem. Clearly there are circumstances in which this problem becomes so bad that the method we have described can cease to be an appropriate method. Now, we have not gotten enough experience in using

higher-order surfaces to know, really, where the method ceases to work. We don't see any solution in geometries which intrinsically require this kind of representation, except perhaps to go to a neutron-tracking method which does not require one to identify boundary crossings. Now, this can be done, and is done in many codes. It is not something which we do because such an approach does not fit into the philosophy of our code. For example, I was looking at a very complicated human phantom some time ago, in which almost all the bodily organs were represented (and very well represented) in mathematical form. This sort of geometry would have been extremely difficult for us to deal with, but it was quite possible to treat using a form of tracking in which one only needed to be able to identify the region in which a point lies, without locating boundary crossings. I think in cases of very hard geometry it may be better to use this kind of a scheme. We don't, and there will be cases for which our scheme is a bit unwieldy.

Gelbard: If you are dealing with a thermal system and are in a resonance band, and if you try to make the cross sections spatially constant, you could have a real mess on your hands; so this approach does not seem to be to be a completely general solution.

Coveyou: Well actually, in very complicated geometries, I would think that the introduction of artificial cross sections would work very well, since you wouldn't have to know where a particle crosses a boundary but only which region it is in.

Gelbard: Yes, but my point was that with ficticious cross sections you have a problem with resonances. First of all, everytime you make a collision, you have got to find the maximum cross section. Secondly, if you happen to be in a resonance the cross sections may be very large, and unless you go to some other scheme in which you use fictious cross sections in some parts of the geometry and not in others, you have to track through the whole geometry with a huge cross section.

Coveyou: I think Bending was talking about problems where there is no multiplication, where there isn't likely to be any uranium.

Kalos: But there are some important resonances in my calcium.

Bending: When we say that our code is a rather general Monte Carlo code, we are saying that it has general capabilities. We are not saying necessarily that our method is the only, or the best, or even a viable method for every conceivable Monte Carlo problem. I think one could certainly construct situations in which the combination of geometry problems, and perhaps thermal tracking problems and resonances are such that perhaps none of us can conceive of a viable route for solving them. In such situations we cannot contribute a great deal. We would like to feel that we are able to make a significant contribution to a wide range of problems, but not necessarily to every one.

Gelbard: You mentioned something about thermal neutron scattering. Do you have a thermal neutron scattering treatment?

Bending: No, we have a slot into which it could fit. We have space reserved for bivariate distributions. Now, the principal object in that space will be some sort of treatment of $S(\alpha,\beta)$ thermal scattering. We have not as

yet done any coding to fit into that space, so we can't contribute to that one at the moment.

Gelbard: You mentioned S(α,β) sampling and a number of other people have. Is it generally felt that this is very much superior to the use of a transfer matrix? It seems to me that the coding for a transfer matrix is much simpler, and I don't see that much advantage to the S(α,β) scheme, with its added complications.

Bending: I can't comment on the technical virtues of S(α,β) sampling relative to a transfer matrix except to say that the philosophy of SPARTAN is that the code writer does not determine the physical laws which are contained in it. If a user wants us to write the code so that it can deal with S(α,β), then we will make it deal with S(α,β). If that makes it cost twice as much, then that is his problem. We obviously would look at alternatives, but we don't feel that it is right for the code's designer or programmer to specifically decide against a particular method, and not make his code capable of using it.

Gelbard: It struck me that your geometry treatment is very similar to MAGI's combinatorial geometry package. I was wondering if anyone has any comment as to the relative merits of the two schemes? Is there any obvious advantage of one over the other?

Kalos: I got the same impression, but I don't think that I could comment on the basis of Bending's presentation. One would have to ask a lot of technical questions.

Steinberg: The combinatorial geometry systems that we use in SAM-CE has been around for quite a while, and has undergone an extensive amount of development and shakeout. We have had a lot experience with it. In fact, as you mentioned, you are putting it in VIM — people have put it into the Morse code. So that it has been used a lot and it has been proved out, and I suppose in that sense it has some advantages. I think that structurally the two schemes look very similar. I see no fundamental difference. The combinatorial package is there and has been shaken down; that is about all I can say. In fact, you can even make motion pictures with it.

Kalos: We have been much more conservative, however, about adding higher-degree surfaces.

Bending: We have added these surfaces as a facility which is available. The code will find the proper roots in the proper places. The fact that we have added this facility does not imply that we necessarily recommend its use all that freely. As we have pointed out, there are cases where it may be needed, but at the same time we need to be aware of the time penalties involved in using that part of the system.

Borgwaldt: I am not so familiar with such types of geometry routines. can you say anymore about it? Is there any documentation on this combinatorial geometry?

Bending: Yes, there is.

Borgwaldt: Fine, I'll see you later.

Kalos: Briefly, there is a simple input, and there is a processing code which translates it into an organization of equations which must be solved. In current versions some of the topological data that you need to know about, the information that region 75 is on the other side of region 31, does not have to be specified as input. It is learned heuristically in the course of running a Monte Carlo. The procedure is self-learning.

Borgwaldt: I have another remark. If I use normal first- and second-order surfaces, and the equations for them, then I have a means of finding out whether some input is complete nonsense because I have a failure mechanism. If one really wants to introduce higher-order equations, can one find failure mechanisms which identify illegal surface equations?

Steinberg: In the combinatorial geometry system, as we use it, there is nothing higher than second degree, so all our checking mechanisms have been based on second-degree equations. We have not implemented the torus. I should mentioned that at MAGI we have one basic version and many specialized versions. Now, in the basic version, the torus does not exist. On the other hand, the torus has been implemented in some of the special versions. The CHECKER code, for example, which was discussed at the last ANS meeting, in the basic version, is only dealing with the quadratic surfaces — quadratic and planar.

Bending: In SPARTAN, to answer your question on behalf of the SPARTAN geometry system, there are three levels of checking. One is a check that the data are internally self-consistent, and there are quite a few mistakes that one can make that will be picked up by the data checking routine and thrown back at you. Beyond that there are facilities for drawing pictures of the layout. These are relatively the same as the ones which were described in connection with MONK earlier. Ultimately, if you get into the situation which was mentioned earlier, of accidentally creating regions in space and then forgetting to say what material they contain or anything like that, you would get some type of failure during run time, and I think a failure is preferable to a wrong answer. The intention is that, to check for errors, one should use the facility for drawing out maps of the surface, of the layout as a whole.

Coveyou: What happens if you have one of these illegal layouts and a particle never gets into it?

Bending: It does not matter, in that case. You get the right answer. Nothing goes wrong: if you accidentally create a region for which you specify no parameters, nothing goes wrong until a particle happens to get into it. When it does, you'll get an error message from one of the geometry routines.

Gelbard: In the discussion of the error checking routine at the Philadelphia ANS Meeting, the point was made that in the original combinatorial geometry package one could define a region multiply (i.e., a region could be included separately in hwo different definitions) and that without the new checking capability, the code would not stop, that it would do something when it got to that region. Now, if the region is not defined, the code has got to stop, but if it is multiply defined, at least in this case it will not. Could that happen in SPARTAN?

Bending: If a region is multiply defined, in SPARTAN, one definition overrides the other one in a systematic way. There are circumstances in which the material in a given region may be specified twice, and if the definitions disagree, this would not be flagged as an error. One will always know which material was actually used by the code.

Gelbard: In fact, it may not be an error; it may be a legitimate way to write input?

Bending: I have regarded this as a legitimate way of presenting the input.

Gelbard: Is this code intended primarily for safety studies and for the study of radiation effects?

Bending: These are two principal uses. I think probably the biggest area of usage we anticipate for the code will be in safety studies. We also hope that it will find a use in shielding calculations: not in regular shielding calculations on fairly straightforward geometries, for which we have already gotten capabilities elsewhere, but for particular shielding problems where the geometry is beyond the range of other codes. One thinks of deep penetration streaming calculations, for example, as typical of the sort of problem to which we may make a useful contribution.

Gelbard: But it does seem to be true, as we pointed out before that there is relatively little use of Monte Carlo, generally, for critical calculations. Your code is not really primarily meant forcritical calculations, and apparently there has not been much use for Monte Carlo in critical calculations. Is that fair to say of other installations? Whatever our own feelings may be, must we concede that Monte Carlo as a tool for analysis of criticals has so far not been very successful?

Gast: I would just like to make a comment on the $S(\alpha,\beta)$ treatment in the thermal group. Following the suggestions Gelbard made while at Bettis, a simple test program was written which had a complete continuum of exit energies and angles and sampled $S(\alpha,\beta)$. We compared results we got with this test code with those given by our standard treatment, in which we use a 25×25 transfer matrix with P_3 scattering. All this was done for hydrogen bound in water. We found, in a test problem which was much more severe than any we would normally encounter in a water lattice, that $S(\alpha,\beta)$ and the transfer matrix gave results, in this case, which were indistinguishable.

Bending: This is interesting to note. In fact, the standard transfer matrix approach could probably be accommodated within our existing framework without extending our physical system at all, I think.

Gelbard: I would like to make one more comment before we leave the subject. There seem to be several Monte Carlo codes which either use some form of thermal multi-energy treatment, or in which such a treatment is proposed. On the other hand we hear very little about the treatment of the epi-thermal band, and the proper treatment of epi-thermal neutrons is a fairly tricky business. The epi-thermal binding in water does have a significant effect on the flux spectrum up to 2, 3, and 4 eV, and this is sometimes important. So it seems to me that if you are going to go to the trouble of putting in a

thermal binding treatment of some sort, that it really would be necessary to make some sort of provision for binding in the epi-thermal range also.

Kalos: I wanted to add a very simple-minded remark about geometry checking: namely that in complex situations (situations where, for example, you are dealing with multinomial surfaces) the device of picking a couple of points at random, tracking from one to another and then backwards, and then checking that you got back to where you came from, can be a powerful, though, naive scheme of data checking.

ON RANDOM NUMBER GENERATION

R. R. Coveyou

Union Carbide Corporation
Oak Ridge, Tennessee 37830, U.S.A.

ON RANDOM NUMBER GENERATION

Robert R. Coveyou

Union Carbide Corporation Nuclear Division
Oak Ridge, Tennessee 37830, U.S.A.

ABSTRACT

The subject of random number generation is currently controversial. Differing opinions on this subject seem to stem from implicit or explicit differences in philosophy; in particular, from differing ideas concerning the role of probability in the real world of physical processes, electronic computers, and Monte Carlo calculations.

An attempt is made here to reconcile these views. We propose not to discuss the role of stochastic ideas in the real world, but rather to discuss their role in our mathematical models. In illustration of these ideas, we construct a mathematical model of the use of random number generators in Monte Carlo calculations, and use the analysis of this model to set up criteria for the comparison and evaluation of random number generators.

INTRODUCTION

It is not easy to understand why the design, used in Monte Carlo calculations, and analysis of random number generators (RNG) is so hard to understand. Or, put another way, why it is so easy to understand in several different, and mutually inconsistent, ways.

Nevertheless, the set of all expressed understandings of the subject slightly outnumbers the set of all workers who have devoted a significant amount of attention to it (people sometimes change their minds).

It is my belief that some of the current variant opinions concerning the design and use, in Monte Carlo calculations, of RNG stem from equally variant opinions, explicit and implicit, concerning the wider question of the logic of the application of stochastic (indeed, any mathematical) methods to the analysis of physical (and other scientific) problems.

This is neither the time nor the place for an extended discussion of these wider questions. We will attempt to deal with them here only insofar as is necessary for the partial understanding of the subject at hand.

The philosophical questions referred to here are those revolving around the relations between physical reality and our mathematical comprehension of this reality. Just what these relations really are, I do not pretend to know. And, so far as I can tell, the physicists don't really pretend to know either.

Here, we can and must restrict ourselves to narrower issues. These issues are those which revolve around the relations between the physical realities which we seek to model in Monte Carlo calculations, and our mathematical models of these realities.

Suppose that we wish to analyze some physical (or other) problem. Let us say that the problem is that of the effective prediction of the future behavior of some physical system. For whatever reason, we desire to use a stochastic model of the system in our analysis. It may well be that our choice is motivated by experience; that we know that stochastic models of systems similar to the system of interest have been effective in the past. Or it may be motivated by reasons of personal preference for stochastic models. The point is that the reasons for the choice of a stochastic model are not so much logical as aesthetic.

So, we construct a mathematical model, of stochastic character, for the physical system which gives rise to our problem, and for its operation. This model takes the form of a <u>probability space</u>, each member of which is interpreted as (a possible record of) a possible outcome of the operation of the system. There is a <u>distinguished</u> single member of the probability space, whose identity we may or may not know, which is interpreted as (the record of) the actual performance of the system. We may also interpret observations of the actual performance of the system as an effort to identify this distinguished element, as precisely as is feasible.

Having constructed our model, we decide, for whatever reason, to use the Monte Carlo method in our analysis (perhaps we wish to present our results at this meeting).

Now, an actual Monte Carlo calculation is, itself, a <u>physical</u> process, involving the physical (and, sometimes, mental and emotional) behavior of computers; human, electronic, and other. Hence, discussion of the (stochastic or deterministic) character of the calculation is subject to much the same difficulties as before. These difficulties can be bypassed, if not resolved, by the same (model theoretic) technique.

So, we here discuss, not actual Monte Carlo calculations, but mathematical models of such calculations. And, in particular, a specific model for such calculations, chosen so that the properties of the RNG, used in Monte Carlo calculations, are emphasized in the analysis of the model.

THE MODEL

We will base our analysis of RNG on consideration of their performance in the evaluation of definite integrals of integrable functions defined on the real unit interval and on higher-dimensional analogues of the unit interval. All of the results here can be and will be extended to evaluation of such integrals defined on compact separable abelian topological groups, which are the most general manageable by the methods here used.

Let

$$U = [u \in R: \ 0 \leq u < 1]$$

denote the real unit interval, considered as a compact, separable, abelian topological group, with group operation ordinary addition (mod 1), furnished with Haar measure equivalent to ordinary Lebesque measure.

For the positive integral n, let

$$U^n = [\vec{U}_n = \{U_1,\ldots,U_n\}: \ U_j \in U; \ 1 \leq j \leq n]$$

denote the n-dimensional real unit cube, furnished with appropriate algebraic, topological, and measure theoretic paraphenalia. These too are as described above.

Let Ω denote the space of all real (or complex) valued integralbe functions

$$\omega: \ U^n \to R: \ \omega(\vec{u})$$

defined on U^n.

We will make Ω into a probability space, in accordance with rules set down later. Then each $\omega \in \Omega$ will be a <u>random function</u>.

Definition: a random number generator (RNG) is (a procedure designed to produce) a sequence

$$\Gamma\{\gamma_0,\gamma_1,\ldots,\gamma_{p-1},\gamma_p,\gamma_{p+1}\ldots\},$$

with each $\gamma_j \in U$.

For specific ω, Γ, and positive integers N, let

$$I(\omega) = \int_{U^n} \omega(\vec{u})\, d\vec{u}$$

denote the integral of ω over U^n, and let

$$\hat{I}(\omega,\Gamma,N) = \frac{1}{N} \sum_{j=0}^{N-1} \omega(\vec{\gamma}_j)$$

$$= \frac{1}{N} \sum_{j=0}^{N-1} \omega(\gamma_j,\gamma_{j+1},\cdots,\gamma_{j+n-1})$$

denote the familiar Monte Carlo estimate of the integral of ω of sample size N firnished by the RNG Γ.

Our interpretation of this model for the use of RNG in Monte Carlo calculations is that, being armed with a specific RNG Γ, we contract, given any function $\omega \in \Omega$, to provide an estimate $\hat{I}$ of the integral I of ω.

We also assume that the functions ω, presented to us for estimate of their integrals, are chosen at random from the probability distribution we have imposed on Ω.

Of course, the specification of the model is not yet complete. We have not yet specified the required probability distribution on the space Ω of <u>problems</u>.

For a really detailed analysis, it would be necessary to specify this distribution exactly. But it would not be wise to follow this course, since it would inevitably lead to the suspicion that the resulting judgments concerning the evaluation of specific RNG, or specific classes of RNG, would be quite sensitive to the choice of problem distribution, and would, possibly be quite different for different distributions.

So, we will <u>not</u> specify the problem distribution in complete detail here. In fact, we will specify the problem distribution to that extent here, and to that extent only, that allows us to make some little progress with the analysis.

So, we ask just what we need to know about the problem distribution before we can say anything meaningful about the comparison of RNG in the light of our chosen criterium of mean square error.

Since we have allowed <u>any</u> sequence Γ to call itself a RNG, we must also set forth criteria which will serve to separate "good" RNG from "bad." It seems to me that a very natural figure of merit for specific RNG is furnished by the mean square error $Q_N(\Gamma,\Omega)$

$$= \left\langle \omega \in \Omega: \ |\hat{I}(\omega,\Gamma,N) - I(\omega)|^2 \right\rangle .$$

An easy calculation yields that

$$Q_N(\Gamma,\Omega)$$

$$= \frac{1}{N^2} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \left\langle \omega \in \Omega: \ \omega(\vec{\gamma}_j)\omega(\vec{u}) \right\rangle d\vec{u}$$

$$+ \int_{U^n} \int_{U^n} \left\langle \omega \in \Omega: \ \omega(\vec{u})\ \omega(\vec{v}) \right\rangle d\vec{u}\ d\vec{v}$$

$$= \frac{1}{N^2} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} A(\vec{\gamma}_j,\vec{\gamma}_k)$$

$$- \frac{2}{N} \sum_{j=0}^{N-1} \int_{U^n} A(\vec{\gamma}_j\vec{u})\ d\vec{u}$$

$$+ \int_{U^n} \int_{U^n} A(\vec{u},\vec{v})\ d\vec{u}\ d\vec{v},$$

if we assume the existence of the <u>correlation function</u>

$$A(u,v) = \left\langle \omega \in \Omega: \ \omega(\vec{u})\ \omega(\vec{v}) \right\rangle .$$

It is quite natural to assume, as we do, that

$$A(\vec{u},\vec{v}) = A(\vec{u} - \vec{v});$$

that is, that A is translation invariant. Then

$$Q_N(\Gamma,\Omega) = \frac{1}{N^2} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} A(\vec{\gamma}_j - \vec{\gamma}_k) - \int_{U^n} A(\vec{u})\ d\vec{u}.$$

Therefore, in order to make <u>some</u> progress in our analysis, we explicitly make the following assumptions concerning the problem distribution, and for the moment, these only:

For each $\vec{u} \in U^n$, (0)

$$\langle \omega \in \Omega: \ \omega(\vec{u}) \rangle = 0.$$

These exists a function (1)

$$A: \ U^n \to R: \ A(\vec{u})$$

such that, for each $\vec{u}$ and $\vec{v}$ in U^n,

$$\langle \omega \in \Omega: \ \omega(\vec{u}) \ \omega(\vec{u} + \vec{v}) \rangle = A(\vec{v}).$$

also,

$$\langle \omega \in \Omega: \ |\omega(\vec{u})|^2 \rangle = A(\vec{0}) = 1. \qquad (2)$$

It is known that, given an otherwise appropriate function A defined on U^n, the <u>formal</u> necessary and sufficient condition that A be the correlation (or covariance) function of a distribution of random functions defined on U^n is that, for each set of N points u_j; $1 \leq j \leq N$; of U^n, and each set of N real constant 2_k; $1 \leq k \leq N$, we have

$$\sum_{j=1}^{N} \sum_{k=1}^{N} \lambda_j A(u_j - u_k) \lambda_k \geq 0.$$

That is, A must be a real, positive definite, function defined on U^n.

The analysis thus far furnishes criteria for the <u>comparison</u> of RNG; it is natural to say that, of the two RNG Γ, and Γ_1 is preferable to Γ_2 iff $Q(\Omega,\Gamma_1) \leq Q(\Omega,\Gamma_2)$. In this case, of course, all we can yet say is that Γ_2 with respect to the problem distribution Ω.

But another natural question, and one for which we must, if we can, provide some sort of answer, is that of determining whether or not a specific RNG, of those RNG of u specific class, or the best RNG of u specific class may be regarded as satisfactory at all.

Now, it might seem natural to assume that the best RNG Γ are those which most clearly resemble a sequence formed by successive random independent choices from U^n with probability distribution that specified by Haar (or Lebesque) measure. We will not find this to be true. But it does seem natural to define a RNG to be <u>satisfactory</u> if it is at least as good as the average such random sequence.

Now this average is quite easily calculated to be

$$\frac{1-\int A(\vec{u})\, d\vec{u}}{N}$$

and we <u>can</u> say that the RNG Γ is satisfactory (at least for sample size N, starting with $\vec{u}_j$) iff

$$Q(\Omega,\Gamma) \leq \frac{1-\int A(\vec{u})\, d\vec{u}}{N} .$$

DISCUSSION

I have carried this analysis much further and, in particular, have carried out the indicated generalization to compact separable abelian groups, and have used the theory of harmonic analysis on such groups to get much more extensive results.

In particular, I have some interesting results about <u>periodic</u> RNG. I have found quite strong evidence that periodic RNG whose output vectors fill a lattice (finite subgroup) of the unit cube may be very good RNG and, in particular, may be quite a bit better than the "random" RNG briefly discussed above. I should also say that some are much worse! This has a bearing on the current, widespread, and, I think, unmotivated, distrust of such lattice filling periodic RNG.

DISCUSSION

Gelbard: Is it not true that the first objection to conventional random number generators came from Marsaglia, who was playing poker on the computer and found that he was losing consistently? Doesn't this suggest a significant defect in the generator?

Kalos: It was dice!

Coveyou: In Marsaglia's original paper, to which my paper on the Fourier analysis was a partial reply, he said that the conventional method was bad. Actually what he had found out was that this choice $\lambda = \sqrt{p}$ is a very bad choice. It's bad because, while it makes two random consecutive numbers practically independent, it makes three consecutive numbers almost completely dependent, in a very simple way. The fact is that each one of Marsaglia's examples (the first time he wrote about this subject) was a random number generator in which the choice of parameters was faulty and, the fact that it was faulty was known in the business at the time. What I did in my paper was to apply my test to each of the generators he had analyzed, and I found that my test predicted that they would be faulty. On the other hand, Marsaglia's test, a particular form of an χ-square test that he used, showed nothing at all to be wrong with his generator. I said at the time that he had simply chosen random number generators that were not typical of the performance of the best congruential generators. If you chose the multiplier at random, your chance of getting a random number generator as bad as the ones that Marsaglia analyzed would be practically zero.

Cashwell: There does seem to be something funny about Marsaglia's tests, because Whitesides said that with our generator he could not duplicate Marsaglia's results at all. He got good results.

Kalos: Maybe he just had a bad run.

Coveyou: Marsaglia had a bad congruential random number generator, and he simply assumed that he had a typical congruential random generator. He did not — he had a bad one. It was one which, at the time, people in the field knew was not good. In fact, it was not one — it was three or four like this that he tried. All of them had this characteristic.

Gast: In his talk Coveyou discussed a figure of merit, q, which he used to test random number generators. Why isn't it sufficient to use the serial coefficient of various lags, directly, to test the quality of a random number generator?

Coveyou: Because, as I pointed out yesterday, the serial correlation coefficient is not invariant under translations of the addend. Now, what that means is that effectively identical random number generators can have different correlation coefficients. This is a characteristic of correlation coefficients of lag 1, and all other lags for that matter. This fact was first pointed out in Berger Johnson's book. He didn't stress it, but he did point out that the serial correlation coefficient is distinctly not a good test. It is a good test in the sense that, if a random number generator has a high serial correlation coefficient, then it is not a good generator. However, a random number generator that is not good can also have a low serial correlation coefficient.

Indeed it has been suggested (and I was guilty of making the suggestion) that you can reduce the serial correlation coefficient of an additive, linear, congruential random number generator simply by changing the addend. In fact, there is a choice for the addend that actually makes the serial correlation coefficient very small and Berger Johnson pointed out that changing the addend does not change the distribution of the pairs or triplets at all. Without actually going through the mathematics it is a little difficult to see how this can happen: but it does happen, and therefore the serial correlation coefficient is just not a complete description of the correlation.

Kalos: But it must be true that the amount by which the serial correlation coefficient can be changed by this translation is bounded.

Cashwell: Yes.

Kalos: If you know that they are bounded, you know that serial correlation coefficients below some value are all equally good. So if you use this information and try to make as many serial correlation coefficients as small as possible, that sounds like a sensible test.

Coveyou: Perhaps it could be. Maybe you can actually exploit this fact I talked about by essentially adjusting the addend not to make the serial correlation coefficient as small as possible, but to make it as large as possible. The maximized correlation coefficients might be used as a test for the original random number generator. You wouldn't actually use the random number generator that you constructed to have as high a correlation as possible. You just use it as an example of this whole class of generators which really all have the same distribution of pairs and triplets.

Kalos: But what you are telling us is that minimum serial correlation coefficient is not the right criterion.

Coveyou: Right.

Kalos: But at the same time there is a bound which you ought to attain, if the generator is any good.

Coveyou: Yes, that is true.

Kalos: So I think that you dismissed Gast's suggestion a little bit too quickly.

Borgwaldt: I have heard the argument, in connection with application of random number generators, that the main point is the character of the problem that you are treating. In other words, that a well-behaved problem will get along with a bad random number generator, if one could specify what a bad random number generator is. Now, I would say that if this is true, and if one can say whether a problem is well behaved or not, that this should enter into your definition of a probability space. Badly behaved problems should have very low measure in your probability space. Could you comment on the definition of the probability space, R, which you introduced. I think you skipped some important points in your argument at that time.

Coveyou: Well, what I have to say about such an approach is this. I am thinking about doing large Monte Carlo calculations on computers, and I believe that the only reasonable choice for my purpose is a random number generator that is, generally, essentially as good as you can make it. In point of fact, I do not agree with the argument. It is true that, in certain problems, you can afford a sloppy random number generator, but, I do not think that should be relevant. The point is that one thing that we all do not want to do is to wonder, everytime we do a Monte Carlo calculation, what random number generator we can get away with. The Monte Carlo user wants a generator he can depend upon.

Kalos: I would like to make two comments: An example, an amusing example, of a situation in which you can get away with a sloppy random number generator is a linear transport problem in which importance sampling has been carried out to the ultimate so that, in fact, you use a zero variance estimation procedure. Then the answer is the same, independent of what random number you use, independent of all properties of the random generator. Thus, in some sense you can get away with sloppy random number generation. The second remark is that I assume that everybody does soluble test problems, from time to time, of a general character. One such problem, for example, is the linear straight-ahead model Boltzman equation. One solves such a problem and looks for the right answer. And one does a few integrals here and there to make sure he gets the right answer for those integrals. You don't tell your friends about this, but you want to make sure that, as installed and as you use it, your random number generator can be relied upon.

Coveyou: I think that is probably correct, the right way to do business. But I suspect that, when people do this, they are not particularly thinking about the random number generator. They ary thinking about the logic of their Monte Carlo code.

Gelbard: Yes, I think that you are agreeing that you would like a random number generator that you can rely on even when you are doing rather peculiar, over-simplified problems, in which the systematics of the generator might become important; I think that was Kalos' point.

Coveyou: I should have mentioned further, the magnitudes of the errors we are talking about. Let me be more specific. If you choose the multiplier at random, then the largest magnitude you expect in the mean square error is proportional to $1/\sqrt{\rho}$ for a randomly chosen random number generator, and that means that the errors introduced by the random number generator, unless it is very bad, are quite small. In fact they are, in almost every case, far smaller than the statistical error of the calculation.

Borgwaldt: Well, I must come back again to the problem of specification. A random number generator which creates random numbers between zero and one may have a period of p or p over four. But, if I use this random number generator in a specific problem, and call it up in different situations, I will not get this period, but a period which is considerably larger. At one time, for example, we used the following system. We used one random number generator to drive four secondary random generators for exponential distribution, isotropic angular distribution and so on. Each was initialized by a separate initial random number. But then we found that there was a big advantage in having all of these random number generators driven by one primary generator,

and this should, according to my understanding, give a period which is far higher than that of the fundamental random number generator with a period of p over four. Is that right?

Coveyou: But my one comment is that the fact that the random number generator is periodic, and the magnitude of the period, are just simply not relevant to Monte Carlo problems. Perhaps I should make an exception of very extensive Monte Carlo calculations done on the IBM-360 with a single precision random number generator. There the period is a billion. For almost every other machine it is far higher than that. It is just not relevant what that period is; it is just too big already.

Kalos: As I understand the comment, there is a conjecture that when you go through the period, and come back to the same random numbers, they are being used for a different purpose, so that the period does not matter.

Coveyou: I agree with that.

Kalos: However, that strikes me as being something that I would not want to rely on. It is perfectly possible, but I don't want to rely upon it. I suggest, for everybody who is engaged in this sort of work, the following calculation that I make. Consider the period of the random number generator you have, and suppose that your machine is generating nothing but random numbers as fast as it can generate random numbers. On the 6600 in assembly language, since there are two multipliers and there is an instruction which generates one random number, you can generate random numbers at the rate of two per microsecond. If you write a little program, an assembly language program, to generate the whole sequence of random numbers, it will repeat itself after a year. So you run your 6600 for a year doing nothing but generating random numbers before you exhaust the period. I was somewhat reassured by that. On the 7600 the calculation was a little bit harder because of the fact that the multiplier is pipelined; exactly what the cycle time is, I do not know, but it is the order of months. Now, you should do that calculation with your IBM-360 using, I would hope, a double precision generator.

Coveyou: A double precision generator on the IBM-360 would last forever.

SESSION III

THE BELL TRANSFORM

C. J. Everett and E. D. Cashwell

Los Alamos Scientific Laboratory
Los Alamos, New Mexico 87544, U.S.A.

* * *

Presented by E. D. Cashwell

THE BELL TRANSFORM

C. J. Everett and E. D. Cashwell

Los Alamos Scientific Laboratory
Los Alamos, New Mexico 87544, U.S.A.

ABSTRACT

An exposition and extension is given of some work of M. Bell at CERN, involving the construction of a radially symmetric probability density in the plane which has a given marginal density. The problem is of importance in certain Monte Carlo procedures. While our results are more extensive and precise, the crucial step relies on a remarkable integral transform discovered by Bell.

INTRODUCTION

Let Y be the set of all strictly decreasing functions $y(r)$ on the interval $[0,R]$, with $y(0) = 1$, $y(R) = 0$, which are of class C^1 on $(0,R)$, and F the set of continuous functions $f(r) \geqslant 0$ on $(0,R)$, not identically zero on any subinterval, with

$$\int_0^R 2\pi r\, f(r)\, dr = 1 \quad . \tag{1}$$

The functions $f(r)$ of F define a class of radially symmetric densities $\rho(x,y) \equiv f(\sqrt{x^2 + y^2})$ on the circle $C = \{(x,y);\ x^2 + y^2 \leqslant R^2\}$. In particular it follows from Eq. (1) that

$$\iint_C f(\sqrt{x^2 + y^2})\, dxdy = \int_0^{2\pi}\int_0^R f(r) r dr d\theta = 1 \quad . \tag{2}$$

Let P be the set of all marginal densities

$$p(x) = 2\int_0^{\sqrt{R^2 - x^2}} f(\sqrt{x^2 + y^2}) dy \tag{3}$$

on (-R,R) for the functions f(r) of F. Finally, define Y^* as the set of all functions $y^*(r)$ on [0,R] arising from the marginal densities p(x) of P by way of the integral transform

$$y^*(r) = 2\int_0^{\sqrt{R^2 - r^2}} p(\sqrt{r^2 + x^2})dx \quad . \tag{4}$$

This transform, which has the <u>appearance</u> of a marginal density, was introduced and exploited by M. Bell [1], and is the key to all that follows.

Now consider the mappings

$$y(r) \to f(r) \to p(x) \to y^*(r) \tag{5}$$

for these sets Y, F, P, Y^*, defined respectively by

$$f(r) = -\frac{1}{\pi}\, dy/dr^2 \quad ; \quad 0 < r < R \tag{6}$$

and Eqs. (3) and (4). It is shown below that $y^*(r) = y(r)$ for every y(r) in Y, and that all three of the correspondences in (5) are one-to-one, and onto all of the sets F, P, Y^* respectively.

It follows that a given marginal density p(x) uniquely determines its radially symmetric density f(r). It is also proved that a given continuous density p(x) on (-R,R), which is an even function, not zero on any subinterval, is indeed the marginal density of some f(r) of F if and only if its Bell transform $y^*(r)$ defined by Eq. (4) is strictly decreasing and of class C^1 on (0,R). The unique f(r) for such a density p(x) is then explicitly constructible via the mappings of (5), in the order

$$p(x) \to y^*(r) \equiv y(r) \to f(r) \tag{7}$$

THE MAPPING $y(r) \to f(r)$

We consider as basic here the more familiar correspondence $f(r) \to y(r)$ defined on F to Y by the relation

$$y(r) = 1 - \int_0^r 2\pi f(r) r dr \quad ; \quad 0 \leqslant r \leqslant R \tag{8}$$

where the integral is simply the distribution function of the density $2\pi f(r)r$ on $(0,R)$, and y, or $1 - y$, plays the role of the random number in straightforward Monte Carlo sampling of this radial density for r on $(0,R)$.

Theorem 1

The correspondence $f(r) \to y(r)$ defined by Eq. (8) is one-one on F to all of Y, its inverse $y(r) \to f(r)$ being given by the relation

$$f(r) = -\frac{1}{\pi}\, dy/dr^2 \quad ; \quad 0 < r < R \quad . \tag{9}$$

Proof: (a) The function $y(r)$ in Eq. (8) is well defined on $[0,R]$, with $y(0) = 1$, $y(R) = 0$, and of class C^1 on $(0,R)$, since $f(r)$ is continuous, with integral (1); moreover $y(r)$ is strictly decreasing since $f(r) \geqslant 0$ and not zero on any subinterval. Hence $y(r)$ is in the set Y. (b) The correspondence is one-to-one, since $f \to y$ and $g \to y$ implies by Eq. (8) that

$$\int_0^r f(r)rdr \equiv \int_0^r g(r)rdr \quad . \tag{10}$$

Differentiation in Eq. (10) shows that $f(r)r \equiv g(r)r$ and hence $f(r) \equiv g(r)$ on $(0,R)$. (c) The mapping (8) carries F onto all of Y. In fact, if $y(r)$ is a given function in Y, then the function (6), namely

$$f(r) \equiv -\frac{1}{\pi}\, dy/dr^2$$

is well defined and continuous on $(0,R)$ since $y(r)$ is of class C^1; $f(r) \geqslant 0$ and not zero on any interval since $y(r)$ is strictly decreasing; finally we verify the integral property (1),

$$\int_0^R 2\pi f(r)rdr = -\int_0^R (dy/dr^2)(2rdr) = -\int_0^R (dy/dr)dr = y(0) - y(R)$$

$$= 1 - 0 = 1 \quad .$$

Hence the $f(r)$ so defined is in F. Finally, we see that this $f(r)$ maps into the given $y(r)$ under (8), and therefore that Eq. (6) defines the mapping $y(r) \to f(r)$ which is inverse to $f(r) \to y(r)$. In fact, we have for r on $[0,R]$,

$$1 - \int_0^r 2\pi f(r)rdr = 1 + \int_0^r (dy/dr^2)2rdr = 1 + \int_0^r (dy/dr)dr$$

$$= 1 + y(r) - y(0) = y(r).$$

THE ITERATED MAPPING $y \to f \to p \to y^*$

It is obvious by definition that the correspondences $f \to p$ and $p \to y^*$ given by Eqs. (3) and (4) are respectively on F to all of P, and on P to all of Y^*. The remarkable transform (4) now serves to show that, in the iterated mapping (5), namely

$$y(r) \to f(r) \to p(x) \to y^*(r)$$

we have actually come full circle.

Theorem 2

The correspondence (5) takes Y into Y identically, i.e., $y^*(r) = y(r)$ for every $y(r)$ of Y, and hence $Y^* = Y$. In fact, under the mapping

$$f(r) \to p(x) \to y^*(r) \tag{11}$$

it appears that

$$y^*(r) = 1 - \int_0^r 2\pi f(r)rdr = y(r) \quad ; \quad 0 \leqslant r \leqslant R \quad . \tag{12}$$

Moreover, all three of the correspondences in (5) are one-to-one and onto all of the sets F, P, Y^*, respectively.

Proof: Evaluation of $y^*(r)$ in terms of $f(r)$ via the mapping (11) shows that

$$y^*(r) = 2\int_0^{\sqrt{R^2 - r^2}} p(\sqrt{r^2 + x^2})dx$$

$$= 2\int_0^{\sqrt{R^2 - r^2}} dx\ 2\int_0^{\sqrt{R^2 - (r^2 + x^2)}} f(\sqrt{(r^2 + x^2) + y^2})dy$$

$$= 4\int_0^b dx \int_0^{\sqrt{b^2 - x^2}} f(\sqrt{r^2 + (x^2 + y^2)})dy$$

where we have set the constant $R^2 - r^2 = b^2$ temporarily. We thus obtain an integral over the circle $x^2 + y^2 \leqslant b^2$ in the first quadrant, and introducing polar coordinates via the transformation

$$x = \rho \cos \theta$$

$$y = \rho \sin \theta$$

shows that

$$y^*(r) = 4\int_0^{\pi/2} \int_0^b f(\sqrt{r^2 + \rho^2})\rho d\rho d\theta$$

$$= \pi \int_0^{\sqrt{R^2 - r^2}} f(\sqrt{r^2 + \rho^2})2\rho d\rho \quad .$$

Changing from ρ to the variable $s = \sqrt{r^2 + \rho^2}$ then yields

$$y^*(r) = \int_r^R 2\pi f(s)sds = 1 - \int_0^r 2\pi f(s)sds = y(r) \quad ,$$

as characterized in terms of f by Eq. (8). We have still to show that the two correspondences in (11) are one-to-one. This will follow at once if we prove that the net mapping $f(r) \to y^*(r)$ induced by (11) is one-to-one. But this is obvious from Eq. (12), since $f(r) \to y(r)$ is one-to-one by Theorem 1.

Corollary 1

A given density $p(x)$ on $(-R,R)$ which is in fact a marginal density of some $f(r)$ in F, uniquely determines its source density $f(r)$ via the correspondences

$$p(x) \to y^*(r) \equiv y(r) \to f(r)$$

in the explicit form

$$f(r) = -\frac{1}{\pi}\frac{d}{dr^2}\left[2\int_0^{\sqrt{R^2 - r^2}} p(\sqrt{r^2 + x^2})dx\right] \tag{13}$$

which thus provides the inverse of the mapping $f(r) \to p(x)$.

Proof: The result follows from Theorem 2, using the relations (4) and (6).

The formula (13) is due to Bell [1].

DENSITIES WHICH ARE INDEED MARGINAL DENSITIES

A density $p(x)$ on $(-R,R)$ which is in fact a marginal density of some $f(r)$ in F certainly belongs to the set E of continuous even functions $e(x)$ on $(-R,R)$ with $e(x) \geqslant 0$, not zero on any subinterval, and of course having

$$2\int_0^R e(x)dx = 1 \quad .$$

We have seen in Corollary 1 how the ancestral density $f(r)$ may be constructed for a density $p(x)$ known to be marginal. There remains the problem of determining which functions $e(x)$ of the set E are actually marginal densities. An answer is provided by

Theorem 3

A density $e(x)$ of the set E is a marginal density in P if and only if its transform $\bar{y}(r)$, defined on $[0,R]$ by

$$\bar{y}(r) = 2\int_0^{\sqrt{R^2 - r^2}} e(\sqrt{r^2 + x^2})dx \tag{14}$$

is strictly decreasing and of class C^1 on $(0,R)$.

Proof: If $e(x)$ is in P, its transform $\bar{y}(r)$ in Eq. (14) coincides with its transform $y^*(r)$ in Eq. (4) which we know is in the set $Y^* = Y$ by Theorem 2. Hence $\bar{y}(r)$ is decreasing and of class C^1 on $(0,R)$ by definition of Y. Conversely, suppose $\bar{y}(r)$ in Eq. (14) has the stipulated properties, for a given function $e(x)$ in the set E. Then clearly $\bar{y}(r)$ is itself in the set Y, since $\bar{y}(0) = 1$ and $\bar{y}(R) = 0$ by (14). We may therefore define for this $\bar{y}$ of Y the functions $f(r)$, $p(x)$, $y^*(r)$ by the standard mappings of (5), i.e., with

$$\bar{y}(r) \to f(r) \to p(x) \to y^*(r) \quad .$$

By Theorem 2, this implies $y^*(r) = \bar{y}(r)$ and hence we should have

$$\int_0^{\sqrt{R^2 - r^2}} p(\sqrt{r^2 + x^2})dx \equiv \int_0^{\sqrt{R^2 - r^2}} e(\sqrt{r^2 + x^2})dx \tag{15}$$

where $p(x)$ is in P. Now the even densities $p(x)$ and $e(x)$ on $(-R,R)$ may be regarded as defining radial functions $p(r)$ and $e(r)$ on $(0,R)$, and hence when suitably renormed to satisfy the requirement (1), as functions in the set F. But then Eq. (15) asserts that the radial densities so defined have a common marginal density, and we know from the one-to-one character of the mapping $f(r) \to p(x)$ proved in Theorem 2, that this implies $p(r) \equiv e(r)$ on $(0,R)$, and hence $p(x) \equiv e(x)$ on $(-R,R)$, since both are even functions. Hence $e(x)$ is indeed in the set P of marginal densities.

For example, Theorem 3 shows that the functions $e(x) = 1/\pi(1 - x^2)^{1/2}$ and $e(x) = \frac{3}{2}x^2$ of the set E ($R = 1$) are not marginal densities. (Note that both are increasing on $(0,1)$). The former has the transform $\bar{y}(r) \equiv 1$, while the latter has a non-monotone $\bar{y}(r) = 3r^2(1 - r^2)^{1/2} + (1 - r^2)^{3/2}$, with a maximum at $r = 1/\sqrt{2}$.

The test of Theorem 3 may be difficult to apply in practice, and simpler criteria are desirable. A trivial sufficient condition may be mentioned.

Corollary 2

A density $e(x)$ of E, of class C^1 on $(0,R)$, is a marginal density in P provided $e(x)$ is non-increasing on $(0,R)$.

Proof: For $0 \leqslant r < s \leqslant R$, one has

$$\frac{1}{2}\bar{y}(r) = \int_0^{\sqrt{R^2 - r^2}} e(\sqrt{r^2 + x^2})dx > \int_0^{\sqrt{R^2 - s^2}} e(\sqrt{r^2 + x^2})dx$$

$$\geqslant \int_0^{\sqrt{R^2 - s^2}} e(\sqrt{s^2 + x^2})dx = \frac{1}{2}\bar{y}(s)$$

and the result follows from Theorem 3.

Examples 1 and 2 below exhibit marginal densities which are respectively strictly decreasing, and constant, in accord with Corollary 2, while Examples 3 and 4, displaying such densities which are non-monotone and even strictly increasing, show that the sufficient condition of the Corollary is far from necessary.

SOME EXAMPLES

We include here some examples (for the case $R = 1$) which illustrate various features of the mappings (5).

Example 1

$y(r) = 1 - r^2 \to f(r) = 1/\pi \to p(x) = \frac{2}{\pi}(1 - x^2)^{1/2} \to y^*(r) = 1 - r^2 = y(r)$; $p(x)$ decreasing on $(0,1)$.

Example 2

$y(r) = (1 - r^2)^{1/2} \to f(r) = 1/2\pi(1 - r^2)^{1/2} \to p(x) = 1/2 \to y^*(r) = y(r)$; $p(x)$ constant.

Example 3

$y(r) = 1 - r^4 \to f(r) = \frac{2}{\pi}r^2 \to p(x) = \frac{4}{\pi}\{x^2(1 - x^2)^{1/2} + \frac{1}{3}(1 - x^2)^{3/2}\} \to y^*(r) = y(r)$; $p(x)$ non-monotone on $(0,1)$, maximum at $x = 1/\sqrt{2}$.

Example 4

$$y(r) = \frac{1}{2}(1 - r^2)^{1/2}(2 + r^2) \to f(r) = 3r^2/4\pi(1 - r^2)^{1/2} \to p(x)$$
$$= \frac{3}{8}(1 + x^2) \to y^*(r) = y(r),\ p(x) \text{ increasing on } (0,1).$$

Note that all four functions $p(x)$ necessarily have decreasing transforms $y^*(r)$.

The construction of $f(r)$ from $p(x)$ by the formula (13) is easily verified in the above cases.

THE INFINITE CASE $R = \infty$

The previous results for the case of radially symmetric densities on the circle of finite radius R extend in an obvious way to such densities defined on the infinite plane, and indeed with simpler underlying formulas. Rather than repeat the steps in a formal way, we simply state the analogous methods without proof.

We now take for Y and F the sets defined above with R replaced by ∞; P is then the set of marginal densities

$$p(x) = 2\int_0^\infty f(\sqrt{x^2 + y^2})dy$$

on $(-\infty,\infty)$ for the radial function $f(r)$ of F, and Y^* the set of Bell transforms

$$y^*(r) = 2\int_0^\infty p(\sqrt{r^2 + x^2})dx$$

on $[0,\infty)$ arising from the densities $p(x)$ in P. The mappings (5),

$$y(r) \to f(r) \to p(x) \to y^*(r)$$

are defined as before with (5) unchanged.

Theorem 1 on the correspondence $y(r) \to f(r)$ then applies verbatim, as does Theorem 2 on the iterated mapping (5).

The analogue of Corollary 1 gives the inverse of the mapping $f(r) \to p(x)$ in the form

$$f(r) = -\frac{1}{\pi}\frac{d}{dr^2}\left[2\int_0^\infty p(\sqrt{r^2+x^2})dx\right] \tag{16}$$

and provides the construction of the (unique) radial density with a given marginal density $p(x)$ on $(-\infty,\infty)$.

A density $e(x)$ of the set E, defined as before with R replaced by ∞, is found to be a marginal density iff its Bell transform

$$\bar{y}(r) = 2\int_0^\infty e(\sqrt{r^2+x^2})dx \quad ; \quad 0 \leqslant r < \infty$$

is strictly decreasing and of class C^1 on $(0,\infty)$, just as in Theorem 3, and the non-increasing property on $(0,\infty)$ of a class C^1 function $e(x)$ suffices to identify it as a marginal density.

We conclude with two illustrations of Eq. (16).

Example 5

For the normal density

$$p(x) = \frac{1}{\sqrt{2\pi}}\, e^{-x^2/2}$$

on $(-\infty,\infty)$, (16) yields the ancestral radial density $f(r) = \frac{1}{2\pi}\, e^{-r^2/2}$ on $(0,\infty)$ (it had better!), with the accompanying density

$$\rho(x,y) = f(\sqrt{x^2+y^2}) = \frac{1}{2\pi}\, e^{-\frac{x^2+y^2}{2}}$$

on the entire plane.

Example 6

From (16), one finds that the Cauchy density $p(x) = 1/\pi(1 + x^2)$ on $(-\infty,\infty)$ is the marginal density of the less well known radial function $f(r) = 1/2\pi(1 + r^2)^{3/2}$.

Note that the functions $p(x)$ in both examples have decreasing Bell transforms on $(0,\infty)$, as required by Theorem 3.

REFERENCE

1. M. Bell, "Creating Bunches of Random Points in a Phase Plane with a Given Projected Distribution," CERN/MPS/DL 73-3, April 1973.

DISCUSSION

Kalos: It looks to me as though the problem that was originally put was that of inferring the radial distribution of some quantity from measured line integrals. Is that correct? If you have, for example, a plasma which is known to be radially symmetric and you measure the density, with a laser for example, along a line, then p of x is precisely such a line integral. You now are asked to infer the radial distribution of the density of the plasma, which is the y or r. So this is the sort of problem you are talking about. Now this problem is one that I have been acquainted with for a long time. Sometime around 1960 a mathematician by the name of Mort Slater solved a more general problem, which I don't think he was the first one to solve. The problem can be formulated as follows: x and y have some unknown distribution, p of x and y. Given the ability to measure line integrals as a function of distance of closest approach and angle, how do you infer p of x and y? He showed that it was possible to do this, and that the solution involves an Able integral equation. So there is still more that can be done with the problem you posed. I don't think that his proof was as elegant and as clearcut as yours, but there is more here to be mined.

Cashwell: Well, I am glad that you told me because obviously, with a result like ours, we certainly felt there had to be a lot of work that had been done on this problem, work that we were not aware of. Our analysis started, simply, as an attempt to understand what went on in the Bell paper. This was done by Slater?

Kalos: Slater, Mort Slater. I don't know if he published this in any mathematical journal. I know of Slater's work only through a memorandum at United Nuclear Corporation. It was made the basis of a device for measuring densities by gamma-ray transmission, and was actually applied to the measurement of the density distribution of space capsules.

Cashwell: I see, very interesting.

Coveyou: The radial distributions you are talking about here are in two dimensions?

Cashwell: Two dimensions, right.

Coveyou: I am going to stick my neck out here. Can you find any one-dimensional distribution which is not a marginal of a radial distribution? I have a feeling that any one-dimensional function which is a distribution is the marginal distribution of some radial distribution, and I think that is true regardless of the number of dimensions.

Cashwell: You don't believe my necessary and sufficient conditions?

Coveyou: No, I think that you could find a case where they were not satisfied.

Cashwell: But there are examples in my paper.

Coveyou: Examples of one-dimensional distributions which are not marginal distributions of any radial distribution? My original intuitive feeling

is that simply any one-dimensional distribution (any one-dimensional distribution that is symmetric, of course) has got to be the marginal of some radial distribution. Now you can find functions which do not satisfy your condition, but are they distribution functions?

Cashwell: Can you produce the radial distribution for any given one-dimensional function?

Coveyou: I think I can write down a formula for it. In fact, I suspect that more than this is true, that if you take any two dimension numbers, one smaller than the other, and take a distribution in the smaller-dimensional space, then you can represent it as a marginal distribution of a radial distribution in the larger dimensional space. I don't see how it can fail. I think you can write down a formula for it.

Cashwell: I would be very interested in seeing it.

Coveyou: Another comment is that I think this has been done in one dimension versus three. Isn't this what is usually called the Blizzard transform, or something of the sort? And I believe that case, the transformation from one dimension to three dimension, is discussed in Feller. I think I have seen it there, either that or something that is very closely related.

Kalos: Suppose the one-dimensional function is 0 at the origin and rises as x increases. How can that be the marginal distribution of a radial distribution function? It is out of the question.

Coveyou: The density may have to be decreasing.

Kalos: That is right.

Summary of Subsequent Discussion: Cashwell pointed out, and discussed, counterexamples in his paper, i.e. examples of one-dimensional distribution functions which are marginal distributions of any radial distribution.

Gelbard: I am curious whether Bell said anything about where he encountered this problem?

Cashwell: Very little is said in the paper but it is apparently in connection with a beam of particles.

Gelbard: That is essentially what Mal Kalos said.

Cashwell: It is what Kalos said. They had a projected density and wanted to infer the radial density function.

Gelbard: Dr. Borgwaldt wanted to make some points on another subject.

Borgwaldt: Bob Gast asked to hear from other people how they are dealing with the eigenfunction strategy. I will try to explain the scheme we are implementing now. We have decided to have a fixed, rather small, number of fission neutron sites with varying weight. I think that, if we had a stochastic operator, we could perform an eigenfunction iteration with a single neutron.

Now, our operator is not stochastic, so what we are trying to do is to modify the operator in such a fashion as to give it a more stochastic character. We are trying to decouple the source iteration procedure from the estimation procedure. We are trying to develop an estimation procedure which uses each source neutron site once, and only once as a contributor to an estimate, and to estimate errors along the lines which Kalos has explained. The source iteration procedure we are developing is a random walk in a $4 \times n$ dimensional space, where n is the number of the fission neutron sites and 4 is the number of the coordinates. This random walk procedure is formulated as a controlled procedure where the control parameter is the sum of the neutron weights. We try to control this process in such a fashion that the total neutron weight of these n fission neutrons is equal to one. And, on the other hand, we try to formulate the process in such a fashion that this control of the neutron weights will not lead to a bias which is higher than that which we can tolerate. The control procedure will hopefully show when the neutron number in our calculation is too small, because then the controller will hunt. We are using a limited controller which will warn us when the requirements for control are too high on account of the low number of neutrons. Essentially what we are doing is this. If we have, let us say, 100 neutrons, we start them one after another. If 85 neutrons are sufficient to create a new generation of 100 new source sites, then we prevent the remaining 15 neutrons from creating progeny, but they are used for the estimation procedure. The difference between these numbers, 100 and 85, is used to create a new weight for the following neutron generation, for giving a weight to the new generation. On the other hand, if we start the 100 neutrons and still have a deficit in progeny, we start a few neutrons a second time. These additional starters are not used for the estimation procedure, but only for creating progeny. We hope to be able to prove that this scheme will have a bias which is a second-order effect, that the first-order effect will cancel out. The implementation has not yet been completed. Our first tests show that this method works quite nicely, but I still have to add some minor corrections. The complete scheme which will guarantee that each neutron is used once and only once, for estimation is not yet in operation, but our remaining problems are simply programming problems.

Kalos: By a second-order bias you mean a bias that is of order $1/n^2$? Is that right?

Borgwaldt: Yes.

Kalos: That would be very nice.

Gelbard: But you don't mean to assert that your scheme would be unbiased in the limit where you used only one history per generation?

Borgwaldt: I would say that one cannot use such a scheme on a one-neutron basis. That is clear. But I believe that with such a procedure the number of neutrons can be significantly smaller than the 500 neutrons mentioned by Gast. But I must concede that the situation in a large thermal reactor is different from the situation in a fast reactor. We are thinking about 50 to 100 neutrons per generation in fast criticals, like the ZPR-348, for example.

COARSE-MESH REBALANCING ACCELERATION FOR EIGENVALUE PROBLEMS

T. Asaoka, Y. Nakahara, S. Miyasaka, K. Horikami,
T. Suzuki, T. Nishida, Y. Taji, K. Tsuchihashi,
H. Gotoh, Y. Seki, M. Nakagawa and J. Hirota

Japan Atomic Energy Research Institute
Tokai, Ibaraki, Japan

* * *

Paper was not presented.

COARSE-MESH REBALANCING ACCELERATION FOR EIGENVALUE PROBLEMS

T. Asaoka, Y. Nakahara, S. Miyasaka, K. Horikami,
T. Suzuki, T. Nishida, Y. Taji, K. Tsuchihashi,
H. Gotoh, Y. Seki, M. Nakagawa and J. Hirota

Japan Atomic Energy Research Institute
Tokai, Ibaraki, Japan

ABSTRACT

The coarse-mesh rebalance method is adopted for Monte Carlo schemes for aiming at accelerating the convergence of a source iteration process. At every completion of the Monte Carlo game for one batch of neutron histories, the scaling factor for the neutron flux is calculated to achieve the neutron balance in each coarse-mesh zone into which the total system is divided. This rebalance factor is multiplied to the weight of each fission source neutron in the coarse-mesh zone for playing the next Monte Carlo game.

The numerical examples have shown that the coarse-mesh rebalance Monte Carlo calculation gives a good estimate of the eigenvalue already after several batches with a negligible extra computer time compared to the standard Monte Carlo.

INTRODUCTION

The Monte Carlo method has recently become more and more to play an important role in a wide field of theoretical studies of reactor physics because of the development of the computer softwares and hardwares. In addition, it gives a benchmark to check other mathematical models as well as the physics data attached to them. This tool is, however, not yet handled as skilfully as one might do, mainly due to a large calculation time required for obtaining the result within a reasonably small statistical error. To improve this situation, an effort will be required for developing further new biasing methods and their theoretical foundations, as well as practical methods for obtaining usable sampling functions [1].

In eigenvalue problems, also the acceleration of the convergence rate of iterative processes is essential for enlarging the range of applicability

of the method. This situation is the same as for deterministic methods for which the outer or source iterations are required for solving problems. An acceleration algorithm of iteration processes having been developed for deterministic methods can therefore be applied in principle also to Monte Carlo calculations with iterative batch processes.

CONVERGENCE ACCELERATION ALGORITHMS

Many acceleration techniques have been suggested for use with iterative deterministic methods for solving the steady state transport equation. The Chebychev acceleration is used most often in one-dimensional problems and is usually quite effective. In two-dimensional geometries, however, the technique is less satisfactory by reasons which are not throughly understood. The most effective acceleration method is generally the coarse-mesh rebalance method though there are problems for which the use of this technique leads to an unstable algorithm. The synthetic method as an acceleration technique has similar characteristics and these two methods are in fact equivalent, provided that the balance equation is used as a low-order equation for the synthetic algorithm [2].

For applying the acceleration algorithm of the outer iterations to the Monte Carlo method, consideration should be taken not to consume an extra time for executing the algorithm in computer codes. This suggests that the coarse-mesh acceleration technique is most suited to the Monte Carlo calculation, if the homogeneous regions of the reactor system are chosen as coarse-mesh zones.

The source iteration method can be written in the form:

$$L\varphi_{n+1} = S\varphi_n , \tag{1}$$

where φ_n stands for the angular flux at the n-th outer iteration. The low-order equation to it is a balance equation written as

$$\tilde{L}_n \tilde{\varphi}_{n+1} = \tilde{S}_n \tilde{\varphi}_n . \tag{2}$$

For practical applications, the source operator S can be assumed to be isotropic and hence $S\varphi_n$ depends only on the scalar flux ϕ_n. The geometrical system is divided into the coarse-mesh zones and the transport equation is integrated over all directions and over a coarse-mesh zone. The resulting equations, one for each coarse-mesh zone, involve the scalar fluxes within the coarse-mesh zone and the currents across the coarse-mesh boundaries:

$$I(r) = \int d\underline{\Omega}\, (\underline{\Omega}\cdot\underline{n})\, \varphi(r, \underline{\Omega}). \tag{3}$$

Therefore, a low-order approximation to the equation for the scalar flux ϕ is sought instead of Eq.(1) for φ. Upon assuming that an approximation $\tilde{\varphi}_n$ (and therefore $\tilde{\phi}_n$) is available and the flux shape within a coarse-mesh zone is that of $\tilde{\phi}_n$, then the additional assumption that the current obtained from a relationship of $\tilde{\phi}_n$ must match $I(r)$ from $\tilde{\varphi}$ leads to a set of equations

for the flux amplitude or a set of balance equations.

The iteration proceeds as follows. Given Φ_n (therefore $S\Phi_n$), $\tilde{\Phi}_n$ is then obtained from

$$L\tilde{\Phi}_n = S\Phi_n . \tag{4}$$

The currents $I(r)$ are computed from appropriate integrals of $\tilde{\Phi}_n$. A set of balance equations are formulated for these fluxes and currents, and the scalar flux Φ_{n+1} is obtained by solving these equations.

APPLICATION OF COARSE-MESH REBALANCE METHOD

Since in Monte Carlo games neutrons crossing the boundaries between different media are sampled for the determination of their flight distances, the homogeneous regions can conveniently be chosen as the coarse-mesh zones. The balance equation for obtaining the scale factor f_ℓ, by which all fluxes in the region ℓ are multiplied to insure the neutron balance, may then be written as

$$f_\ell \left[\sum_{\ell'} (CUR)_{\ell\to\ell'} + (AB)_\ell + (SD)_\ell \right] = (QQ)_\ell + \sum_k f_k (CUR)_{k\to\ell} , \tag{5}$$

where, upon denoting the weight of the h-th history neutron at the j-th collision point by W_{hj},

$(CUR)_{\ell\to\ell'} = \sum_{h=1}^{H} \sum_j W_{hj}^{\ell\to\ell'}$; the total neutron current across the boundary from the region ℓ to ℓ',

$(AB)_\ell = \sum_h \sum_j (\Sigma_a/\Sigma_t)_\ell W_{hj}$; the total number of neutrons absorbed in the region ℓ,

$(SD)_\ell = \sum_h W_{hJ}^{SD}$; the total number of neutrons slowed down below the cut-off energy in the region ℓ (J stands for the last collision number for the h-th history neutron),

$(QQ)_\ell$ = the total number of fission source neutrons for the region ℓ, which is normalized as

$\sum_\ell (QQ)_\ell$ = (the total number of histories for the present batch, H) $\times$ (initial weight of each source neutron, W_0) = 1.

The source term can be rewritten as

$$(QQ)_\ell = f_\ell (FS)_\ell / k_{eff} , \tag{6}$$

where $(FS)_\ell = \sum_h \sum_j (\nu\Sigma_f/\Sigma_t)_\ell W_{hj}$; the total number of fission neutrons produced in the region ℓ.

In this case, a source iteration is required for determining the rebalance factors and it gives a better estimate of the eigenvalue k_{eff}.

Equation (5) is solved by using a forward elimination and backward substitution [3]. In one-dimensional geometries, for example, by assuming

$$\sum_{\ell'} (CUR)_{\ell\to\ell'} = (FL)_\ell + (FR)_{\ell+1},$$
$$\sum_{k} (CUR)_{k\to\ell} = (FR)_\ell + (FL)_{\ell+1}, \tag{7}$$

where $(FL)_\ell$ is the total neutron current crossing leftwards the left boundary of the region ℓ from the ℓ to $(\ell-1)$-th region, and $(FR)_{\ell+1}$ is the current crossing rightwards the right boundary of ℓ from ℓ to $\ell+1$, the set of Eq.(5) can be rewritten as

$$\begin{bmatrix} a_1 & -b_1 & & & 0 \\ -c_2 & a_2 & -b_2 & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & -b_{L-1} \\ 0 & & & -c_L & a_L \end{bmatrix} \begin{bmatrix} f_1 \\ f_2 \\ \vdots \\ \vdots \\ f_L \end{bmatrix} = \begin{bmatrix} QQ_1 \\ QQ_2 \\ \vdots \\ \vdots \\ QQ_L \end{bmatrix}, \tag{8}$$

where

$$a_\ell = (FL)_\ell + (FR)_{\ell+1} + (AB)_\ell + (SD)_\ell,$$

$$b_\ell = (FL)_{\ell+1}, \qquad c_\ell = (FR)_\ell.$$

The forward elimination is therefore performed as (with $H_0 = G_0 = 0$)

$$T_\ell = a_\ell - c_\ell H_{\ell-1}, \quad G_\ell = (QQ_\ell + c_\ell G_{\ell-1})/T_\ell,$$

$$H_\ell = b_\ell / T_\ell, \quad \ell = 1, 2, ----, L, \tag{9}$$

and then the backward substitution is carried out as follows:

$$f_L = G_L,$$

$$f_\ell = G_\ell + H_\ell f_{\ell+1}, \quad \ell = L-1, ----, 2, 1. \tag{10}$$

If the source iteration for Eq. (5) with (6) fails to converge, if any factor f_ℓ becomes negative, or, if any denominator T_ℓ becomes zero, all f_ℓ's are set equal to the single-system factor f given by

$$f = \sum_\ell (QQ)_\ell / [\, TCR + \sum_\ell (AB)_\ell + (SD)_\ell \,], \qquad (11)$$

where TCR stands for the total neutron flow leaking out of the whole system.

The rebalancing based on Eq. (5) with (6) will be required only for the first few outer iterations (batches) of the Monte Carlo game and thereafter Eq. (5) with the fixed source $(QQ)_\ell = (FS)_\ell / k_{eff}$ will be enough for obtaining the rebalance factors. As the number of batches increases, the factors all usually approach unity and the deviations from unity will become much smaller than the standard deviations of the quantities of Eq. (5), for example, the standard deviation of k_{eff} as a typical representative of the quantities:

$$(\overline{k_{eff}^2} - \bar{k}_{eff}^2)^{1/2} / \sqrt{N} \gg \max_\ell |1 - f_\ell|$$

where N stands for the total number of batches. In this case, Eq. (11) will also be applied for achieving the whole system rebalance or the rebalance technique will not be necessary any more.

NUMERICAL EXAMPLES AND DISCUSSIONS

The present coarse-mesh rebalance technique has been incorporated in the MORSE computer code [4] by attaching two new subroutines. One is the COARSE subroutine for evaluating the quantities of Eq. (5) and the other is the REBAL for computing the rebalance factors according to Eqs. (9) and (10), or (11) if required. After obtained the converged values of f_ℓ's as a result of i iterations in the REBAL, the weight of each fission neutron produced in the region ℓ, stored in the fission bank in the MORSE code, is multiplied by $\prod_{k=1}^{i} f_\ell^{(k)} k_{eff}^{(0)} / k_{eff}^{(1)}$, where $k_{eff}^{(0)}$ and $k_{eff}^{(1)}$ are respectively the k_{eff} values obtained in the MORSE prior to the rebalance calculation and in the REBAL after the i iterations.

In order to test the effectiveness of the coarse-mesh rebalance acceleration applied to Monte Carlo eigenvalue calculations, the k_{eff} of a homogeneous spherical fast reactor of 55 cm radius was calculated with a 5-energy-group model by dividing the system into two coarse-mesh zones at a radius of 35 cm.

The results for the average values of k_{eff} and their standard deviations are shown in Fig. 1 as a function of the total number of neutron histories, where the first batch is excluded from calculating the averages. As seen from Fig. 1, after several batches the values of $k_{eff}^{(1)}$ obtained from the coarse-mesh rebalance Monte Carlo (white circles in Fig. 1) are already very close to the S_{12} value calculated by the ANISN computer code [5], while the standard Monte Carlo by using the original MORSE (black circles) gives still a slightly too low value after about 4,000 histories. It should be noted that the coarse-mesh rebalance Monte Carlo requires only a negligible extra computer time (only 2 % more per collision) compared to the standard

calculation.

In Table I are shown the final results for the total number of collisions suffered in each coarse-mesh zone and the number of fission neutrons produced in each zone. The Monte Carlo results obtained from about 4,000 histories agree well with the S_{12} values.

As another example, the calculations were performed on a two-region spherical fast reactor composed of the core of 35 cm radius and the blanket of 55 cm outer radius. The coarse-mesh zones are chosen as the same as these two regions.

It is seen from Fig. 2 that in this case the coarse-mesh rebalance version is much better than the standard version, though neither of two calculations did not converge yet to the S_{12} value. The difference between the Monte Carlo results for k_{eff} and the S_{12} value comes mainly from the underestimation of the fourth group contribution in the Monte Carlo calculations (see Table II). This fact indicates that it is desirable to sample more neutron histories per batch than used here. It may be worthwhile to mention here again that the coarse-mesh rebalance Monte Carlo calculation takes more computer time by only 2 % per collision than that of the standard calculation.

CONCLUSIONS

It has been shown in this article that the coarse-mesh rebalance method can successfully be applied to Monte Carlo calculations for accelerating the convergence of a source iterative process. Even for the present numerical examples on spherical systems with only two coarse-mesh zones, the eigenvalue obtained from the coarse-mesh rebalance Monte Carlo calculation approaches quickly to the correct value. The computer time required for the extra rebalance calculation is negligibly small compared to that for the standard Monte Carlo calculation.

It is therefore expected that the coarse-mesh rebalance technique is certainly more effective for complex systems composed of different media. In addition, the technique will be useful also for estimating the neutron flux at a point detector position.

REFERENCES

1. F. A. R. SCHMIDT, "Status of the Monte Carlo Development," in Proc. Seminar on Numerical Reactor Calculations, International Atomic Energy Agency, Vienna (1972), p. 699.

2. Wm. H. REED, "The Effectiveness of Acceleration Techniques for Iterative Methods in Transport Theory," in Proc. of the Second Conference on Neutron Transport Theory, USAEC Rep. CONF-710107, Los Alamos Scientific Laboratory (1971), p. 391.

3. K. D. LATHROP and F. W. BRINKLEY, "TWOTRAN-II: An Interfaced, Exportable Version of the TWOTRAN Code for Two-Dimensional Transport," USAEC Rep. LA-4848-MS, Los Alamos Scientific Laboratory (1973).

4. E. A. STRAKER et al., "The MORSE Code — A Multigroup Neutron and Gamma-Ray Monte Carlo Transport Code," USAEC Rep. ORNL-4585, Oak Ridge National Laboratory (1970).

5. W. W. ENGLE, Jr., "A User's Manual for ANISN, A One Dimensional Discrete Ordinates Transport Code with Anisotropic Scattering," USAEC Rep. K-1693, Union Carbide Corporation (1967).

TABLE I

The Total Number of Collisions and the Number of Fission Neutrons Produced in Each Coarse-Mesh Zone of a Homogeneous Spherical Fast Reactor with a Unit Neutron Source

Calculation Method	Energy-Group Index	$\int dV\ \Sigma_t\phi$		$\int dV\ \nu\Sigma_f\phi$	
		First Zone	Second Zone	First Zone	Second Zone
Standard Monte Carlo	1	0.043	0.049	0.0073	0.0083
	2	1.234	1.210	0.1148	0.1128
	3	5.256	5.111	0.1938	0.1880
	4	6.075	5.738	0.2199	0.2068
	5	0.733	0.574	0.0309	0.0242
Coarse-Mesh Rebalance Monte Carlo	1	0.046	0.050	0.0079	0.0086
	2	1.284	1.251	0.1193	0.1160
	3	5.298	5.250	0.1957	0.1938
	4	5.871	0.537	0.2114	0.1995
	5	0.759	0.611	0.0320	0.0257
S_{12}	1	0.049	0.052	0.0085	0.0089
	2	1.181	1.230	0.1099	0.1144
	3	5.172	5.243	0.1905	0.1931
	4	5.867	5.723	0.2120	0.2068
	5	0.710	0.679	0.0299	0.0286

TABLE II

The Total Number of Collisions and the Number of Fission Neutrons Produced in Each Region of a Two-Region Spherical Fast Reactor with a Unit Neutron Source

Calculation Method	Energy-Group Index	$\int dV\ \Sigma_t\phi$		$\int dV\ \nu\Sigma_f\phi$	
		Core	Blanket	Core	Blanket
Standard Monte Carlo	1	0.057	0.035	0.0097	0.0111
	2	1.634	0.759	0.1510	0.0929
	3	7.276	7.976	0.2677	0.0188
	4	8.100	10.104	0.2924	0.0240
	5	1.013	0.090	0.0427	0.0004
Coarse-Mesh Rebalance Monte Carlo	1	0.060	0.032	0.0103	0.0103
	2	1.666	0.696	0.1553	0.0851
	3	7.675	7.940	0.2827	0.0187
	4	8.432	10.048	0.3055	0.0239
	5	0.996	0.083	0.0419	0.0004
S_{12}	1	0.078	0.032	0.0133	0.0103
	2	1.778	0.610	0.1653	0.0747
	3	8.021	7.303	0.2955	0.0172
	4	9.032	9.728	0.3264	0.0231
	5	1.003	0.110	0.0422	0.0005

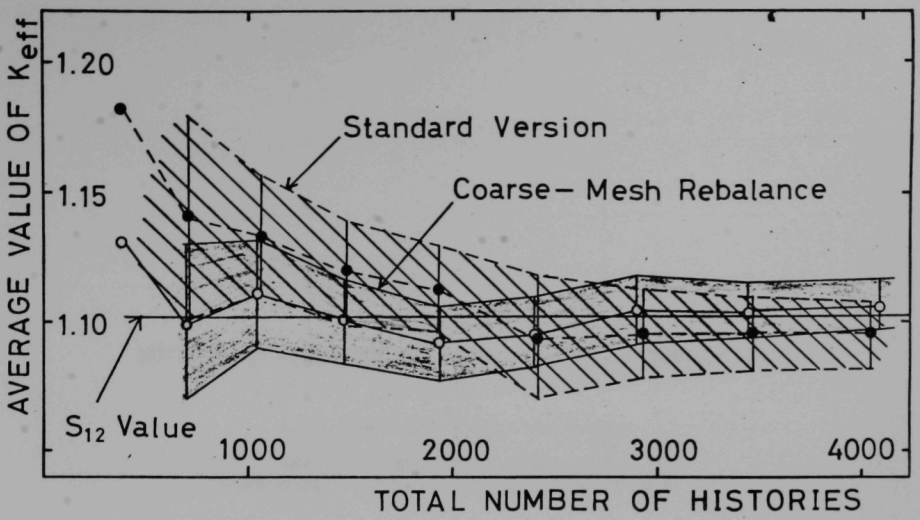

Fig. 1. The average value of k_{eff} and its standard deviation of a homogeneous spherical fast reactor.

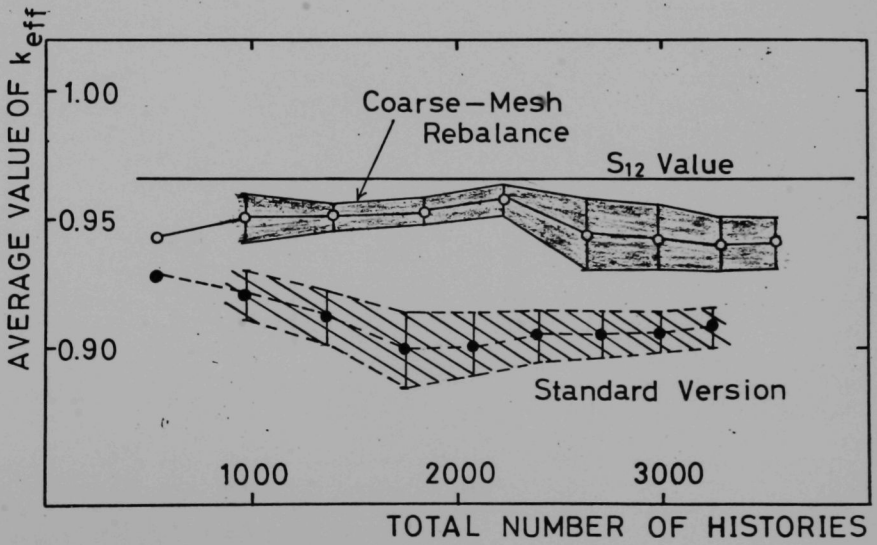

Fig. 2. The average value of k_{eff} and its standard deviation of a two-region spherical fast reactor.

MONTE CARLO EIGENFUNCTION STRATEGIES AND UNCERTAINTIES

R. C. Gast and N. R. Candelore

Bettis Atomic Power Laboratory
West Mifflin, Pennsylvania 15122, U.S.A.

* * *

Presented by R. C. Gast

MONTE CARLO EIGENFUNCTION STRATEGIES AND UNCERTAINTIES

R. C. Gast and N. R. Candelore

Bettis Atomic Power Laboratory
West Mifflin, Pennsylvania 15122, U.S.A.

ABSTRACT

Comparisons of convergence rates for several possible eigenfunction source strategies led to the selection of the "straight" analog of the analytic power method as the source strategy for Monte Carlo eigenfunction calculations. To insure a fair game strategy, the number of histories per iteration increases with increasing iteration number. The estimate of eigenfunction uncertainty is obtained from a modification of a proposal by D. B. MacMillan and involves only estimates of the usual purely statistical component of uncertainty and a serial correlation coefficient of lag one.

I. INTRODUCTION

It is convenient to consider each possible Monte Carlo eigenfunction strategy as composed of two basic parts. The first part is the totality of the Monte Carlo procedures that would be involved in a fixed source calculation. The second part is the eigenfunction source shape strategy that connects successive generations.

The construction of an efficient overall eigenfunction strategy requires that the fixed source procedures, and the method of connecting successive generations, each be made efficient. Methods for improving the efficiency, i.e. reducing the standard deviation per unit machine time, in fixed source procedures have been frequently investigated and developed. In contrast, methods for reducing the variance arising from the source shape strategy have received relatively little attention; even the basic requirements for a "fair game" strategy appear to have been little understood.

All source shape strategies considered in this report are required by the authors to be "fair game" strategies. Thus, the first subsection of Section II is a discussion of "fair game" requirements. This subsection also includes a description of a "reference" strategy which is used to provide a comparison of relative convergence rates of all other strategies. The remainder of Section II describes various source shape strategies considered; they include strategies using Green's function, strategies using

extrapolation factors, and strategies using importance sampling. The final subsection of Section II discusses the details of the source strategy selected.

Having selected a definite strategy, Section III of the report then describes the approach that led to the means of assigning eigenfunction uncertainties. The resulting uncertainty formula is partly empirical, and for that reason it required extensive numerical testing. The results of these tests are also presented in Section III.

II. EIGENFUNCTION SOURCE STRATEGY

The investigation that led to the selected source shape strategy is described in this section. First, however, the "fair game" requirements for a general eigenfunction strategy are discussed and a reference strategy is defined.

a. Requirements for Fair Game and a Reference Strategy

The presentation in this subsection will summarize some results given in detail in Ref. [1].

First, we introduce for convenience the term "limit in probability" denoted by the symbol "PLim." Thus $\underset{N\to\infty}{\text{PLim}}\, X_N = X$ means that for the sequence of random variables X_N, and for some parameter X, the condition $\text{Lim}\,[P\{|X_N - X| < \epsilon\}] = 1$ is met for any arbitrarily chosen positive ϵ. In this case we say that the sequence of random variables, X_N, converges in probability to X. The concept of repeatability involved in determining $P\{|X_N - X| < \epsilon\}$ for a given value of $N = N_o$, where X_N is a sequence of dependent random variables, is to repeatedly return to the calculation beginning with new sets of random numbers and obtain a sequence of values for X_{N_o}. The fraction of such values for which $|X_{N_o} - X| < \epsilon$ is an estimate of the desired probability.

We may write any overall Monte Carlo eigenfunction strategy using a neutron production Green's function as

$$\psi_i^{n+1} = \sum_j G_{ij}^n \, K^n(S_j^n).$$

The three random variables in this equation are defined as follows:

1) G_{ij}^n is an estimate of the fission neutron production rate in elementary volume i due to a unit fission source in elementary volume j in the nth iteration. Possible values of this random variable are determined by starting one or more neutrons in elementary volume j and determining the neutron production rate per start neutron in elementary volume i by Monte Carlo tracking. Consequently, the number of neutron histories in a value of G_{ij}^n will be the number of starting neutrons in elementary volume j in the

n^{th} iteration. If the number of histories in a value of G^n_{ij} approaches infinity, then the limit in probability of G^n_{ij} is G_{ij} where G_{ij} is the analytic Green's function over elementary volumes.

2) ψ^{n+1}_i is an estimate of the fission neutron production rate in elementary volume i at the nth iteration of the iteration process used. As implied by the above equation, ψ^{n+1}_i is a dependent random variable.

3) $K^n(S^n_j)$ is another estimate of the fission neutron production rate in elementary volume j at the nth iteration, but in general, is only relative to other elementary volumes. S^n_j alone is another estimate of the fission neutron production rate in elementary volume j at the nth iteration, and is obtained in some manner, yet to be specified, from previous iteration results. One possible means of doing this is the analog of the analytic power method, i.e. set $S^n_j = \psi^n_j$. The symbol $K^n(S^n_j)$ denotes an operator K^n acting of S^n_j. In general, the operator K^n involves a sampling process and a normalization process, although one or both of these processes may be absent. The sampling operator K^n in the Monte Carlo eigenfunction iteration process defines the means by which neutron starting sites and weights are determined in the nth iteration. For example, a possible definition of the operator K^n which involves both sampling and normalization processes is as follows: First a finite number of neutron starting sites are distributed among the elementary volumes using $\dfrac{S^n_j}{\sum S^n_j}$ as the frequency distribution function and using a set of random numbers designated by the index n. Second, the sum of neutron weights (say unity per site) over the starting sites are normalized to unity. If we let N^n be the number of neutron histories in iteration n, and if we let S_j be a source strength per elementary volume j that is independent of n (i.e. a fixed source), then we may write

$$\underset{n\to\infty}{\mathrm{PLim}}\left(\frac{\sum_{n=1}^{n} N^n K^n(S_j)}{\sum_{n=1}^{n} N^n}\right) = \bar{S}_j,$$

where $\bar{S}_j$ are normalized values of S_j, i.e. $\sum_j \bar{S}_j = 1$.

The only sampling operator used in this report is that given above as an example. In the discussion of this particular operator, we did not identify the precise means by which starting sites are selected using

$\dfrac{S_j^n}{\sum_j S_j^n}$ as the frequency distribution function. For example, a purely random sampling could be used, or one that is purely systematic, or some combination of both (semi-systematic). We do not need to be concerned with these details, since all that is required to demonstrate a fair game is that the above PLim equation is satisfied; that is, the sampling operator will properly treat a fixed source.

We will now exhibit the condition to be met for a fair game in a Monte Carlo eigenfunction strategy. Let Ψ_i be the fundamental mode eigenfunction neutron production rate over elementary volume i from analytic transport theory. Thus, Ψ_i satisfies the equation

$$\Psi_i = \sum_j G_{ij} \overline{\Psi}_j,$$

where G_{ij} is the analytic Green's function over elementary volumes, and where $\overline{\Psi}_j = \dfrac{\Psi_j}{\sum_j \Psi_j}$. Here the eigenvalue, λ, is a normalization factor, i.e. $\lambda = \sum_j \Psi_j$. Using PLimits, the condition to be met for a fair game in an eigenfunction strategy where iteration results are accumulated with weights equal to number of histories per iteration, N^n, is given by

$$\underset{n\to\infty}{\mathrm{PLim}}\left(\frac{\sum_{n=1}^{n} N^n \Psi_i^n}{\sum_{n=1}^{n} N^n}\right) = \Psi_i.$$

As discussed previously, Ψ_i^n satisfies an equation of the form

$$\Psi_i^{n+1} = \sum_j G_{ij}^n K^n(S_j^n).$$

It may be observed at this point that it will generally be desirable to accumulate iteration results with weights equal to number of histories per iteration, since this will result in minimum variance in a fixed source

calculation.

The simplest Monte Carlo eigenfunction strategy is the analog of the analytic power method, i.e. set $S_j^n = \Psi_j^n$. This strategy will be taken as the reference strategy, i.e. all other strategies will be compared to the power method to judge relative convergence rates. As shown in Ref. [1], this power method strategy is not a fair game if the number of histories per iteration, N^n, is a constant independent of n. (The appendix of this report gives a weighting scheme for neutron starting weights such that a fair game is obtained in this case). A sufficient set of conditions to insure a fair game in the power method strategy are:

(1) $\lim_{n\to\infty} N^n = \infty$.

(2) The sampling operator properly treats a fixed source. Since the number of histories per iteration becomes large as n increases (condition 1), we must have $\mathrm{PLim}_{n\to\infty} K^n(S_j) = \overline{S}_j$.

(3) Because the strategy is a Markov process, we must require that all iterations used in the calculation produce progeny. This is easily done with the device that should an iteration not produce progeny, then the first prior iteration with progeny is used to continue the process.

(4) The physical problem under consideration must be one for which the analytic neutron transport equation has a unique fundamental mode eigenfunction. That is, certain sufficient "connectivity" conditions among the elementary volumes are met which insure such a unique eigenfunction. These connectivity conditions will be met for all reactor problems of practical interest.

With these conditions, it follows essentially by definition that the resulting "power method" strategy is a fair game. This is easily seen by considering the representation of the nth iteration of the Monte Carlo calculation, i.e.

$$\Psi_i^{n+1} = \sum_j G_{ij}^n \, K^n(\Psi_j^n).$$

Because of conditions (3) and (4), we know that the $\mathrm{PLim}\, \Psi_i^n$ does exist. Let its value be Ψ_i. Taking PLimits of both sides of the above equation, we immediately obtain

$$\Psi_i = \sum_j G_{ij} \, \overline{\Psi}_j .$$

This is the analytic neutron transport equation over elementary volumes. Since negative neutrons are not permitted to be born at fission sites, and since the trivial zero solution is not possible because of condition (3),

then the above Ψ_i must be the fundamental mode eigenfunction. Thus the general condition for a fair game is met, i.e.

$$\underset{n\to\infty}{\mathrm{PLim}}\left(\frac{\sum_{n=1}^{n} N^n \Psi_i^n}{\sum_{n=1}^{n} N^n}\right) = \Psi_i,$$

where Ψ_i is the fundamental mode eigenfunction over elementary volumes from analytic transport theory.

All of the source shape strategies considered in the remainder of this section are made fair games by the same approach used for the power method described above. That is, the number of histories per iteration will increase with increasing iteration number.

The subsequent subsections will describe the new eigenfunction strategies that were considered and the results of comparing their convergence rates to the reference strategy. In order to make these comparisons more meaningful, a source guess accuracy requirement was adopted. However, there is another obvious reason for adopting such a requirement. That is, in any conceivable eigenfunction strategy, one cannot permit an arbitrarily poor source guess in a large core and obtain satisfactory convergence in reasonable computing times. Thus in all strategies considered, we will require that the eigenfunction guess be within about 2 P.E.'s of the true eigenfunction, where the P.E. is computed for the number of histories in the 1st iteration. This condition can be easily met, since such a guess is not highly accurate. A workable method of verifying that the source guess is of this accuracy for practical size cores is to compare the eigenfunction estimate at the end of the problem with the guess. Should this comparison reveal that the guess was not within 2 P.E.'s, then a sequence of initial iterations should be omitted until the source accuracy condition is met. This procedure will be carried out in all subsequent testing of eigenfunction strategies.

b. Source Strategies Using Green's Function

An accumulated neutron production Green's function in an eigenfunction Monte Carlo calculation, say $G^n_{jj'}$, is defined as the estimate of the fission neutron production rate in region j due to a unit fission source in region j′. This matrix is the result of normalizing an accumulative matrix formed by adding to a corresponding matrix element the fission neutron production as each neutron collision occurs. The normalization is performed by row; the normalization factor for the j′ row being the reciprocal of the total number of neutrons born in region j′. The index n indicates the total number of iterations contributing to the accumulation. Thus if the Monte Carlo eigenfunction calculation is a fair game, i.e. capable of converging to the true transport theory solution, then at any stage of such a calculation an estimate of the region integrated eigenfunction, say P_j, and an

estimate of the eigenvalue, say $\hat{\lambda}$, may be made using the above Green's function. Of course these estimates would be in addition to those already available directly from the Monte Carlo calculation. The relative merits of these two types of estimates will be discussed later.

Knowing $\hat{G}^{n-1}_{jj'}$, at the end of iteration n-1, we may obtain the estimated, i.e. projected, eigenvector $\hat{P}_j$ and the projected eigenvalue $\hat{\lambda}$ by the analytic power method. That is, we define $\hat{P}^{\ell}_j$ as,

$$\hat{P}^{\ell}_j = \sum_{j'} \hat{G}^{n-1}_{jj'} \, \bar{\hat{P}}^{\ell-1}_{j'}, \text{ where}$$

$$\bar{\hat{P}}^{\ell}_j = \frac{\hat{P}^{\ell}_j}{\sum_j \hat{P}^{\ell}_j} .$$

The above iterative process is carried out by starting with some eigenvector guess and continuing until some convergence criterion for the eigenvector is met. Let L be the total number of iterations so performed. Then $\hat{P}^{L}_j$ is the projected eigenvector and

$$\sum_j \hat{P}^{L}_j \quad \text{is the projected eigenvalue.}$$

Using the representation of a general Monte Carlo eigenfunction strategy from subsection (a), i.e.

$$\psi^{n+1}_i = \sum_j G^n_{ij} \, K^n(S^n_j),$$

we may hopefully define an accelerated strategy by setting the source for the nth iteration as

$$S^n_j = \hat{P}^{L}_j .$$

Of course, this could be done with some chosen regularity during the course of the calculation rather than every iteration.

Use has been made of the projected eigenvalue, and the quantity L in conjunction with the O5R Monte Carlo program. For example, G. W. Morrison, J. T. Mihalczo, and D. C. Irving [2] use the quantity L to decide how many initial Monte Carlo iterations do not have a sufficiently converged source in order that they may be excluded from the eigenfunction accumulation. Of

course, in this application, the eigenvector guess, by region j, would be taken as the original source guess in the Monte Carlo calculation.

The projected, i.e. matrix, eigenvalue has been used in connection with 05R criticality calculations. This estimated eigenvalue is of interest since it is a separate estimate from that obtainable directly from the Monte Carlo calculation, i.e. the ratio of total neutron production to total source. Of course these two estimates are not independent. Mihalczo, in his 05R calculations of delayed-critical assemblies [3], found that the matrix eigenvalue was always within a standard deviation of that computed from the ratio of total neutron production to total source. A similar conclusion was arrived at by M. R. Mendelson in his Monte Carlo criticality calculations using the KAPL version of 05R [4]. Mendelson found that if a reasonable source guess was used, then the matrix eigenvalue does not appear to be any better than that obtained from a ratio.

The idea of using a Green's function in Monte Carlo eigenfunction calculations to obtain better estimates of the eigenfunction and eigenvalue has been frequently discussed by Monte Carlo workers. It appears to have originated with K. W. Morton [5] and was further developed by E. L. Kaplan [6]; however, it also appears that their work was purely theoretical. This idea is especially appealing if one views it as follows: Consider a reactor broken up into a finite number of gross regions over which a Green's function matrix is tabulated. In applying the analytic analog of the power method to such a problem, the eigenfunction error would consist principally of the first overtone component.

It intuitively seems that the matrix eigenfunction should enable rapid removal of the lower overtone modes, since they would be easily sensed over gross regions. After local normalization, deviations from the eigenfunction interior to a gross region would be more dependent on the higher overtones, which die out rapidly. Thus it would seem that if the projected eigenfunction over gross regions at iteration n were used to determine the source for iteration n+1, then a considerable gain in convergence rate would be obtained. This differs from the approach used with 05R calculations, since knowledge of the projected eigenfunction was not used to determine the source for the subsequent iteration.

This idea has been more recently discussed by L. L. Carter and N. J. McCormick [7]. Again, however, it was not actually tried in a Monte Carlo calculation. Instead, they constructed what was believed to be a reasonable analytic representation of the Monte Carlo procedure using diffusion theory. In their analysis of convergence rates with this model, they considered the straight power method, the use of the Green's function as in 05R, and the use of the Green's function at iteration n to determine the source shape over gross regions for iteration n+1. They found that the 05R approach gave faster convergence than the power method and that the last approach above gave much faster convergence than the power method. It should be noted that the first part of their conclusion is in conflict with actual Monte Carlo calculations as performed by Mihalczo [3] and Mendelson [4].

Because of the intuitive appeal of using the Green's function at iteration n to determine source shape over gross regions for iteration n+1, the technique was tried. The projected eigenfunction as described earlier was

used to determine the source shape over gross regions in the power method. Semi-systematic sampling of actual sites within a gross region was used to accurately obtain the desired source strength by gross region. Also, the projected eigenfunction was made with a variety of periods, that is every iteration, every 2nd iteration, every 5th iteration, every 10th iteration.

In all cases that were tested, which covered a variety of core sizes, no gain in eigenfunction convergence rates could be detected relative to the straight power method. Since this Green's function procedure was more elaborate than the corresponding procedure used for O5R by Mihalczo and Mendelson, then by inference we have verified their conclusion of no gain in convergence rates with the simpler procedure. We must also conclude that the analytic modeling of Monte Carlo by Carter and McCormick was much too crude to permit reasonable estimates of gains in convergence rates using Green's functions.

In view of the complete failure of the Green's function approach, one might ask what is wrong with the intuitive argument that was given to support it. The best way to answer this is to state a proper view of the Green's function approach which is as follows: Within the straight Monte Carlo analog of the power method, the eigenfunction by gross region may be estimated by either the total cumulated source or the total cumulated fission neutron production. Since in the power method, the neutron production for the nth iteration becomes, within statistics resulting from the semi-systematic sampling, the source for the $n+1^{st}$ iteration, then it follows that the total cumulated source at the nth iteration will be very close in shape to the cumulated fission neutron production. This will occur even for n small if a reasonable source guess is made. The Green's function matrix elements are formed from the same neutron histories that contribute to the cumulated neutron production. The eigenvector from this matrix, by definition, is a source that will yield a neutron production of the same shape. But based on the information content of the matrix, i.e. all histories, this is already nearly the case since the cumulated source is very close in shape to the cumulated neutron production. Thus the eigenvector from the matrix will always be close to that obtained by the power method after very few iterations. That is, the Green's function approach cannot provide a worthwhile improvement is convergence rate relative to the power method in Monte Carlo.

c. Source Strategies Using Extrapolation Factors

The success of extrapolation factors in accelerating eigenfunction convergence in analytic diffusion theory calculations is well known; this suggests the application of such factors in Monte Carlo eigenfunction calculations.

First we will consider the use of a single linear extrapolation factor in the reference strategy, that is, a straight analog of the analytic power method. This will be done by obtaining the form of this factor for the analytic power method and applying it directly in a series of Monte Carlo test calculations. Making use of the neutron production Green's function, G_{ij}, as in subsection (a), we may write the analytic power method for the fundamental mode, using matrix notation and using subscripts for region indices, as

$$\lambda_0^{\ell} \bar{\psi}_i^{\ell} = G_{ij} \bar{\psi}_j^{\ell-1} \text{ for the } \ell\text{-}1^{st} \text{ iteration,}$$

where λ_0^ℓ is the multiplication factor for the ℓth iteration, i.e. $\lambda_0^\ell = \sum \psi_i^\ell$, and where bars denote normalization, e.g. $\bar{\psi}_i^\ell = \frac{\psi_i^\ell}{\Sigma\psi_i^\ell}$. When a single linear extrapolation factor, ω, is used in the power method to obtain the fundamental mode eigenfunction, then the above procedure is altered as follows. After obtaining $G_{ij}\psi_j^{\ell-1}$ from the ℓ-1st iteration, we normalize it by dividing by the multiplication factor, λ_0^ℓ, and define

$$\psi^\ell = \bar{\psi}^{\ell-1} + \omega\left(\frac{G_{ij}\bar{\psi}_j^{\ell-1}}{\lambda_0^\ell} - \bar{\psi}^{\ell-1}\right),$$

$$= \bar{\psi}^{\ell-1}(1-\omega) + \omega\,\frac{G_{ij}\psi^{\ell-1}}{\lambda_0^\ell}.$$

This ψ_i^ℓ is then used on the right hand side of the equation for the ℓth iteration, i.e.

$$\lambda_0^{\ell+1}\,\psi_i^{\ell+1} = G_{ij}\bar{\psi}_j^\ell.$$

For the Monte Carlo eigenfunction calculation, a corresponding procedure may be defined as follows. First write a general eigenfunction strategy as in subsection (a), i.e.

$$\Psi_i^{n+1} = \sum_j G_{ij}^n\, K^n(S_j^n).$$

After obtaining Ψ_i^n from the n-1st iteration, we define

S_i^n for the nth iteration as,

$$S_i^n = \bar{\Psi}_i^{n-1} + \omega\,(\bar{\Psi}_i^n - \bar{\Psi}_i^{n-1}).$$

The value of the extrapolation factor, ω, in the analytic calculation is defined by the following considerations. Let Ψ_i be a general eigenfunction of $\lambda\Psi_i = G_{ij}\Psi_j$. When the extrapolation procedure is used, the corresponding eigenfunction equation with eigenvalues μ is,

$$\mu_i = \Psi_i(1-\omega) + \omega \frac{G_{ij}\Psi_j}{\lambda_0},$$

where λ_0 is the fundamental mode eigenvalue, i.e. $\lambda_0 = \lim_{\ell \to \infty} \lambda_0^{\ell}$.

We may rewrite this equation as

$$\left(\frac{\mu-1+\omega}{\omega/\lambda_0} \right) \Psi_i = G_{ij}\Psi_j.$$

That is, the eigenfunction is the same but the eigenvalues μ have undergone a linear transformation,

$$\mu = \omega \frac{\lambda}{\lambda_0} + 1-\omega.$$

Let λ_n be the eigenvalue associated with the nth mode eigenfunction of $\lambda\Psi_i = G_{ij}\Psi_j$. We know that $\lambda_0 > \lambda_1 > \lambda_2 > \ldots \lambda_n > \ldots > 0$ and for large $n, \lambda_n \to 0$ (this is known from physical considerations; not mathematical ones). Since the fundamental mode eigenfunction error for $\omega = 1$ and large ℓ is proportional to $\left(\frac{\lambda_1}{\lambda_0}\right)^{\ell}$ then we may make use of the linear transformation to reduce the size of this quantity. If n can take on large values then the optimum transformation occurs when $\mu_1 = -\mu_n$ for n large, since in this case we will have the minimum value of Max $\left(\frac{\mu_n}{\mu_0}, n=1,2,\ldots\right)$. The value of ω to achieve this transformation is obtained as follows. For n large $\mu_n \to 1-\omega (\omega > 1)$. Also we know that $\mu_1 = \omega \frac{\lambda_1}{\lambda_0} + 1-\omega$. Thus the desired ω is given by $\omega-1 = \omega \frac{\lambda_1}{\lambda_0} + 1-\omega$ or $\omega = \dfrac{2}{2 - \frac{\lambda_1}{\lambda_0}}$ ($\frac{\lambda_1}{\lambda_0}$ is called the dominance ratio).

The fundamental mode eigenfunction error for large ℓ in the analytic procedure is now proportional to $\left(\frac{\mu_1}{\mu_0}\right)^{\ell} = (\omega-1)^{\ell} = \left(\dfrac{\frac{\lambda_1}{\lambda_0}}{2-\frac{\lambda_1}{\lambda_0}}\right)^{\ell}$ instead of $\left(\frac{\lambda_1}{\lambda_0}\right)^{\ell}$.

A series of test problems were carried out using the single linear extrapolation factor. The indices i and j refer to different fuel compositions. The values of ω used were set by the dominance ratio as described

above and the ω was applied every iteration. The test problems demonstrated that no gain in convergence rates could be obtained relative to the reference strategy. This failure occurs because the difference required in the extrapolation, $(\overline{\Psi}_i^n - \overline{\Psi}_i^{n-1})$, is dominated by statistical fluctuations; this happens even when only a few gross regions (i.e. compositions) are used and with as many as 2000 histories per iteration. In all tests, the source accuracy conditions were met as described in subsection (a).

In spite of the failure of the use of a single linear extrapolation factor in the reference strategy, the question still remains as to whether a more elaborate extrapolation scheme could provide a gain in Monte Carlo eigenfunction convergence rates. One of the more elaborate extrapolation schemes involves the use of multiple linear extrapolation factors, i.e. the Chebyshev polynomial scheme such as that applied in diffusion theory calculations [8]. In this scheme the above extrapolation equation,

$$\Psi_i^{\ell} = \overline{\Psi}_i^{\ell-1} + \omega \left(\frac{G_{ij}\overline{\Psi}_j^{\ell-1}}{\lambda_0^{\ell}} - \overline{\Psi}_i^{\ell-1} \right)$$

is replaced by,

$$\Psi_i = \overline{\Psi}_i^{\ell-1} + \alpha_{\ell} \left(\frac{G_{ij}\overline{\Psi}_j^{\ell-1}}{\lambda_0^{\ell}} - \overline{\Psi}_i^{\ell-1} \right) + \beta \left(\overline{\Psi}_i^{\ell-1} - \overline{\Psi}_i^{\ell-2} \right).$$

Here α_{ℓ} and β_{ℓ} are functions of the iteration number ℓ and of the dominance ratio $\frac{\lambda_1}{\lambda_0}$.

The expressions for α_{ℓ} and β_{ℓ} are

$$\alpha = \frac{2}{2 - \frac{\lambda_1}{\lambda_0}}, \quad \beta_1 = 0,$$

$$\alpha_{\ell} = \frac{4}{\left(\frac{\lambda_1}{\lambda_0}\right)} \left[\frac{\cosh\left[(\ell-1)\gamma\right]}{\cosh\left[\ell\gamma\right]} \right], \quad \beta_{\ell} = \frac{\cosh\left[(\ell-2)\gamma\right]}{\cosh\left[\ell\gamma\right]}; \quad \ell \geq 2,$$

where $\gamma = \cosh^{-1}\left(\frac{2}{\left(\frac{\lambda_1}{\lambda_0}\right)} - 1 \right)$.

Thus, if repeated applications of Chebyshev polynomials of degree one are made then the scheme is identical to the previous one with a single extrapolation factor. In general applications α_{ℓ} and β_{ℓ} would approach values given by

$\lim_{\ell\to\infty} \alpha_\ell = \dfrac{4}{\left(\dfrac{\lambda_1}{\lambda_0}\right)}$, $\lim_{\ell\to\infty} \beta_\ell = 1$. It is now clear that if the use of the single extrapolation factor fails because of statistical fluctuations in the difference

$$\left(\frac{G_{ij}\bar{\psi}_j^{\ell-1}}{\lambda_0^\ell} - \psi_i^{\ell-1} \right),$$

then the Chebyshev polynomial approach must fail also (it will fail all the more because $\alpha_\ell > \omega$).

In rejecting both of the above extrapolation schemes for Monte Carlo calculations, we have rejected successful schemes for analytic calculations. One still might raise the question as to whether an extrapolation scheme could be devised for Monte Carlo that could overcome the difficulties caused by statistical fluctuations. In attempting to answer this question, it is convenient to use some concepts from information theory. First, however, let us recall that the Chebyshev polynomial scheme originated by asking the following question: Is it possible when forming ψ^ℓ for the ℓth iteration to take as the value of ψ^ℓ a linear combination of ψ^ℓ and all previous ψ's and thereby accelerate convergence? That is, if we write

$\psi^\ell = \sum_{n=0}^{\ell} a_{\ell n}\psi^n$, then can we choose constants $a_{\ell n}$, independent of source guess, such that this ψ^ℓ is closer to the fundamental mode eigenfunction? The answer is yes, and although the $a_{\ell n}$ are not obtained explicitly, the resulting procedure is the three-term Chebyshev polynomial extrapolation described above.

There is an analogous question we may ask when considering a Monte Carlo eigenfunction procedure, and as stated earlier it can be developed using some notation from information theory. Let the eigenfunction estimate from the ℓth iteration be written as, $\psi^\ell = S^\ell + N^\ell$, where S^ℓ represents a wanted signal from the ℓth iteration and N^ℓ is a random noise. Consider now a prediction of the eigenfunction at iteration $\ell + n$, say $\hat{S}^{\ell+n}$, written as a linear combination of ψ^m, $m=1,\ldots,\ell$. That is:

$$\hat{S}^{\ell+n} = \sum_{m=1}^{\ell} h_m \psi^{\ell-m+1},$$

where the constants h_m are to be chosen such that $\hat{S}^{\ell+n}$ is closer in some sense to the fundamental mode eigenfunction. Here the h_m are called a linear filter. The problem of finding an optimum linear filter has been considered by N. Wiener [9]. In particular, he solves for h_m which will minimize

$$E\left\{\left[S^{\ell+n} - \hat{S}^{\ell+n}\right]^2\right\},$$

where E denotes expected value. What is required then, is an extension of the linear filter approach which will have the property that if the noise approaches zero (i.e. number of histories per iteration becomes large), then the h_m's become similar to those that are obtained from the Chebyshev polynomial approach. If the signal is constant (i.e. a fixed source problem), then clearly minimum variance will occur for $\hat{S}^{\ell+n}$ if the h_m's are proportional to the number of histories in $\psi^{\ell-m+1}$.

Initial iterations, with on the order of 500 histories each, will be closer to the latter case than to the former. That is, the initial iterations will be more like constant signal with random noise than like variable signal with zero noise. But the reference strategy already weighs iteration results by the number of histories per iteration. Of course, the noise to signal ratio does diminish with increasing iteration number because of the increasing number of histories per iteration.

If the above extension of linear filter theory were obtained, its use might not result in large gains in Monte Carlo eigenfunction convergence rates; nevertheless this approach has some promise [10], but is outside the scope of this study.

d. Source Strategies Using Importance Sampling

Importance sampling offers a general means of reducing the Monte Carlo uncertainty in a specified neutron reaction rate. For example, if an approximate solution of the adjoint transport equation is known, where the adjoint source is the space-energy dependent cross section for the desired reaction rate, then one may alter the frequency distribution functions in the normal Monte Carlo process (and introduce weights to retain a fair game) in such a way that the standard deviation for the desired reaction rate is decreased per history. As the approximate adjoint solution approaches the true adjoint solution, this standard deviation per history approaches zero; if the approximate solution becomes too crude, the standard deviation can become larger than that from a normal Monte Carlo calculation.

In addition to the difficulty of obtaining a sufficiently accurate approximate solution of the adjoint transport equation, the importance sampling approach described above has another serious drawback. Namely, one may choose only one neutron reaction rate per calculation to be obtained with small standard deviation. In general, importance sampling may be viewed as a process where the reduction of variance for a reaction of interest is obtained at the cost of increasing the variance for other reactions. Thus, such a process would not be generally acceptable when attempting to obtain the eigenfunction shape by Monte Carlo.

However, one may still raise the question as to whether there are importance sampling techniques that could significantly improve the definition of the overall eigenfunction shape by reducing the ratio of the source shape uncertainty to the statistical uncertainty and by reducing the total

uncertainty. A definitive answer to this question is outside the scope of this study, but the answer is conjectured to be no. For example, consider the following limited form of importance sampling, a form that may be called "source" importance sampling. Start with the observation that neutrons from the high energy part of the fission spectrum are more important in determining overall eigenfunction shape than those from the low energy part. Thus, source importance sampling that would start more neutrons at high energy (with adjusted weights to retain the proper source spectrum) could result in some reduction of source shape uncertainty, but would increase statistical uncertainty because of the non-constant starting weights. Intuitively, we feel that such a process could not result in a significant reduction of total eigenfunction uncertainty.

More generally, several specialized forms of importance sampling are frequently available for fixed source problems that could be used in eigenfunction calculations. These are available in splitting and Russian Roulette. Such forms of importance sampling have not been used in test calculations because again we intuitively believed that significant reductions in total eigenfunction uncertainty would not occur. An exception, which is used routinely in eigenfunction calculation, is Russian Roulette in the thermal group in large water reflectors.

3. A Reference Strategy

In the preceding subsections, we have discussed several possible source strategies including a reference strategy, ie. the straight analog of the analytic power method with an increasing number of histories per iteration. The strategies discussed include: 1) those using a Green's function approach, 2) those using extrapolation factors, and 3) those using importance sampling. The test calculations performed in this study did not reveal any strategy superior to a reference strategy, which will be described. Also from this study, the approach that has been deemed most likely to offer some improvement in Monte Carlo eigenfunction convergence rates relative to the reference strategy is the extension of linear filter theory discussed in subsection (c); but, as stated earlier, this is beyond the scope of this work.

Some important operational details of the reference strategy are as follows: Because all neutrons in a given iteration are processed through an energy group before proceeding to the next energy group, then the calculational time per history becomes shorter if the problem has fewer iterations with more histories per iteration. On the other hand, it is desirable in an eigenfunction calculation to have as many iterations as possible. Thus, reducing the number of histories as far as possible without losing too much machine efficiency in initial iterations led the authors to adopt 500 as the minimum number of histories per iteration. Because of correlated sampling, it is desirable that all eigenfunction problems have the same number of histories per iteration.

The reference strategy uses an increase of 10 histories per iteration in successive iterations. Thus in the Nth iteration, the number of histories per iteration would be $500 + 10\,N$ and the total number of histories through the Nth iteration would be $500N + 5N(N+1)$.

The rate of increase of 10 histories per iteration was arrived at as follows: From the work of Lieberoth [11], it is known that the eigenfunction error that results from an infinite number of power iterations, each with a constant number of histories per iteration, decreases as $\frac{1}{N}$. Thus, the desire to keep the total number of iterations as large as possible for a fixed total number of histories, and at the same time to insure a reasonable removal rate for the eigenfunction error that would result from a fixed number of histories, led the authors to adopt the rate increase of 10 histories per iteration. From the discussion of this error removal in Ref. [1], it is evident that this error will tend to cause the approach to the eigenfunction to be on the side in which there is some power shift toward regions of low k_∞, and that the eigenvalue will tend to converge from below rather than above. However, calculations for realistic types of reactors have not shown such a discernible trend. This suggests, for realistic calculations, that in general such a component of the source shape uncertainty is small relative to the statistical uncertainty.

There is one final important operational detail connected with the eigenfunction strategy adopted; the procedure to be used to obtain an estimate of the total eigenfunction uncertainty. The remainder of this report will deal with arriving at such a procedure.

III. EIGENFUNCTION UNCERTAINY

This section presents the arguments that led to a formula for total eigenfunction uncertainty for the strategy adopted. As in previous sections, it will again be convenient to consider the overall strategy as composed of the same two parts. That is, the first part is the totality of the Monte Carlo procedures that would be involved in a fixed source calculation; the uncertainties in this part are called statistical. The second part is the eigenfunction source strategy that connects successive generations; the uncertainties in this part are called source shape uncertainties. It is obvious that in any eigenfunction calculation, the uncertainties from these two parts are connected. However, this breakdown is convenient here simply because uncertainties that arise from fixed source calculations are very easily obtained and are an integral part of all fixed source Monte Carlo calculations. Some insights as to the nature of the source shape uncertainties can be obtained by considering some properties of the statistical uncertainties. This is done in subsection (a). Subsection (b) formalizes these observations through use of serial correlation coefficients and arrives at a formula for total eigenfunction uncertainty by a modification of an approach proposed by D. B. MacMillan. Because the overall formula for eigenfunction uncertainty is partly empirical, subsection (c) then gives some results of its verification by numerical testing.

a. Some Properties of the Statistical Uncertainties

There are two commonly used methods in Monte Carlo to provide estimates of statistical uncertainties. The first is a variance calculation based on a history by history basis; i.e. each experiment is taken to be one history. This estimate is an accurate estimate of the purely statistical uncertainty, since it yields the identical result that one would obtain from a

hypothetical fixed source problem where the total sampling from the fixed source yields the total set of source sites from all iterations in the eigenfunction calculation.

The second estimate of statistical uncertainty that is commonly used appears superficially to be a correct combination of source shape uncertainty and statistical uncertainty. This estimate is that obtained from a variance calculation on an iteration by iteration basis. This estimate appears to include the source shape uncertainty, because within the power method the iteration source change is by iteration. However, test calculations which compare these two estimates of uncertainty have shown that their magnitudes are reasonably close to each other. Since the first estimate is without question the purely statistical uncertainty, then we must consider the second estimate the same uncertainty as the first.

The above observations suggest the following argument. Since the uncertainty calculation on an iteration by iteration basis does sense the changing source shape, this uncertainty could be of the same magnitude as the statistical uncertainty only if there exists a correlation among iteration source shapes such that there is almost a cancellation of the source shape component of uncertainty. But the existence of such a correlation suggests that asymptotically (i.e. after a large number of iterations) the source shape uncertainty is some constant times the statistical uncertainty. Stated differently, after a large number of iterations we can visualize that the sampling of possible source shapes lags behind the sampling of possible neutron production shapes based on these source shapes, and that the ratio of sampling rates tends to some constant.

Of course, the above argument, which suggests that the ratio of source shape and statistical uncertainties tends to a constant for a large number of iterations, is of little value unless the magnitude of this constant can be estimated. This ratio strongly depends on the type of quantity being edited in the Monte Carlo calculation; for example, let us consider the region dependence of the ratio for neutron production rates. If the region size starts to approach the whole reactor, then because the neutron production shape error will most likely consist of the lower overtones which will roughly integrate to zero, the ratio of source shape and statistical uncertainty decreases. Stated differently, there are negative correlations among neutron production rates over some large subregions of a reactor. At the other extreme of a region size becoming small, it is obvious that again the ratio of source shape and statistical uncertainties decreases. Thus, one concludes that the ratio becomes a maximum for some intermediate sized region.

An early approach taken by the authors for obtaining the asymptotic source shape uncertainty was to obtain a maximum ratio of source shape and statistical uncertainties from test calculations and use this constant universally [12]. Thus, in practice we computed the usual statistical uncertainty, and applied a constant multiplicative factor to obtain the source shape uncertainty. When continuing calculations revealed that this factor was greater than 10 for some edit quantities, this approach was dropped since it would force unreasonable conservatism in the uncertainties for most edit quantities.

b. Eigenfunction Uncertainty Using Serial Correlation

The conjecture of the previous subsection, that in the asymptotic range the source shape uncertainty is some constant times the statistical uncertainty, may be formalized as follows: Let x_n be an estimate of some neutron reaction rate from iteration n in an eigenfunction calculation. Accumulating iteration results for N iterations yields the Monte Carlo estimate of this reaction rate, say R_N, where

$$R_N = \sum_{n=1}^{N} a_n x_n,$$

where $a_n = \dfrac{N_n}{\sum_{n=1}^{N} N_n}$, and N_n is the number of neutron histories in iternation n.

As discussed in Section II, the eigenfunction strategy insures that $\underset{N\to\infty}{\text{PLim}}\, R_N$ is the transport theory value of the reaction rate. Following a well-known derivation using expected value notation, we may write

$$\sigma^2(R_N) = E\,[R_N - E(R_N)]^2.$$

Let $\mu_n = E(x_n)$, then we have

$$\sigma^2(R_N) = E\,[a_1(x_1-\mu_1) + \dots a_N(x_N-\mu_N)]^2$$

$$= \sum_{n=1}^{N} (a_n^2\, E(x_n-\mu_n)^2 + 2\sum_{n<m} a_n a_m\, E[(x_n-\mu_n)(x_m-\mu_m)].$$

But $\sigma^2(x_n) = E(x_n-\mu_n)^2$, and the correlation coefficient between x_n and x_m, say ρ_{nm}, is defined by,

$$\rho_{nm}\, \sigma(x_n)\, \sigma\,(x_m) = E[(x_n-\mu_n)(x_m-\mu_m)],$$

Thus we may write,

$$\sigma\,(R_N) = \sum_{n=1}^{N} \left(a_n^2\, \sigma^2(x_n) + 2 \sum_{n<m} a_n a_m\, \rho_{nm}\sigma(x_n)\sigma(x_m) \right).$$

Let $\sigma_s^2(R_N)$ be the purely statistical component of variance for R_N, after N iterations; then in the asymptotic range we must have

$$\sigma^2(x_n) = \frac{1}{a_n}\,\sigma_s^2(R_N).$$

We may write

$$\sigma^2(R_N) = \sigma_s^2(R_N)\left[1 + 2\sum_{n=1}^{N}\sum_{n<m}\sqrt{a_n}\,\sqrt{a_m}\,\rho_{nm}\right].$$

Now if the source shape uncertainty is some constant times the statistical uncertainty, then this only implies that

$$\underset{N\to\infty}{\mathrm{PLim}}\left[\sum_{n=1}^{N}\sum_{n<m}\sqrt{a_n}\,\sqrt{a_m}\,\rho_{nm}\right]$$ exists. However, if the number of histories per iteration were a constant, that is, $a_n = a_m = \frac{1}{N}$, then it is also obvious that in the asymptotic range the correlation coefficient, ρ_{nm}, is a function only of the difference in n and m. The expectation that this should be true even when the number of histories per iteration increases with iteration number has been verified by test calculations. Thus, if we set k = n-m, then we may write

$$\rho_{nm} = \rho_k,$$

where ρ_k is called the serial correlation coefficient of lag k.

The expression for $\sigma^2(R_N)$ now becomes

$$\sigma^2(R_N) = \sigma_s^2(R_N)\left[1 + 2\sum_{k=1}^{N-1}\left(\sum_{n=1}^{N-k}\sqrt{a_n}\,\sqrt{a_{n+k}}\,\rho_k\right)\right].$$

Let us define

$$A = \sum_{k=1}^{N-1}\left(\sum_{n=1}^{N-k}\sqrt{a_n}\,\sqrt{a_{n+k}}\,\rho_k\right).$$

We may note in passing, that if we assume that the purely statistical uncertainty and source shape uncertainty are independent, then $\sqrt{2A}\,\sigma_s(R_N)$ is the magnitude of the source shape uncertainty.

Attempts by the authors to use the above expression for $\sigma^2(R_N)$ in test eigenfunction calculations did not prove to be workable. For example, for some edit quantities the series that defines A was slowly converging, and the ρ_k involved appeared erratic for larger values of k even after a few hundred iterations. Also, there are practical difficulties in this approach because of the machine storage requirements in obtaining a large number of serial

correlation coefficients for every edit quantity.

A workable alternative to this approach was proposed by D. B. MacMillan [13]. He proposed that only the serial correlation coefficient of lag 1 be calculated explicitly and that those of higher labs be inferred. He argues on a physical basis that an upper bound on the value of ρ_k may be obtained by writing

$$\rho_k = \rho_1 \left(\frac{\lambda_1}{\lambda_0}\right)^{k-1} ,$$

where $\frac{\lambda_1}{\lambda_0}$ is the dominance ratio.

Test calculations by MacMillan as well as subsequent tests by the authors have shown that this argument is sound. Thus with MacMillan's approach we have upper estimates for all ρ_k once we have estimated ρ_1 and supplied a value of $\frac{\lambda_1}{\lambda_0}$. We will now consider $\lim_{N\to\infty} A$ with values of ρ_k obtained from MacMillan's conjecture. Because of the fixed increase in number of histories per iteration, we may write

$$\lim_{N\to\infty} \sum_{n=1}^{N-k} \sqrt{a_n} \quad \sqrt{a_{n+k}} = 1$$

for any value of k. Thus we have

$$\lim_{N\to\infty} A = \sum_{k=1}^{\infty} \rho_k = \rho_1 \sum_{k=1}^{\infty} \left(\frac{\lambda_1}{\lambda_0}\right)^{k-1} = \frac{\rho_1}{1 - \left(\frac{\lambda_1}{\lambda_0}\right)}$$

MacMillan's formula for total eigenfunction uncertainty in reaction rate R_N becomes

$$\sigma(R_N) = \sigma_s(R_N) \sqrt{1 + \frac{2\rho_1}{1 - \left(\frac{\lambda_1}{\lambda_0}\right)}} .$$

A modification of this approach for estimating $\sigma(R_N)$ was developed by the authors for two reasons. First, the above approach requires supplying the dominance ratio, which is not always easy to obtain. Second, test calculations reveal that the estimated $\sigma(R_N)$ is conservative if ρ_1 is small in a core with large $\frac{\lambda_1}{\lambda_0}$. Recall from the discussion in subsection (a) that we can expect, for a core with a given dominance ratio, there will be reaction rates for which the ratio of source shape and statistical uncertainties will be a maximum and hence ρ_1 will be a maximum. Further, this maximum value of

ρ_1 will increase as $\frac{\lambda_1}{\lambda_0}$ increases. The authors concluded by test calculations that MacMillan's formula is not overly conservative in its estimate of $\sigma(R_N)$ for reaction rates for which ρ_1 is a maximum. These same test calculations have revealed that for reaction rates for which ρ_1 is much less than its maximum value, the values of ρ_k decline with increasing k more like ρ_1^{k-1} than like $\left(\frac{\lambda_1}{\lambda_0}\right)^{k-1}$. That is,

$$\rho_k = \rho_1(\rho_1)^{k-1} \text{ (for small } \rho_1).$$

This would lead to

$$\lim_{N\to\infty} A \approx \frac{\rho_1}{1-\rho_1} \text{ (for small } \rho_1).$$

But since the maximum value of ρ_1 depends on $\frac{\lambda_1}{\lambda_0}$, we may hopefully scale this expression such that we obtain a reasonable value of $\lim_{N\to\infty} A$ for the maximum value of ρ_1. This scale factor was empirically determined to be about 5. Thus we have

$$\lim_{N\to\infty} A \approx \frac{5\rho_1}{1-\rho_1} \text{ (for maximum } \rho_1).$$

We may now linearly connect the values of $\lim_{N\to\infty} A$ for maximum ρ_1 and for small ρ_1 at the values $\rho_1 = 1$ and $\rho_1 = 0.2$ by writing

$$\lim A \approx \frac{5\rho_1^2}{1-\rho_1}.$$

This gives an alternate formula for $\sigma(R_N)$, i.e.

$$\sigma(R_N) = \sigma_s(R_N)\sqrt{1 + \frac{10\rho_1^2}{1-\rho_1}}.$$

As a result of test calculations which will be described in the next subsection, this formula for $\sigma(R_N)$ was adopted.

c. Numerical Testing of Formula for Total Eigenfunction Uncertainty

The formula for total eigenfunction uncertainty in a reaction rate R_N is given in the previous subsection as

$$\sigma(R_N) = \sigma_s(R_N)\sqrt{1 + \frac{10\rho_1^2}{1-\rho_1}},$$

where $\sigma_s(R_N)$ is the purely statistical standard deviation, and ρ_1 is the serial correlation coefficient of lag 1. Some details of the calculational steps are as follows: If the estimate of the reaction rate R_N is given by

$$R_N = \sum_{n=1}^{N} a_n x_n,$$

then the purely statistical variance $\sigma_s^2(R_N)$ is given by

$$\sigma_s^2(R_N) = \frac{1}{N} \sum_{n=1}^{N} a_n (x_n - R_N)^2.$$

Since for convenience we desire to use the above R_N and $\sigma_s^2(R_N)$ in obtaining ρ_1, we may introduce the number of histories per iteration, N_n, as weights in the definition of ρ_1 as follows:

$$\rho_1 = \frac{1}{N\sigma_s^2(R_N)} \sum_{n=2}^{N} \sqrt{a'_{n-1}} \sqrt{a''_n} (x_{n-1} - R'_N)(x_n - R''_N)$$

where $a'_{n-1} = \dfrac{N_{n-1}}{\sum_{n=2}^{N} N_{n-1}}$, $a''_n = \dfrac{N_n}{\sum_{n=2}^{N} N_n}$

$$R'_N = \sum_{n=2}^{N} a'_{n-1} x_{n-1},$$

$$R''_N = \sum_{n=2}^{N} a''_n x_n .$$

By making use of the approximations

$$R'_N \cong \sum_{n=2}^{N} \sqrt{a'_{n-1}} \sqrt{a''_n} x_{n-1} \left(\frac{1}{\sum_{n=2}^{N} \sqrt{a'_{n-1}} \sqrt{a''_n}} \right),$$

$$\text{and } R_N'' \cong \sum_{n=2}^{N} \sqrt{a'_{n-1}}\ \sqrt{a''_n}\ x_n \left(\frac{1}{\sum_{n=2}^{N} \sqrt{a'_{n-1}}\ \sqrt{a''_n}} \right),$$

we obtain the expression used, that is

$$\rho_1 = \frac{1}{N\sigma_s^2(R_N)} \left[\left(\sum_{n=2}^{N} \sqrt{a'_{n-1}}\ \sqrt{a''_n}\ x_{n-1}x_n \right) - \left(R_N' R_N'' \sum_{n=2}^{N} \sqrt{a'_{n-1}}\ \sqrt{a''_n} \right) \right].$$

The eigenfunction test calculations may be described as follows: The cross sections and fixed source procedures are such that a zero variance results in an infinite medium calculation. In particular, the cross sections are energy independent and purely absorbing. The Monte Carlo eigenfunction calculations were performed for a series of eight slab cores with dominance ratios ranging from 0.36 to 0.995. Although the dominance ratio, $\frac{\lambda_1}{\lambda_0}$, is not required in the uncertainty formula adopted, it is estimated for each of the slab cores in the following manner: The neutron flux distribution, say $\varphi(x)$, due to a plane unit source at x=0 in an infinite purely absorbing medium with constant cross section is given by $\varphi(x) = \frac{1}{2} E_1(\Sigma|x|)$. The neutron flux age, J, is then $\frac{1}{3\Sigma^2}$.

Using well known results from two-group diffusion theory, we may write the nth mode eigenvalue as proportional to $\frac{1}{1+JB_n^2}$, where B_n^2 is the geometric buckling for the nth mode eigenfunctions. For a slab reactor, $B_0^2 = (\frac{\pi}{H})^2$ and $B_1^2 = (\frac{2\pi}{H})$, where H is the effective slab width, given as the actual slab width plus twice the extrapolation distance. The extrapolation distance may be taken as $\frac{2}{3\Sigma}$. Thus the dominance ratio may be estimated as

$$\frac{\lambda_1}{\lambda_0} = \frac{1 + \frac{1}{3}(\frac{\pi}{H\Sigma})^2}{1 + \frac{1}{3}(\frac{2\pi}{H\Sigma})^2}.$$

For each slab core calculated, 17 different region edits of the neutron production rate were obtained and compared to accurate production rates. The 17 regions are: 2 symmetrically placed slab regions of volume fraction 0.01 each, 10 uniform slabs of volume fraction 0.1 each, 2 symmetrically placed slab regions of volume fraction 0.3 each, 2 symmetrically placed slab regions of volume fraction 0.3 each, 2 symmetrically placed slab regions of volume

fraction 0.5 each, and the whole core. The accurate production rates were obtained by averaging results in symmetrically placed regions after 400 iterations and a total of 10^6 histories. In all calculations, the source guess accuracy requirement was met as discussed in subsection (b) of Part II.

The eigenfunction uncertainties for the five different sizes of regions, in each of the eight cores, are tabulated at multiples of 50 iterations up to 250 iterations in Table I. The eigenfunction uncertainties are tabulated using naively the purely statistical standard deviation as well as using the standard deviation obtained with the formula. The method of uncertainty tabulation is by use of a figure of merit, f_m. For example, the figure of merit for the regions of size 0.1 is the number of the ten regions with a neutron production rate error greater than one probable error, plus the number of the ten regions with an error greater than two probable errors, etc. to an infinite sum. Thus, in this case the expected figure of merit, $\overline{f}_m$, is 7.58. A value much greater than this indicates that the uncertainty estimates are too small, and a smaller value indicates a conservative estimate. Shown in Table I is the expected figure of merit plus 2σ. An acceptable means of estimating eigenfunction uncertainties is defined to be one with a figure of merit less than this.

Table I shows that for a very small core, $\frac{\lambda_1}{\lambda_0} = .36$, an acceptable estimate of uncertainties is provided by the purely statistical standard deviation, and that the formula gives only a slightly more conservative result. As the cores become larger, we see that the purely statistical estimate of uncertainty becomes, in general, unacceptable, whereas the formula gives acceptable results. We may note that in the special case of the neutron production over the whole core, i.e. the eigenvalue, the purely statistical standard deviation for all cores is acceptable and that the formula gives essentially the same standard deviation. For the case of the largest core, $\frac{\lambda_1}{\lambda_0} = .995$, the unacceptable results from the formula when the number of iterations is 50 or 100 is because the magnitude of the serial correlation coefficient, ρ_1, is not reasonably estimated for a large core unless more than 100 iterations are performed. Although more than 100 iterations appears to be necessary to obtain acceptable eigenfunction shape uncertainties, this is not required if only an acceptable eigenvalue uncertainty is desired.

One final observation will be made concerning the approach of an estimated Monte Carlo eigenfunction shape to the true eigenfunction shape. When viewing an accumulation by iteration of a neutron reaction rate, one almost naturally looks for an oscillation about the true reaction rate as the number of iterations increases. Of course this tendency also exists when viewing a fixed source calculation. The properties of accumulated results by iteration in a fixed source calculation are those of the classical random walk problem. W. Feller [14] points out that in such a problem, the most probable number of times that the estimated reaction rate crosses the true reaction rate is zero, and that zero times is more probable than one time, etc. Further, the mean number of iterations between crossings approaches infinity as the number of iterations increases.

These same properties exist in a somewhat more severe form in a Monte Carlo eigenfunction calculation even though the number of histories increases

with increasing iteration number. This is because of the positive correlation between successive iteration reaction rates. Thus, if the neutron source shape at one iteration tends to make a single iteration reaction rate too high, then more likely than not the source shape of the next iteration will tend to make the next single iteration reaction rate too high. In contrast, for a fixed source calculation, each single iteration result, assuming a symmetrical frequency distribution function, is equally likely to be above or below the true reaction rate. By analogy to the classical random walk problem, we conclude for the eigenfunction calculation that the mean number of iterations between crossings of the estimated and true reaction rates approaches infinity as the number of iterations increases.

REFERENCES

1. R. C. GAST, "Monte Carlo Eigenfunction Iteration Strategies that are and are not Fair Games," WAPD-TM-878, (1969).

2. G. W. MORRISON, J. T. MIHALCZO, and D. C. IRVING, "REACT and CONVRG-Fortran Subroutines for Determining Source Convergence for the 05R Monte Carlo Neutron Transport Code," ORNL-TM-1325, (1966).

3. J. T. MIHALCZO, "Multiplication Factor of Uranium Metal by One Velocity Monte Carlo Calculations," Nuc. Sci. and Eng., Vol. 27, pg. 557, (1967).

4. M. R. MENDELSON, "Monte Carlo Criticality Calculations for Thermal Reactors," Nuc. Sci. and Eng., Vol. 32, pg. 319, (1968).

5. J. B. PARKER and E. R. WOODCOCK, "Monte Carlo Criticality Calculations," Progress in Nuclear Energy, Series IV, Vol. 4, pg 441, Pergamon Press, (1961).

6. E. L. KAPLAN, "Monte Carlo Methods for Equilibrium Solutions in Neutron Multiplication," UCRL-5275-T, (1958).

7. L. L. CARTER and N. J. McCORMICK, "Source Convergence in Monte Carlo Calculations," Nuc. Sci. and Eng., Vol. 36, pg 438, (1969).

8. L. A. HAGEMAN, "Numerical Methods and Techniques Used in the Two-Dimensional Neutron-Diffusion Program PDQ-5," WAPD-TM-364, pg 25, (1963).

9. W. B. DAVENPORT, Jr. and W. L. ROOT, An Introduction to the Theory of Random Signals and Noise, McGraw-Hill Book Co., Inc., pg 222, 1958.

10. Subsequent to the writing of this report, a 1972 PhD thesis at Carnegie Mellon University by R. E. Swaja, "Application of Stochastic Optimal Estimation Principles to Monte Carlo Eigenfunction Convergence," has shown that a linear filter theory approach he develops results in gains in eigenfunction convergence for beginning iterations, but asymptotically appears to approach the convergence rate of the power method.

11. J. LIEBEROTH, "A Monte Carlo Technique to Solve the Static Eigenvalue Problem of the Boltzmann Transport Equation," Nukleonik, Vol. 11, pg. 213, (1968).

12. R. C. GAST and N. R. CANDELORE, "Monte Carlo Eigenfunction Uncertainties," Trans. Am. Nucl. Soc., Vol. 14, No. 1, pg 219, (1971).

13. D. B. MACMILLAN, "Monte Carlo Confidence Limits for Iterated-Source Calculations," Nuc. Sci. and Eng., Vol. 50, pg. 73, (1973).

14. W. FELLER, An Introduction to Probability Theory and Its Applications, John Wiley and Sons, Inc., Chapter III, (1957).

APPENDIX

As discussed in Section II, the strategy adopted is an analog of the analytic power method, with an increasing number of histories per iteration. A question that naturally arises is whether neutron weights could be introduced in an analog of the power method with a fixed number of histories per iteration so as to make it a fair game. This appendix gives a procedure for doing this and a proof that the result is a fair game.

We know from Section II.a that the equation describing the analog of the power method with an increasing number of histories per iteration, i.e.

$$\psi_i^{n+1} = \sum_j G_{ij}^n K^n (\psi_j^n)$$

results in the analytic transport equation,

$$\psi_i = \sum_j G_{ij} \bar{\psi}_j ,$$

upon taking the P Lim as $n \to \infty$.

Suppose now we have a system that is supercritical and follow all progeny such that the number of histories increases with iteration number (continuation of the Markov process must be assured as discussed in Section II.a. This process may be represented in a power method analog with a fixed number of histories per iteration by the use of weights as follows: Let k_n be the multiplication factor for the nth iteration, i.e. the total neutron production divided by the total starting weight in the nth iteration. Then the supercritical process could be represented by having the weight for iteration n, say W_n determined by $W_o = 1$, $W_n = \prod_{i=0}^{n-1} k_i$. Suppose for the moment there are no statistical fluctuations in the multiplication factor k_n, then asymptotically (n large) k_n becomes a constant, say $\bar{k}$. In this case the weight W_n will have an exponential increase, i.e., we may write $W_{n+1} = \bar{k}^n = e^{\alpha n}$ where $\alpha = \ln \bar{k}$. With such a weighting, the effective number of histories contributing to any estimated reaction rate remains finite as n increases and hence the process cannot be a fair game. This is easily seen by observing that the effective number of histories contributing to a reaction rate is proportional to the mean value of n for a fixed number of histories per iteration; that is the mean value of n, say $\bar{n}$, over previous iterations up to iteration N. But,

$$\overline{n} = \int_0^N e^{\alpha(N-n)} (N-n)dn = \frac{1}{\alpha^2} [1 - e^{-\alpha N}(\alpha N+1)],$$

and

$$\lim_{N\to\infty} \overline{n} = \frac{1}{\alpha^2} ;$$

hence the effective number of histories contributing to a reaction rate is finite as $N\to\infty$ (of course the same occurs in a subcritical reactor).

If we did not have statistical fluctuations in the asymptotic multiplication factor $\overline{k}$, then a scaling of ν (neutrons per fission) such that $\overline{k} = 1$ and hence $\alpha = 0$ will produce the basic requirement that the effective number of histories contributing to a reaction rate approaches ∞ as $N\to\infty$. In this case the weight W_n would be,

$$W_n = \frac{\prod_{i=0}^{n-1} k_i}{(\overline{k})^{n-1}} .$$

Because of the statistical fluctuations in k_i, the estimated multiplication factor for iteration i, caused by a finite number of histories in the Monte Carlo calculation, a reactor that is just critical as defined by the analytic transport equation is subcritical in reality and the neutron density will approach zero as the iterations approach ∞ (this also occurs in a real reactor). This may be easily seen as follows: A reactor that is just critical mathematically (i.e. continuous neutron density in the linear transport equation) will meet the condition:

$$\underset{N\to\infty}{\mathrm{PLim}} \frac{1}{N} \sum_{i=1}^{N} k_i = 1.$$

We may write $k_i = 1 - \Delta_i$ where Δ_i are the statistical fluctuations and assume without loss of generality

$$\sum_{i=1}^{N} \Delta_i = 0.$$

The amplitude of the spatially integrated neutron density after N iterations is then proportional to

$$\prod_{i=1}^{N} k_i .$$

We may write

$$\prod_{i=1}^{N} (1-\Delta_i) = 1 - \sum_{i=1}^{N} \Delta_i + \Delta_1\left(\sum_{i=2}^{N} \Delta_i\right) + \Delta_2\left(\sum_{i=3}^{N} \Delta_i\right) + \ldots + O(\Delta^3).$$

Making use of

$$(\Delta_1+\Delta_2+\ldots+\Delta_N) = \Delta_1^2 + \Delta_2^2 + \ldots + \Delta_N^2 + 2\Delta_1\left(\sum_{i=2}^{N} \Delta_i\right) + 2\Delta_2\left(\sum_{i=3}^{N} \Delta_i\right) + \ldots,$$

we may write

$$\prod_{i=1}^{N} (1-\Delta_i) = 1 - \frac{1}{2}\left[\sum_{i=1}^{N} \Delta_i^2\right] + O(\Delta^3) < 1.$$

Hence it is clear that the expected neutron density amplitude will approach zero as $N \to \infty$. The degree of supercriticality required to keep the expected neutron density amplitude at a constant for a large number of generations, N, in a reactor with finite neutron density, is obtained by scaling ν by

$$\frac{1}{\left(\prod_{i=1}^{N} k_i\right)^{\frac{1}{N}}}.$$

Thus the desired neutron weight at iteration n in a problem with a total of N iterations is given by

$$W_n(N) = \frac{\prod_{i=1}^{n-1} k_i}{K^{n-1}}, \text{ where } K = \left(\prod_{i=1}^{N} k_i\right)^{\frac{1}{N}}.$$

(This weighting was first suggested by D. B. MacMillan, private communication).

The proof that the use of the above weight, $W_n(N)$, in the Monte Carlo process produces a fair game may be summarized as follows: Using the definitions of Section II.a, we may write the Monte Carlo process as

$$\psi_i^{n+1} = \sum_j G_{ij}^n W_n(N) K^n(\psi_j^n).$$

Summing both sides of this equation over n, we have

$$\frac{1}{N}\sum_{n=1}^{N}\psi_i^{n-1} = \sum_j \left[\left(\frac{\sum_{n=1}^{N} G_{ij}^{n}\, W_n(N) K^n(\psi_j^n)}{\sum_{n=1}^{N} W_n(N) K^n(\psi_j^n)} \right) \frac{1}{N}\sum_{n=1}^{N} W_n(N) K^n(\psi_j^n) \right].$$

The quantity in the parenthesis above is simply a weighted average of the estimated neutron production in elementary volume i due to a neutron born in elementary volume j; hence, its PLim as $N \to \infty$ is the analytic neutron production Green's function G_{ij}. All that is additionally required to show that the above equation reduces to the analytic transport equation,

$$\psi_i = \frac{1}{\lambda}\sum_j G_{ij}\psi_j, \text{ where } \psi_i = \underset{N\to\infty}{\mathrm{PLim}}\ \frac{1}{N}\sum_{n=1}^{N}\psi_i^{n+1}$$

is to show that

$$\underset{N\to\infty}{\mathrm{PLim}}\ \frac{1}{N}\sum_{n=1}^{N} W_n(N) K^n(\psi_j^n) = \underset{N\to\infty}{\mathrm{PLim}}\ \frac{1}{\lambda}\frac{1}{N}\sum_{n=1}^{N}\psi_j^{n+1},$$

where

$$\lambda = \frac{\underset{N\to\infty}{\mathrm{PLim}}\ \frac{1}{N}\sum_{n=1}^{N}\sum_i \psi_i^{n+1}}{\underset{N\to\infty}{\mathrm{PLim}}\ \frac{1}{N}\sum_{n=1}^{N} W_n(N)}.$$

The desired result is immediately obtained by recognizing that the product, $W_n(N)$ times the sampling operator K^n, defines simply a Russian Roulette operator.

The weights $W_n(N)$ and a fixed number of histories per iteration were not used, since the weight for the nth iteration continually changes as the total number of iterations increases, thereby destroying the ability to use accumulative counters for internal editing of reaction rates (this difficulty would not exist for post editing).

TABLE I

Slab Core Eigenfunction Uncertainties

Values of f_m for $\frac{\lambda_1}{\lambda_0} = .36$

Region Data		No. of Iterations											
		50		100		150		200		250			
No.	Fract. Size	Stat.	For	Stat.	For	Stat.	For	Stat.	For	Stat.	For	$\overline{f}_m$	$\overline{f}_m+2\sigma$
2	.01	1	1	1	1	2	2	3	3	1	1	1.51	4.44
10	.1	10	9	9	8	12	12	5	5	10	10	7.58	14.12
2	.3	1	1	2	2	3	3	0	0	1	1	1.51	4.44
2	.5	0	0	2	2	3	3	1	1	1	1	1.51	4.44
1	1.0	0	0	1	1	1	1	1	1	1	1	.76	2.82

Values of f_m for $\frac{\lambda_1}{\lambda_0} = .68$

Region Data		No. of Iterations											
		50		100		150		200		250			
No.	Fract. Size	Stat.	For	Stat.	For	Stat.	For	Stat.	For	Stat.	For	$\overline{f}_m$	$\overline{f}_m+2\sigma$
2	.01	1	1	0	0	2	2	2	2	3	3	1.51	4.44
10	.1	10	4	10	6	16	9	15	9	13	7	7.58	14.12
2	.3	2	0	2	0	4	1	5	2	5	2	1.51	4.44
2	.5	1	0	0	0	1	0	1	0	5	2	1.51	4.44
1	1.0	1	1	1	1	0	0	0	0	0	0	.76	2.82

Values of f_m for $\frac{\lambda_1}{\lambda_0} = .84$

Region Data		No. of Iterations											
		50		100		150		200		250			
No.	Fract. Size	Stat.	For	Stat.	For	Stat.	For	Stat.	For	Stat.	For	$\overline{f}_m$	$\overline{f}_m+2\sigma$
2	.01	2	0	0	0	1	1	0	0	1	1	1.51	4.44
10	.1	13	6	19	7	17	6	13	6	17	8	7.58	14.12
2	.3	3	0	4	1	6	1	3	1	6	1	1.51	4.44
2	.5	0	0	6	0	1	0	0	0	1	0	1.51	4.44
1	1.0	0	0	1	1	2	2	1	1	1	1	.76	2.82

TABLE I - (Cont'd)

Values of f_m for $\frac{\lambda_1}{\lambda_0} = .92$

Region Data		No. of Iterations											
		50		100		150		200		250			
No.	Fract. Size	Stat.	For	Stat.	For	Stat.	For	Stat.	For	Stat.	For	$\bar{f}_m$	$\bar{f}_m+2\sigma$
2	.01	4	3	7	6	4	3	4	4	5	4	1.51	4.44
10	.1	16	6	28	13	27	11	38	14	33	10	7.58	14.12
2	.3	12	3	19	6	21	6	25	7	10	1	1.51	4.44
2	.5	9	2	16	4	11	2	23	4	8	0	1.51	4.44
1	1.0	0	0	1	0	1	1	1	1	1	1	.76	2.82

Values of f_m for $\frac{\lambda_1}{\lambda_0} = .96$

Region Data		No. of Iterations											
		50		100		150		200		250			
No.	Fract. Size	Stat.	For	Stat.	For	Stat.	For	Stat.	For	Stat.	For	$\bar{f}_m$	$\bar{f}_m+2\sigma$
2	.01	2	2	2	2	1	1	3	3	4	3	1.51	4.44
10	.1	36	6	43	9	41	8	41	9	36	9	7.58	14.12
2	.3	12	1	13	1	4	0	1	0	5	0	1.51	4.44
2	.5	18	0	28	2	18	1	19	2	26	2	1.51	4.44
1	1.0	1	1	1	1	0	0	0	0	1	1	.76	2.82

Values of f_m for $\frac{\lambda_1}{\lambda_0} = .98$

Region Data		No. of Iterations											
		50		100		150		200		250			
No	Fract. Size	Stat.	For	Stat.	For	Stat.	For	Stat.	For	Stat.	For	$\bar{f}_m$	$\bar{f}_m+2\sigma$
2	.01	1	1	2	2	3	2	3	2	3	2	1.51	4.44
10	.1	61	16	53	12	45	8	26	6	40	10	7.58	14.12
2	.3	33	7	19	2	14	1	12	1	18	2	1.51	4.44
2	.5	36	13	24	4	12	0	2	0	0	0	1.51	4.44
1	1.0	1	1	1	1	0	0	0	0	0	0	.76	2.82

TABLE I - (Cont'd)

Values of f_m for $\frac{\lambda_1}{\lambda_0} = .99$

Region Data		No. of Iterations											
		50		100		150		200		250			
No.	Fract. Size	Stat.	For	Stat.	For	Stat.	For	Stat.	For	Stat.	For	$\bar{f}_m$	$\bar{f}_m+2\sigma$
2	.01	5	2	1	0	3	1	4	2	6	3	1.51	4.44
10	.1	50	9	36	3	42	5	41	4	56	8	7.58	14.12
2	.3	24	2	2	0	8	0	14	1	17	2	1.51	4.44
2	.5	20	2	0	0	10	0	17	0	24	2	1.51	4.44
1	1.0	0	0	0	0	0	0	0	0	0	0	.76	2.82

Values of f_m for $\frac{\lambda_1}{\lambda_0} = .995$

Region Data		No. of Iterations											
		50		100		150		200		250			
No	Fract. Size	Stat.	For	Stat.	For	Stat.	For	Stat.	For	Stat.	For	$\bar{f}_m$	$\bar{f}_m+2\sigma$
2	.01	5	3	5	3	7	1	6	4	7	3	1.51	4.44
10	.1	81	23	124	26	90	13	95	13	69	10	7.58	14.12
2	.3	46	9	29	3	13	1	8	0	16	1	1.51	4.44
2	.5	38	5	18	2	6	1	25	2	24	2	1.51	4.44
1	1.0	0	0	1	1	1	1	1	1	1	0	.76	2.82

DISCUSSION

Gelbard: You say that in one case run at Knolls Atomic Power Laboratory (KAPL) it was observed that there was a tenth of a percent bias in the eigenvalue. How many histories were run per generation in that case?

Gast: Five hundred. The 0.1% bias was not considered acceptable by people at KAPL. Their approach was simply to go to approximately 2000 histories per iteration on the assumption that the bias would then drop to 0.025%.

Gelbard: Did they have other information about other problems? For example, was this 0.1% particularly large or was it typical of the problems they ran?

Gast: Their conclusion was that it could be typical.

Kalos: I was there at the time. They were trying to compute k's to about 0.1% and they wanted their errors to be small compared with that.

Gelbard: Any questions or comments?

Kalos: I have a great many comments, because I think that this was a very stimulating, careful, and interesting paper. First of all, once you have decided for various reasons that you must make your simulated system supercritical, why insist on any restrictions on the number of histories in successive generations, or on the number of fissions? Why not simply let these numbers vary freely and prepare yourself to contend with the mild programming problem of handling lists of variable length. The lists will probably go out on discs anyway, and you totally remove the bias. The bias associated with increasing n is removed from the problem, and the game is fairer still. It seems to me that the use of an increasing n presents complications in understanding and analyzing errors. As I remarked to you before, a scheme for analyzing errors that I find seems to work in situations of this kind is to do the following. Assume that I am going to do 120 iterations, of which the first twenty are to be ignored. I then divide the remaining 100 into groups of 20, and I calculate a mean for each group. This, of course, is a biased mean, and the bias is different from the bias of the grand ensemble mean. Then I compute the standard deviation of the individual means, estimate serial correlation coefficients between each group, and infer a statistical error for the whole problem. It is a rather simple strategy.

Gelbard: Let me comment on one of the suggestions that Kalos was making. The suggestion seems to be essentially that, when you average over generations, you do this without normalizing the eigenvectors per generation. You then end up with an arbitrarily normalized eigenvector which you may normalize, of course, after you are finished. Suppose that is what you do. Now that you have normalized you will introduce a bias again, because the bias comes from the fluctuation in the denominator when you normalize. That is the conclusion that comes out of the analysis that we are going to talk about in our next paper. The bias comes from nonlinear terms in the iteration equations, the nonlinear terms being introduced by the denominator in the normalization process. So, if the bias is to be decreased by the process that you described, the amplitude of the fluctuations in the denominator must be decreased. Now, you might expect that they would be decreased because the denominator now

comes, not from one generation, but from a sum over all of the generations which are being averaged. In other words, if you take as your estimate of the eigenfunction a sum over, let us say 100 generations, then to normalize, you divide by the volume integral of this sum. You might expect that the denominator, being a mean over many generations, might not fluctuate much. Mal, do you disagree?

Kalos: I disagree with the whole argument. There is no fundamental normalization here for the eigenfunction anyway.

Gelbard: That's the real question. What I am asserting is that when you go to use the eigenvector, you always find yourself forced to normalize it in some way. You may want to estimate ratios between the eigenvector at one point and another, and when you find these ratios, you are adopting a particular normalization procedure. The simplest normalization to adopt, in principle, is the one where you normalize the final average to one.

Kalos: But in estimating ratios my procedure introduces no bias.

Gelbard: That is where I disagree. I disagree because any estimate you make eventually is based on a normalized eigenvector.

Kalos: But the normalization drops out.

Gelbard: Why does it drop out? You divide the vector you are getting by an integral of that vector.

Kalos: Suppose I decided to divide the eigenvector by 73.

Gelbard: If you always divided it by 73 you would again have an unnormalized eigenfunction.

Kalos: Right!

Gelbard: And you could not compare the value at one point with the value at another point. To make a comparison of values at one point with values at another point, you always must, in one way or another, introduce a normalization.

Kalos: I am afraid I don't understand that remark at all.

Gelbard: What I am saying is that you cannot get away from normalizing the eigenvector that comes out of a Monte Carlo code calculation. I am saying that the value of the eigenvector at a point, unnormalized, is not useable, is not what you are after. You are after a value at a point normalized in some way, or else the value at one point divided by the value at another point.

Kalos: Well, if it is the value at one point divided by value at another point, then obviously the normalization is irrelevant.

Gelbard: Well, I would say that the normalization then is the normalization to the value at the point at which you are getting the ratio. But another way of doing the same thing, of comparing different eigenvectors, or eigenvectors at different points, is to normalize the eigenvector so that its

integral is one. This is the most convenient way of normalizing; and once you normalize in this way, you are taking the eigenfunction and dividing it by a fluctuating denominator, and you have a bias again. So I think that the elimination of the bias in this way, although it sounds attractive at first glance, is really an illusion. I think that you must normalize the eigenvector, eventually, when you use it for any practical purpose, and when you do, you will again introduce a bias. The question is whether you have diminished the bias by normalizing in this modified way.

Coveyou: Are you talking about taking the value of the eigenfunction at two different points and comparing them?

Gelbard: I say that that is what you usually do when you compute the eigenvector.

Cashwell: Isn't it clear then that, when you take the ratio, normalization cancels out?

Kalos: Therefore, the bias that comes from the denominator is totally irrelevant, and if you bring it into the analysis of variance you are making a mistake?

Gelbard: Let me then define an alternative way of normalizing. I will take one point in the problem and normalize to one. That is another way of normalizing. I also say that it is a biased way of normalizing.

Kalos: All ways of normalizing are biased.

Gelbard: So, I am saying that you cannot get away from a normalization procedure of one kind or another and when you introduce a normalization procedure, whether you like it or not, you will come back with a bias.

Kalos: It is true that the ratio of the two estimates of the eigenfunction, being a ratio, is very likely to be a biased quantity. But, one must deal with the bias in that ratio directly, and in fact the bias introduced by another normalization factor, is totally irrelvant.

Gelbard: What I am saying is that some sort of normalization of the eigenvector is inescapable when you go to use or exhibit your results; and whatever normalization you use will introduce bias. Every time that you introduce a normalization procedure, you introduce the bias again. The question is: How large is the bias for different normalization procedures? Suppose you normalize in the way I described, namely you get the sum over many many unnormalized generations, then divide by that sum. You now have an alternate way of normalizing. Now, you may ask, how is this mean biased, compared to the other mean? I think that is a question that you have to look at if you are going to propose this procedure. I think you have to recognize that the end result is going to be normalized; the most convenient way to normalize is to divide by the integral. You can ask: what kind of bias you have once that normalization is carried out. Is it smaller than the bias you would get by normalizing each eigenvector and averaging? This I think is a question one ought to be able to answer.

Cashwell: Yes, but still, the question is: What is biased? The estimate of the normalized eigenfunction is biased, each estimate separately. The ratio of two of the values of the eigenfunction is biased also. But, the second is not biased because the first one is biased. The ratio is biased because it is a ratio.

Gelbard: Normalizing to a value at one point is one way of normalizing. I am saying a more common, and also perfectly satisfactory way, is to normalize so that the integral is one. In any case you must normalize to exhibit your results, and to use them, and then you come back with a bias. The question that you should ask is whether one of these approaches gives you less bias in the normalized result than the other. My suspicion is that the bias in this alternative procedure, which does sound attractive when you think about it first, is not really smaller than the bias in the conventional normalization.

Gast: I would like to make one point about normalization. In our RECAP program the eigenfunction can be represented, either by the accumulated neutron source which is input at the beginning of each iteration, or by the accumulated output of an iteration, the neutron production. You can visualize normalizing either end of the generation. We always normalize the input. To normalize the input one keeps the total starting weight of all neutrons the same in each generation. If you want to normalize the output, the neutron production (and I am not sure whether you were referring to this), the process becomes much more complex.

Gelbard: Another way of avoiding bias in the eigenvector is to keep a fixed number of histories per generation but, when you average, to assign each generation a weight — the weight being equal to the eigenvalue coming from that generation. I think this approach is very closely related to what is in one of Lieberoth's papers. But I suspect, again, that when you try to use the eigenfunction you are forced to normalize it. And, when you normalize it, you are again forced to divide by a fluctuating denominator, and again I think that you will get a bias when you finally do normalize. It is just hard to see what you can do with the unnormalized shape: every time you try to use it you are forced to normalize in one way or another. Well, anymore comments?

Kalos: Yes, I would like to comment on the subject of importance sampling. You conjecture that there is no other form of importance sampling which would have direct application to the reduction of variance of the eigenfunction or, in general, to the acceleration of the iterative process. This is false. There is such a form of importance sampling, at least, in theory. Work on this sort of importance sampling originally goes back to Goad and Johnson. A variant of their method was published by me in a paper presented at the ANS Topical Meeting in Kiamesha Lake. There I pointed out that if you solve the adjoint transport equation, and if you use the fundamental mode of the power distribution as an importance function then, in the limit where ordinary importance sampling is carried out exactly, you get zero variance estimates for the flux at any point you care to choose to start the history. This process converges in one generation and, therefore, you have an acceleration and zero variance technique at the same time. The estimates you need to bias with are the value of the power. Thus you have the possibility of starting with an approximation and, perhaps, doing a multi-stage sampling in which improvements in your guess are fed back. I have no idea how this would work out in practice but at least, in theory, there is an importance sampling method

which reduces variance and accelerates convergence. It is of the same general character as those you have looked at, but the point of view is quite different.

Gast: Our feeling was that because we need the flux or the power over many edit regions, that in trying to get a shape we would be confronted with a sequence of adjoints, and we ....

Kalos: No, it is one adjoint only — there is a single adjoint. There is a single adjoint which gives the flux at any point in the reactor with zero variance.

Gast: At a prechosen point. Then you have to do many points.

Kalos: Yes, yes.

Gast: It is our intuitive feeling (though we intend to look into this in more detail, by the way) that we could probably not gain very much relative to a normal forward run.

Gelbard: Because of the number of points you are interested in?

Gast: Yes, but I am very interested in your feelings as to what the gain could be.

Kalos: I wish I had feelings as to what the gain could be.

Cashwell: My experience has been mainly with shielding problems, not criticality problems, but my guess is that almost any method of getting an approximate adjoint function will help a lot. Whether that is true for the reactor, for the criticality problem, I do not know, but I see no reason why it would not be true.

MONTE CARLO WORK AT ARGONNE NATIONAL LABORATORY

E. M. Gelbard and R. E. Prael

Argonne National Laboratory
Argonne, Illinois 60439, U.S.A.

* * *

Presented by E. M. Gelbard

MONTE CARLO WORK AT ARGONNE NATIONAL LABORATORY

E. M. Gelbard and R. E. Prael

Argonne National Laboratory
Argonne, Illinois 60439, U.S.A.

ABSTRACT

This paper is divided into three parts. In Part I we describe a simple model of the Monte Carlo process and develop a (nonlinear) recursion relation between fission sources in successive generations. From the linearized form of these recursion relations, it is possible to derive expressions for the mean square coefficients of error modes in the iterates and for correlation coefficients between fluctuations in successive generations. In a simple test problem, we find that MacMillan's prescription for computing variances in means in fairly accurate.

Part II deals with first-order nonlinear terms in the recursion relation. From these nonlinear terms an expression for the bias in the eigenvalue estimator is derived, and prescriptions for measuring the bias are formulated.

In Part III we outline plans for the development of the VIM code, and describe briefly the proposed treatment of small sample perturbations in VIM.

Part I

THEORETICAL STUDY OF THE MONTE CARLO EIGENVALUE COMPUTATION, LINEARIZED EQUATIONS

A. Basic Features of the Mathematical Model

Although the mathematical principle underlying the Monte Carlo method in fixed-source computations are well understood, the Monte Carlo methods used in eigenvalue calculations have generally been based primarily on intuition. It is known that many of the methods in use today are biased [1,2]. It is known that correlations between generations tend to complicate the statistical estimation process [3]; but one finds very little theoretical work in the technical literature on the magnitude of the bias, or on the nature of these correlations. We propose here a mathematical model of the Monte Carlo process, a model from which perhaps some of the mathematical properties of the Monte Carlo process can be deduced. For convenience we deal with a $K \times K$ matrix eigenvalue equation rather than an integral equation.

Suppose that, in the Monte Carlo process, we treat precisely N sample particles per generation. One is to imagine that these particles execute a random walk among K boxes. Let $T_{\mu\nu}$ be the probability that a particle in the ν-*th* box will step, next, into the μ-*th*. Suppose, further, that the μ-*th* box has a fixed weight, W_μ ($\mu = 1, 2, \ldots, K$), attached to it. Define a diagonal matrix, W, such that its $\mu\mu$-*th* element is equal to the weight, W_μ. Let $WT = H$. Then our matrix analogue of the reactor eigenvalue equation is the matrix equation

$$H\,\underline{U}_n = \lambda_n \underline{U}_n \,. \tag{1}$$

It should be understood that each box in our model represents an elementary volume of phase-space in the reactor. We are interested in estimating the largest eigenvalue, λ_1, the corresponding eigenvector $\underline{U}_1$ and, perhaps integrals of $\underline{U}_1$ over specified regions in the reactor configuration.

Assume that, at the end of the i-1'*st* generation, N^i_μ particles have terminated their histories in the μ-*th* box. Here, of course,

$$\sum_{\mu=1}^{K} N^i_\mu = N \,,$$

where K, again, in the order of the matrix, H.

To begin the next generation we form the probability distribution

$$p^i_\mu \equiv V^i_\mu \Bigg/ \left[\sum_{\nu=1}^{K} V^i_\nu\right] , \qquad V^i_\mu \equiv N^i_\mu W_\mu \,, \tag{2}$$

and select N starters for the i-*th* generation from this distribution. In vector notation

$$\underline{V}^i = \left(V^i_1, V^i_1, \ldots V^i_\mu, \ldots, V^i_k\right)^T . \tag{3}$$

Note that in our model, each particle of the i-1'*st* generation produces one, and only one, potential fission site for the i-*th*. Thus, our analysis will not apply to a process in which the site of each collision is a potential fission site or, generally, to any process in which the number of fission sites generated, per history, is a random variable. Note that every box in which a history terminates is, in our terminology, a "potential fission site", even if the weight assigned to that box is zero. This peculiarity in terminology is of no importance here, since in order to avoid substantial theoretical complications, we shall assume that all the weights are positive. Although we do not allow any weights to vanish, the W_μ may be arbitrarily small, and it seems reasonable to assume that a Monte Carlo process in which some weights are very small will be practically indistinguishable from a process in which those weights are set to zero.

Define a K-dimensional vector $\underline{\tau}^T = (1, 1, \ldots, 1)$. Clearly $p^i_\mu = V^i_\mu/(\underline{\tau}^T \cdot \underline{V}^i)$. The vectors $\underline{V}^i$ and $\underline{V}^{i+1}$ are connected by the recursion

relation

$$\underline{V}^{i+1} = N \frac{H\,\underline{V}^{i}}{\underline{\tau}^{T} \cdot \underline{V}^{i}} + \underline{\varepsilon}^{i} , \tag{4}$$

$$\underline{V}^{i+1} = \left(1/\lambda^{i}\right) H\,\underline{V}^{i} + \underline{\varepsilon}^{i} , \qquad \lambda^{i} \equiv \left(\underline{\tau}^{T} \cdot \underline{V}^{i}\right) \Big/ N . \tag{5}$$

Here λ^{i} is a familiar i-1'*st* generation estimate of the dominant eigenvalue while $\underline{V}^{i}$ is an estimate of the corresponding eigenvector. In Eq. (5), $\underline{\varepsilon}^{i}$ is the net statistical error committed during the i-*th* generation in transferring particles from their starting locations into the μ-*th* box. More precisely

$$\varepsilon_{k}^{i} \equiv V_{k}^{i+1} - E\left\{V_{k}^{i+1} \middle| V^{i}\right\} = V_{k}^{i+1} - \bar{V}_{k}^{i+1} . \tag{6}$$

It will be seen that $\bar{V}_{k}^{i+1}$ is the conditional mean of V_{k}^{i+1}, given V^{i}. From Eq. (6),

$$E\left\{\varepsilon_{k}^{i}\right\} = E\left\{E\left\{\varepsilon_{k}^{i} \middle| \underline{V}^{i}\right\}\right\} = E\left\{E\left\{V_{k}^{i+1} - \bar{V}_{k}^{i+1} \middle| \underline{V}^{i}\right\}\right\}$$

$$= E\left\{\bar{V}_{k}^{i+1} - \bar{V}_{k}^{i+1}\right\} = 0 . \tag{7}$$

Of course $\underline{V}^{i}$ and $\underline{V}^{i-1}$ are not independent, so that $\underline{\varepsilon}^{i}$ and $\underline{\varepsilon}^{i-1}$ cannot be independent either. But

$$E\left\{\varepsilon_{k}^{i}\varepsilon_{\ell}^{i-1}\right\} \equiv \int P\left(\underline{V}^{i},\underline{V}^{i-1}\right) d\underline{V}^{i}\, d\underline{V}^{i-1} \left[V_{k}^{i} - E\left\{V_{k}^{i} \middle| \underline{V}^{i-1}\right\}\right]$$

$$\times \int P\left(\underline{V}^{i+1} \middle| \underline{V}^{i}\right) d\underline{V}^{i+1} \left[V_{\ell}^{i+1} - E\left\{V_{\ell}^{i+1} \middle| \underline{V}^{i}\right\}\right] .$$

Here $P(\underline{V}^{i},\underline{V}^{i-1})$ is the joint probability distribution function of $\underline{V}^{i}$ and $\underline{V}^{i-1}$, while $P(\underline{V}^{i+1}|\underline{V}^{i})$ is the conditional distribution function of $\underline{V}^{i+1}$ given $\underline{V}^{i}$. By definition

$$\int P\left(\underline{V}^{i+1} \middle| \underline{V}^{i}\right) d\underline{V}^{i+1} \left[V_{\ell}^{i+1} - E\left\{V_{\ell}^{i+1} \middle| \underline{V}^{i}\right\}\right] = 0 ,$$

and it follows that

$$E\left\{\varepsilon_{k}^{i}\varepsilon_{\ell}^{i-1}\right\} = 0 .$$

Thus, although $\underline{\varepsilon}^i$ and $\underline{\varepsilon}^{i-1}$ are not independent, they are, in the technical sense, uncorrelated. It will be necessary, later, to examine the properties of the statistical errors, $\underline{\varepsilon}^i$, more closely, but we need say no more about them at this point. Note that it follows from Eq. (7) that

$$E\left\{\underline{V}^{i+1}\right\} = H\, E\left\{\underline{V}^i/\lambda^i\right\} . \tag{8}$$

Define

$$\underline{V}^i = N\,\underline{U}_1 + \underline{\delta}^i , \tag{9}$$

$\underline{U}_1$ being the main mode eigenvector normalized so that $\underline{\tau}^T \cdot \underline{U}_1 = \lambda_1$. From Eqs. (5) and (9) we see that

$$N\,\underline{U}_1 + \underline{\delta}^{i+1} = N\,\frac{H\left(N\,\underline{U}_1 + \delta^i\right)}{\underline{\tau}^T \cdot \left(N\,\underline{U}_1 + \underline{\delta}^i\right)} + \underline{\varepsilon}^i . \tag{10}$$

Equation (10) is clearly nonlinear in $\underline{\delta}^i$. In our analysis it will be convenient to treat separately those effects which can be deduced from the <u>linearized</u> form of Eq. (10), and those whose study requires retention of higher terms. Intergeneration correlations fall into the first category, biases into the second.

B. <u>The Linearized Iteration Equation and Correlations</u>

Retaining only linear terms in Eq. (10), we get

$$\underline{\delta}^{i+1} \approx \underline{\tilde{\delta}}^{i+1} \equiv \lambda_1^{-1}\left[H - \underline{U}_1\underline{\tau}^T\right]\underline{\tilde{\delta}}^i + \underline{\varepsilon}^i . \tag{11}$$

Taking expectation values of both sides, it is easy to show that $E\{\underline{\tilde{\delta}}^i\} \to 0$ as $i \to \infty$, so that to first order in $\underline{\varepsilon}$ the bias vanishes. In all the work which follows, we will assume that H has a complete set of eigenvectors and that all eigenvalues are real and nonnegative. It is our hope that, even when these assumptions fail, our analysis will still give us some valid insight into the character of the Monte Carlo iterative process. At any rate, assuming completeness, we write

$$\underline{\tilde{\delta}}^i = \sum_n \tilde{\delta}^i_n \underline{U}_n , \tag{12}$$

where, as in Eq. (1), the $\underline{U}_n$ are the eigenvectors of H. We have already specified that $\underline{\tau}^T \cdot \underline{U}_1 = \lambda_1$. It will be convenient to normalize the other eigenvectors similarly if possible. If $\underline{\tau}^T \cdot \underline{U}_n \neq \lambda_n$, let $\underline{\tau}^T \cdot \underline{U}_n = \lambda_n$; otherwise the normalization of $\underline{U}_n$ will remain unspecified. Clearly

$$\underline{\tau}^T \cdot \tilde{\delta}^i = \sum_N{}' \lambda_n \tilde{\delta}_n^i , \tag{13}$$

where the primed summation runs over all n such that $\underline{\tau}^T \cdot \underline{U}_n \neq 0$.

Write Eq. (11) in the form

$$\tilde{\underline{\delta}}^{i+1} = A\tilde{\underline{\delta}}^i + \underline{\varepsilon}^i , \qquad A \equiv \hat{H} - \underline{U}_1 \hat{\underline{\tau}}^T ,$$

$$\hat{H} \equiv H/\lambda_1 , \qquad \hat{\underline{\tau}} \equiv \tau/\lambda_1 . \tag{14}$$

Assume that the iterative process has been in progress for an infinite number of generations. It is easy to show, given our assumptions on H, that $||A|| < 1$. It follows then, from Eq. (14), that

$$\tilde{\underline{\delta}}^{i+1} = \sum_{n=0}^{\infty} A^n \underline{\varepsilon}^{i-n} , \tag{15}$$

where, of course, the superscript on $\underline{\varepsilon}$ is an iteration index, not an exponent.

The reader can verify that

$$A^n = \hat{H}^n - \underline{U}_1 \hat{\underline{\tau}}^T \hat{H}^{n-1} , \qquad n \geq 1 . \tag{16}$$

Putting Eq. (16) into Eq. (15), we see that

$$\tilde{\underline{\delta}}^{i+1} = \underline{\varepsilon}^i + \sum_{n=1}^{\infty}\sum_{k=1}^{K} \rho_k^n \underline{U}_k - \underline{U}_1 \sum_{k=1}^{K}{}' \rho_k^n \varepsilon_k^{i-n}$$

$$= \underline{\varepsilon}^i + \sum_{n=1}^{\infty}\sum_{k=2}^{K} \rho_k^n \underline{U}_k - \underline{U}_1 \sum_{k=2}^{K}{}' \rho_k^n \varepsilon_k^{i-n} . \tag{17}$$

Here $\rho_k \equiv \lambda_k/\lambda_1$, while ε_k^i is the coefficient of the k-*th* eigenvector in the eigenfunction expansion for $\underline{\varepsilon}^i$:

$$\underline{\varepsilon}^i = \sum_{k=1}^{K} \varepsilon_k^i \underline{U}_k . \tag{18}$$

It follows from Eq. (17) that

$$\tilde{\delta}_1^{i+1} = \varepsilon_1^i - \sum_{k=2}^{K}{}' \sum_{n=1}^{\infty} \rho_k^n \varepsilon_k^{i-n} ,$$

$$\tilde{\delta}_k^{i+1} = \sum_{n=0}^{\infty} \rho_k^n \varepsilon_k^{i-n} , \qquad k > 1 . \tag{19}$$

To compute the mean-square coefficients of the error modes, we write

$$\left(\tilde{\delta}_k^{i+1}\right)^2 = \sum_{n=0}^{\infty}\sum_{m=0}^{\infty} \rho_k^n \rho_k^m \varepsilon_k^{i-n} \varepsilon_k^{i-m} , \qquad k \neq 1 . \tag{20}$$

Defining $r_{k\ell} = E\left\{\delta_k^{i+1}\delta_\ell^{i+1}\right\}$, we see from Eq. (20) that

$$r_{kk} = \frac{\varepsilon_{kk}}{1 - \rho_k^2} , \qquad k \neq 1 . \tag{21}$$

For later use, we also cite, here, expressions for the other $r_{k\ell}$'s:

$$r_{11} = \varepsilon_{11} + \sum_{k=2}^{K}{}' \sum_{\ell=2}^{K}{}' \frac{\rho_k \rho_\ell}{\left(1 - \rho_k \rho_\ell\right)} \varepsilon_{k\ell} ,$$

$$r_{1\ell} = r_{\ell 1} = \varepsilon_{1\ell} - \sum_{k=2}^{K}{}' \frac{\rho_k \rho_\ell}{\left(1 - \rho_k \rho_\ell\right)} \varepsilon_{k\ell} , \qquad \ell \neq 1 ,$$

$$r_{k\ell} = \frac{\varepsilon_{k\ell}}{\left(1 - \rho_k \rho_\ell\right)} , \qquad k \neq 1 , \qquad \ell \neq 1 . \tag{22}$$

Above, and in all the work which follows,

$$\varepsilon_{k\ell} \equiv E\left\{\varepsilon_k^i \varepsilon_\ell^i\right\} . \tag{23}$$

Note that we are assuming that the right-hand side of Eq. (23) is independent of i.

The quantity r_{kk} is the mean-square coefficient of the k-*th* eigenvector in an expansion of the "asymptotic" fission source (i.e., the fission source in any single generation preceded by infinitely many iterations). At this point we want to derive a corresponding expression for r_{kk} in a fission source <u>averaged over M successive generations</u>. Define

$$\tilde{\delta}_{k,M}^{i+1} = \frac{1}{M} \sum_{n=1}^{M} \tilde{\delta}_k^{i+2-n} , \tag{24}$$

and

$$r_{kk,M} = E\left\{\left(\tilde{\delta}^{i}_{k,M}\right)^2\right\} , \tag{25}$$

$$r_{kk,M} = \frac{1}{M^2}\sum_{n=1}^{M}\sum_{m=1}^{M} E\left\{\tilde{\delta}^{i+2-n}_{k}\tilde{\delta}^{i+2-m}_{k}\right\} . \tag{26}$$

We consider, for the sake of simplicity, only the variance in the amplitudes of contaminating modes, modes for which $k > 1$. For such modes

$$E\left\{\tilde{\delta}^{i+1-n}_{k}\tilde{\delta}^{i+1-m}_{k}\right\} = \sum_{j=0}^{\infty}\sum_{\ell=0}^{\infty}\rho_k^{j}\rho_k^{\ell}E\left\{\varepsilon_k^{i-n-j}\varepsilon_k^{i-m-\ell}\right\} , \qquad k > 1 . \tag{27}$$

If $n \geq m$, then

$$E\left\{\tilde{\delta}^{i+1-n}_{k}\tilde{\delta}^{i+1-m}_{k}\right\} = \sum_{j=0}^{\infty}\rho_k^{j}\rho_k^{n-m+j}\varepsilon_{kk} , \qquad k > 1 . \tag{28}$$

$$E\left\{\tilde{\delta}^{i+1-n}_{k}\tilde{\delta}^{i+1-m}_{k}\right\} = \frac{\rho_k^{n-m}}{\left(1 - \rho_k^2\right)}\varepsilon_{kk} , \qquad k > 1 . \tag{29}$$

Putting Eq. (29) into Eq. (26), we find after some manipulation that

$$r_{kk,M} = \frac{1}{M}\frac{\varepsilon_{kk}}{\left(1 - \rho_k^2\right)}\left[1 + 2h_M\left(\rho_k\right)\right] , \qquad k > 1 , \tag{30}$$

where

$$h_M\left(\rho_k\right) \equiv \frac{\rho_k}{\left(1 - \rho_k\right)}\left[1 - \frac{1}{M}\frac{\left(1 - \rho_k^M\right)}{\left(1 - \rho_k\right)}\right] . \tag{31}$$

Obviously $h_M(\rho) \to \rho/(1 - \rho)$ as $M \to \infty$. It is easy to show in addition that $h_M(\rho) \approx (M - 1)/2$ if $\rho^M \approx 1$. Thus

$$r_{kk,M} \approx \frac{\varepsilon_{kk}}{\left(1 - \rho_k^2\right)} , \qquad \text{if } M << -\left(\ln \rho_k\right)^{-1} , \qquad k > 1 , \tag{32}$$

and

$$r_{kk,M} \approx \frac{1}{M}\frac{\varepsilon_{kk}}{\left(1-\rho_k^2\right)}\left[1+\frac{2\rho_k}{\left(1-\rho_k\right)}\right] = \frac{1}{M}\frac{\varepsilon_{kk}}{\left(1-\rho_k\right)^2},$$

$$\text{if } M \gg -\left(\ln \rho_k\right)^{-1}, \qquad k > 1 . \tag{33}$$

Note that if successive generations were uncorrelated, one would find that

$$r_{kk,M} = \frac{1}{M}\frac{\varepsilon_{kk}}{\left(1-\rho_k^2\right)} \qquad \text{(without correlation)} . \tag{34}$$

Thus we see that the factor $1 + \left[2\rho_k/(1-\rho_k)\right]$ in Eq. (33) is an amplification factor embodying effects due to intergeneration correlations. A closely related amplification factor is introduced by MacMillan in Ref. 3.

In order to make any use of the expressions derived above, it will be necessary to evaluate the covariances ε_{kk}. We discuss this problem next.

C. Computation of $\varepsilon_{k\ell}$

The quantity $\underline{\varepsilon}^i$ in Eq. (5) is a vector whose spatial components have the form

$$\varepsilon_\nu^i = W_\nu \Delta_\nu^{i+1}, \qquad \Delta_\nu^{i+1} \equiv N_\nu^{i+1} - E\left\{N_\nu^{i+1} \middle| \underline{V}^i\right\} . \tag{35}$$

Expanding $\underline{\varepsilon}^i$ into eigenvectors of H, we find that

$$\varepsilon_k^i = \sum_{\nu=1}^{K} U_{k\nu}^* W_\nu \Delta_\nu^{i+1} . \tag{36}$$

Here $U_{k\nu}^*$ is the ν-*th* components of the adjoint eigenvector $\underline{U}_k^*$. Thus

$$\varepsilon_{k\ell} = \sum_{\nu=1}^{K}\sum_{\mu=1}^{K} U_{k\nu}^* U_{\ell\mu}^* W_\nu W_\mu E\left\{\Delta_\nu^{i+1}\Delta_\mu^{i+1}\right\} . \tag{37}$$

To evaluate $\varepsilon_{k\ell}$, we must now compute the expectation value which appears on the right-hand side of Eq. (37).

Neglecting any bias, we shall assume that the expected number of particles in the ν-*th* box is $\bar{N}_\nu$, where

$$W_\nu \bar{N}_\nu = N\, U_{1\nu} . \tag{38}$$

Further, we suppose that in each generation the probability that a typical particle will be assigned to the ν-*th* box is equal to $\bar{N}_\nu/N$. Under these assumptions one can show that

$$E\left\{\Delta_\nu^{i+1}\Delta_\nu^{i+1}\right\} = \bar{N}_\nu\left(1 - \bar{N}_\nu/N\right) , \tag{39}$$

$$E\left\{\Delta_\nu^{i+1}\Delta_\mu^{i+1}\right\} = -\bar{N}_\nu\bar{N}_\mu/N , \qquad \nu \neq \mu . \tag{40}$$

Putting Eqs. (39) and (40) into Eq. (37), we find that

$$\begin{aligned}\varepsilon_{k\ell} &= \sum_{\nu=1}^{K} U_{k\nu}^* U_{\ell\nu} W_\nu^2 N_\nu\left(1 - \bar{N}_\nu/N\right) \\ &\quad - \sum_{\nu=1}^{K}\sum_{\mu=1}^{K}\left(U_{k\nu}^* U_{\ell\mu}^* W_\nu W_\mu \bar{N}_\nu \bar{N}_\mu\right)\Big/N \\ &\quad + \sum_{\nu=1}^{K}\left(U_{k\nu}^* U_{\ell\nu}^* W_\nu^2 \bar{N}_\nu^2\right)\Big/N ,\end{aligned} \tag{41}$$

$$\begin{aligned}\varepsilon_{k\ell} &= \sum_{\nu=1}^{K} U_{k\nu}^* U_{\ell\nu}^* W_\nu^2 N_\nu \\ &\quad - \sum_{\nu=1}^{K}\sum_{\mu=1}^{K}\left(U_{k\nu}^* U_{\ell\mu}^* W_\nu W_\mu \bar{N}_\nu \bar{N}_\mu\right)\Big/N .\end{aligned} \tag{42}$$

Because of the biorthogonality of eigenvectors and adjoints, the second term on the right-hand side of Eq. (42) vanishes unless $k = \ell = 1$:

$$\varepsilon_{k\ell} = \sum_{\nu=1}^{K} U_{k\nu}^* U_{\ell\nu}^* W_\nu^2 \bar{N}_\nu , \qquad k \neq 1 \quad \text{or} \quad \ell \neq 1 . \tag{43}$$

Again using biorthogonality, one finds that

$$\varepsilon_{11} = \sum_{\nu=1}^{K} \left(U^*_{1\nu} W_\nu\right)^2 N_\nu - N \ . \tag{44}$$

It is interesting to note that the vector

$$\underline{I}_k = W \, \underline{U}^*_k \ , \tag{45}$$

(where W is the weight matrix, defined in Section A) is the importance function for absorption. From Eq. (42), it is easy to show that, in general,

$$\varepsilon_{k\ell} = N \sum_{\nu=1}^{K} \left(I_{k\nu} - \bar{I}_k\right)\left(I_{\ell\nu} - \bar{I}_\ell\right) \Pi_\nu \ , \tag{46}$$

where

$$\bar{I}_k \equiv \sum_{\nu=1}^{K} I_{k\nu} \Pi_\nu \ ,$$

$$\Pi_\nu = \bar{N}_\nu / N \ .$$

Clearly Π_ν is the fundamental-mode absorption probability in the ν-*th* box. Thus the quantity $\left(\varepsilon_{k\ell}/N\right)$ can be interpreted as the covariance of the importance functions, I_k and I_ℓ, with respect to the absorption probability distribution function.

D. Qualitative Features of the Monte Carlo Estimates

Note that we have, so far, been discussing the fluctuations in the net fission source produced in a typical generation by N sample particles. The amplitude of such a fission source would, of course, be proportional to N. It is, however, somewhat more interesting to consider the behavior of a fission source whose normalization is independent of N. Define

$$\underline{v}^{i+1} = \underline{V}^{i+1}/N \ . \tag{47}$$

It will be seen that $\underline{v}^i$ is an estimate of the fission source density produced, at the end of the i-*th* generation, by one fission neutron born at the beginning of the i-*th* generation. Correspondingly, we define

$$\underline{\tilde{Y}}^{i+1} = \underline{\tilde{\delta}}^{i+1}/N \ ,$$

$$\bar{\varepsilon}_{k\ell} = \frac{1}{N^2} \varepsilon_{k\ell} \ ,$$

and

$$\bar{r}_{k\ell} = \frac{1}{N^2} r_{k\ell} = E\left\{\tilde{\gamma}_k^{i+1}\tilde{\gamma}_\ell^{i+1}\right\}. \tag{48}$$

In terms of the newly defined quantities, we may write

$$\bar{r}_{k\ell} = \frac{\bar{\varepsilon}_{k\ell}}{\left(1 - \rho_k\rho_\ell\right)}, \qquad k \neq 1, \qquad \ell \neq 1, \tag{49}$$

$$\bar{r}_{11} = \bar{\varepsilon}_{11} + \sum_{k=2}^{K}{}' \sum_{\ell=1}^{K}{}' \frac{\rho_k\rho_\ell}{\left(1 - \rho_k\rho_\ell\right)} \bar{\varepsilon}_{k\ell}, \tag{50}$$

$$\bar{r}_{1\ell} = \bar{\varepsilon}_{1\ell} - \sum_{k=2}^{K}{}' \frac{\rho_k\rho_\ell}{\left(1 - \rho_k\rho_\ell\right)} \bar{\varepsilon}_{k\ell} = \bar{r}_{\ell 1}, \qquad \ell \neq 1, \tag{51}$$

$$\bar{r}_{kk,M} = \frac{1}{M^2} \frac{\varepsilon_{kk}}{\left(1 - \rho_k^2\right)} \left[1 + 2h_M\left(\rho_k\right)\right], \qquad k \neq 1, \tag{52}$$

$$\bar{r}_{kk,M} \approx \frac{\bar{\varepsilon}_{kk}}{\left(1 - \rho_k^2\right)}, \qquad M \gg -\left(\ln \rho_k^{-1}\right), \qquad k \neq 1, \tag{54}$$

$$\bar{\varepsilon}_{11} = \frac{1}{N} \sum_{\nu=1}^{K} \left(U_{1\nu}^* W_\nu\right)^2 \Pi_\nu - 1, \tag{55}$$

$$\bar{\varepsilon}_{k\ell} = \frac{1}{N} \sum_{\nu=1}^{K} U_{k\nu}^* U_{\ell\nu} W_\nu^2 \Pi_\nu, \qquad k,\ell \neq 1. \tag{56}$$

We see from Eqs. (49)-(51) that the asymptotic fission source contains an "equilibrium" distribution of modes. Although the amplitude of all error modes tends to vanish like 1/N (since the $\bar{\varepsilon}_{k\ell}$ vanish like 1/N), the spectrum of error modes is independent of N.

Generally one would expect that the $\bar{\varepsilon}_{kk}$ would decrease with k, since the number of zeros in U_k^* increases with k. On the other hand there seem to be no reason to suppose that $\bar{\varepsilon}_{kk}$ is very sensitive to k and we shall assume that variations in $\bar{\varepsilon}_{kk}$ are not very important. Certainly this is true in the model problems discussed in Section E. Neglecting variations in $\bar{\varepsilon}_{kk}$, and assuming that $\rho_k \approx 1$, $\rho_\ell \approx 1$, we find that

$$\bar{r}_{kk}/\bar{r}_{\ell\ell} \approx \left(1 - \rho_\ell\right)/\left(1 - \rho_k\right) . \tag{57}$$

Clearly it is not safe to assume (as one does in a deterministic eigenvalue calculation) that only a single error mode survives after many iterations.

Averaging over many generations, we find that

$$\bar{r}_{kk,M}/\bar{r}_{\ell\ell,M} \approx \left(1 - \rho_\ell\right)^2/\left(1 - \rho_k\right)^2 . \tag{58}$$

It will be seen that when M is sufficiently large the spectrum of error modes, which is always independent of N, becomes independent of M as well. In the averaged fission source the importance of error modes with large ρ is considerably enhanced though higher-order error modes persist.

To compute the variance of the eigenvalue estimate, we recall that

$$\lambda^i = \left(\underline{\tau}^T \cdot \underline{V}^i\right) N \approx \lambda_1 + \underline{\tau}^T \cdot \underline{\tilde{\gamma}}^i , \tag{59}$$

$$\Delta\lambda^i \equiv \lambda^i - \lambda_1 \approx \underline{\tau}^T \cdot \underline{\tilde{\gamma}}^i = \sum_{k=1}^{K} \lambda_k \tilde{\gamma}_k^i , \tag{60}$$

$$\left(\Delta\lambda^i/\lambda_1\right) = \sum_{k=1}^{K}{}' \rho_k \tilde{\gamma}_k^i . \tag{61}$$

From Eq. (19) we see that

$$\sum_{k=1}^{K}{}' \rho_k \tilde{\gamma}_k^i = \frac{1}{N}\left[\varepsilon_1^i - \sum_{k=2}^{K}{}' \sum_{n=1}^{\infty} \rho_k^n \varepsilon_k^{i-n} + \sum_{n=0}^{\infty} \sum_{k=2}^{\infty} \rho_k^{n+1} \varepsilon_k^{i-n}\right] , \tag{62}$$

$$\sum_{k=1}^{K}{}' \rho_k \tilde{\gamma}_k^i = \frac{1}{N}\left[\sum_{k=1}^{K}{}' \rho_k \varepsilon_k^i - \sum_{k=2}^{K} \sum_{n=1}^{\infty} \rho_k^n \left(1 - \rho_k\right) \varepsilon_k^i\right] , \tag{63}$$

From Eqs. (61) and (63), it follows that

$$\alpha \equiv E\left\{\left(\Delta\lambda^i/\lambda_1\right)^2\right\} = \sum_{k=1}^{K}{}' \sum_{\ell=1}^{K}{}' \rho_k \rho_\ell \bar{\varepsilon}_{k\ell}$$

$$+ \sum_{k=1}^{K}{}' \sum_{\ell=1}^{K}{}' \frac{\rho_k \rho_\ell \left(1 - \rho_k\right)\left(1 - \rho_\ell\right)}{\left(1 - \rho_k \rho_\ell\right)} \bar{\varepsilon}_{k\ell} . \tag{64}$$

Define

$$\alpha_0 = \sum_{k=1}^{K}\sum_{\ell=1}^{K} \rho_k \rho_\ell \varepsilon_{k\ell} = \frac{1}{N^2} E\left\{\left(\underline{\tau}^T \cdot \underline{\varepsilon}^i\right)^2\right\} . \tag{65}$$

After a little manipulation, one finds that

$$\alpha_0 = \frac{1}{N}\sum_{\nu=1}^{K}\left(W_\nu - \bar{W}\right)^2 \Pi_\nu , \qquad \bar{W} \equiv \sum_{\nu=1}^{K} W_\nu \Pi_\nu , \qquad \Pi_\nu \equiv \frac{\bar{N}_\nu}{N} . \tag{66}$$

This α_0 is equal to 1/N times the variance of the weights. It can be shown that α_0 has, in addition, the following significance. Suppose one were to draw N fission source sites from the main mode fission source distribution, starting one sample particle from each site. Each sample particle is to be followed for <u>one generation</u>. Let N_ν be the number of sample particles which end their histories in box ν, and define

$$\lambda = \sum_{\nu=1}^{K} N_\nu W_\nu . \tag{67}$$

Then α_0 is the variance in (λ/λ_1). The second term on the right of Eq. (66) is, then, the additional variance induced by the iterative process. This second term may be written in the form

$$\alpha_1 = \sum_{n=1}\sum_{k=2}^{K'}\sum_{\ell=2}^{K'} \rho_k^n \rho_\ell^n \left(1 - \rho_n\right)\left(1 - \rho_\ell\right) \bar{\varepsilon}_{k\ell} . \tag{68}$$

Since the $\bar{\varepsilon}_{k\ell}$ are elements of a covariance matrix, and such a matrix must be non-negative definite, α_1 is a sum of non-negative quadratic forms. It is, therefore, non-negative. The iterative process will tend to increase (and will certainly not decrease) the variance of the eigenvalue estimator.

We have seen that in the asymptotic mixture of modes, the modes with $\rho \approx 1$ ("low-order" modes) occur with relatively large mean-square amplitudes. On the other hand, because of the factor $(1 - \rho_k)(1 - \rho_\ell)$, such modes do not play a particularly important role in amplifying the variance of the eigenvalue estimator. In fact, the fluctuations in amplitude of low-order modes has little effect on the eigenvalue estimates precisely because the eigenvalues of such modes are almost the same as the main mode eigenvalue. Because the contributions of low-order modes are suppressed in α_1, one would expect the net amplification in the variance of the eigenvalue to be small.

Suppose, now, that we are interested in estimating the fission rate in a volume, k, within a reactor configuration, From Eq. (48)

$$\underline{v}^i \approx \underline{U}_1 + \tilde{\underline{\gamma}}^i , \qquad \underline{v}^i - \underline{U}_1 = \sum_{k=1}^{K} \tilde{\gamma}^i_k \underline{U}_k , \tag{69}$$

$$\underline{r}^i_R \equiv \int_R \left(\underline{v}^i - \underline{U}_1\right) dV = \sum_{k=1}^{K} C_k \tilde{\underline{\gamma}}^i_k , \qquad C_k \equiv \int_R \underline{U}_k dV , \tag{70}$$

$$\sigma^2_R \equiv E\left\{\left(\underline{r}^i_R\right)^2\right\} = \sum_{k=1}^{K}\sum_{\ell=1}^{K} C_k C_\ell \bar{r}_{k\ell} . \tag{71}$$

Define

$$\sigma^2_{R,M} = E\left\{\left(\frac{1}{M}\sum_{i=1}^{M} \underline{r}^i_R\right)^2\right\} , \tag{72}$$

and

$$\bar{r}^{i,i+j}_{k\ell} \equiv E\left\{\tilde{\gamma}^i_k \tilde{\gamma}^{i+j}_\ell\right\} . \tag{73}$$

By straightforward computation, we find that

$$\bar{r}^{i,i+j}_{11} = -\sum_{\ell=1}^{K} \varepsilon_{1\ell}\rho^j_\ell + \sum_{k=2}^{K}\sum_{\ell=2}^{K} \rho^j_\ell \frac{\rho_k\rho_\ell}{\left(1-\rho_k\rho_\ell\right)} \varepsilon_{k\ell} , \tag{74}$$

$$\bar{r}^{i,i+j}_{1k} = -\bar{\varepsilon}_{1k}\rho^j_k + \sum_{\ell=2}^{K} \rho^j_k \frac{\rho_k\rho_\ell}{\left(1-\rho_k\rho_\ell\right)} \bar{\varepsilon}_{k\ell} , \qquad k \neq 1 , \tag{75}$$

$$\bar{r}^{i,i+j}_{k1} = -\sum_{\ell=2}^{K} \rho^j_\ell \frac{\rho_k\rho_\ell}{\left(1-\rho_k\rho_\ell\right)} \bar{\varepsilon}_{k\ell} , \qquad k \neq 1 , \tag{76}$$

and

$$\bar{r}^{i,i+j}_{k\ell} = \rho^j_\ell \frac{\rho_k\rho_\ell}{\left(1-\rho_k\rho_\ell\right)} \bar{\varepsilon}_{k\ell} , \qquad k \neq 1, \qquad \ell = 1 . \tag{77}$$

Making use of the above expressions we find that

$$E\left\{r_R^i r_R^{i+j}\right\} = \sum_{k=2}^{K} \alpha_k \rho_k^{j-1} , \qquad j \geq 1 , \tag{78}$$

where the α_k are complicated functions of all the ρ's and all the $\bar{\varepsilon}$'s. Using MacMillan's approach, we write

$$\sigma^2_{R,M} = \frac{1}{M}\sigma_R^2 + \frac{2\sum_{i=2}^{M} E\left\{r_R^i r_R^{i-1}\right\} + 2\sum_{i=3}^{M} E\left\{r_R^i r_R^{i-2}\right\} + \cdots}{M^2} , \tag{79}$$

so thus, for large M

$$\sigma^2_{R,M} \approx \frac{1}{M}\left[\sigma_R^2 + 2\sum_{k=1}^{K} \frac{\alpha_k}{\left(1 - \rho_k\right)}\right] . \tag{80}$$

Equation (80) is a generalization of the equation for σ on page 74 of Ref. 3, and is identical with MacMillan's expression if only one ρ_k is different from zero. It is possible that one could use Eq. (80) assuming, perhaps, that only two of the ρ_k (ρ_2 and ρ_3 for example) are different from zero. The four parmeters α_2, α_3, ρ_2, and ρ_3 could, then, be fitted to match observed correlation coefficients of lag 1, 2, 3, and 4. Whether or not such a procedure would be practical is not clear at present.

E. Analysis of Model Problems

In our analysis of the Monte Carlo process we have dealt with a random walk in a discretized phase space, i.e., a random walk among K distinct boxes. It is easy, however, to adapt this analysis to a continuous random walk by letting K go to infinity, while replacing sums over ν by integrals over space and energy. The model problems discussed in this section are problems in which the position and energy variables are regarded as continuous.

We consider, as our model problem configuration, a homogeneous cube of thickness T. Suppose that the scalar flux vanishes on all faces of the cube, and that T is large enough so that diffusion theory is valid. In one group the eigenvectors are, then, given by the expression

$$U_{k\ell m} = C_{k\ell m} \sin B_k X \sin B_\ell Y \sin B_m Z , \qquad B_k \equiv kB ,$$

$$B \equiv \pi/T , \qquad k,\ell,m = 1, 2, \ldots , \tag{81}$$

where the $C_{k\ell m}$ are normalization constants. The adjoints take the form

$$U^*_{k\ell m} = \left(8/T^3 C_{k\ell m}\right) \sin B_k X \sin B_\ell Y \sin B_m Z . \tag{82}$$

$$\bar{\varepsilon}_{k\ell m,k\ell m} = \frac{1}{N}\left(8/T^3C_{k\ell m}\right)^2\left(C_{111}/B^3\right)\alpha_k\alpha_\ell\alpha_m \tag{83}$$

for k, ℓ, and m are not all equal to one, and

$$\alpha_k \equiv 1 + 1/(4k^2 - 1) .$$

Our normalization condition requires that

$$C_{111} = \lambda_{111}(B/2)^3 ,$$

so that we may write

$$\bar{\varepsilon}_{k\ell m,k\ell m} = \frac{1}{N}\left[8\lambda_{111}\Big/\left(T^3C_{k\ell m}\right)^2\right]\alpha_k\alpha_\ell\alpha_m . \tag{84}$$

In the proposed model problem

$$\lambda_{k\ell m} = \left[1 + \tau B^2(k^2 + \ell^2 + m^2)\right]^{-1} . \tag{85}$$

Consequently

$$\rho_{k\ell m} \approx 1 - \tau B^2(k^2 + \ell^2 + m^2 - 3)\big/\lambda_{111} . \tag{86}$$

and, to first order in B^2,

$$1 - \rho^2_{k\ell m} = 2\tau B^2(k^2 + \ell^2 + m^2 - 3)\big/\lambda_{111} . \tag{87}$$

Approximately, then

$$\bar{r}_{k\ell m,k\ell m} = \frac{1}{N}\frac{4\lambda^2_{111}\alpha_k\alpha_\ell\alpha_m}{\left(T\ C_{k\ell m}\right)^2\left(k^2 + \ell^2 + m^2 - 3\right)} . \tag{88}$$

Thus the root-mean-square coefficient of the kℓm-*th* mode is

$$C_{k\ell m}\sqrt{\frac{4\lambda_{111}^2\alpha\ \alpha_\ell\alpha}{N\tau B^2\left(T^3C_{k\ell m}\right)^2\left(k^2+m^2+\ell^2-3\right)}} = \sqrt{\frac{4\lambda_{111}^2\alpha\ \alpha_\ell\alpha}{N\tau B^2T^6\left(k^2+\ell^2+m^2-3\right)}}\ . \tag{89}$$

On the other hand, the expected fission source density produced, at the end of the i-*th* generation, by a single fission neutron of the i-1'*st* generation has the form

$$s = \lambda_{111}(B/2)^3 \sin BX \sin BY \sin BZ$$

$$= \lambda_{111}(\pi/2T)^3 \sin BX \sin BY \sin BZ\ . \tag{90}$$

Therefore the ratio, $f_{k\ell m}$, of the root-mean-square coefficient of the $k\ell m$-*th* error mode, to the expected coefficient of the main mode, is given by the expression

$$f_{k\ell m} = \sqrt{\frac{256\ \alpha_k\alpha_\ell\alpha_m}{N\pi^6\tau B^2(k^2+\ell^2+m^2-3)}}\ . \tag{91}$$

The quantity $f_{k\ell m}$ is a rough measure of the relative amplitude of the $k\ell m$-*th* error mode in the fission source distribution. If the main mode leakage probability is 30% (fairly typical for experimental fast reactors), then $\tau B^2 = 0.1$. In this case we find that roughly 15,000 histories would be required, per generation, to reduce contamination by each of the three principle error modes to 1%. On present computers it is generally unfeasible to work with such a large number of histories per generation.

Averaging over M generations we get

$$f_{k\ell m,M} = \sqrt{\frac{512\ \alpha_k\alpha_\ell\alpha_m\lambda_{111}}{N\ M\pi^6(\tau B^2)^2(k^2+\ell^2+m^2-3)^2}}\ , \tag{92}$$

assuming that M is sufficiently large so that Eq. (54) is valid. Now we find that 100,000 histories would be required (i.e. NM = 100,000), after settling to reduce contaminations by each of the three principle error modes to 1%. This is a rather large number of histories, but a Monte Carlo problem with 100,000 histories can be run in roughly three hours, generally, on the latest IBM or CDC computers.

This simple analysis may give us some idea of the nature of the modal spectrum in the asymptotic fission source, and of the dependence of the spectrum on the number of histories. On the other hand, it tells us very little about fluctuations in the source, or in reaction rates, in specified

subregions of the problem configuration. Certainly it would be possible to study local fluctuations in this first model problem. However, to shorten our computations, we turn now to a model problem which is similar to the first, but still simpler. The second problem configuration is again a cube of thickness T, but with reflecting boundary conditions imposed on each face. It will be convenient, in this case, to designate the main mode with subscripts 000, in place of the subscripts 111. In this modified notation

$$U_{k\ell m} = \sum_{k\ell m} h_k(x) h_\ell(y) h_m(z) \ . \tag{93}$$

Here the $\alpha_{k\ell m}$ are normalization constants; further

$$\begin{aligned} h_k(x) &\equiv 1 \ , & k &= 0 \ , \\ h_k(x) &\equiv \cos kBx \ , & k &\neq 0 \ , \end{aligned} \tag{94}$$

where $B \equiv \pi/T$. As for the adjoints

$$\begin{aligned} U^*_{k\ell m} &= \left(1/\alpha_{k\ell m}\right)\left[\beta_k h_k(x) \beta_\ell h_\ell(y) \beta_m h_m(z)\right] \ , \\ \beta_k &\equiv (2/T) \ , & k &\neq 0 \ , \\ \beta_k &= (1/T) \ , & k &= 0 \ . \end{aligned} \tag{95}$$

From Eqs. (93) and (95) we see that

$$\bar{\varepsilon}_{k\ell m, k\ell m} = \frac{\alpha_{000}}{N} \frac{\beta_k \beta_\ell \beta_m}{\alpha^2_{k\ell m}} \quad \text{unless } k = \ell = m = 0 \ . \tag{96}$$

Further it is clear that

$$\varepsilon_{000,000} = 0 \ , \tag{97}$$

and that

$$\varepsilon_{k\ell m, k'\ell' m'} = 0 \quad \text{unless } k = k', \ \ell = \ell', \ m = m' \ . \tag{98}$$

As in Eq. (70) define

$$\underline{r}^i_R = \int_R \left(\underline{v}^i - \underline{U}_1\right) dV = \sum_{k,\ell,m} C_{k\ell m} \tilde{\gamma}_{k\ell m} \ , \tag{99}$$

$$C_{k\ell m} = \int_R U_{k\ell m}\, dV \ . \tag{99}$$

Let R be the cube bounded by the planes $x = 0$, $x = t$, $y = 0$, $y = t$, $z = 0$, and $z = t$. Then

$$C_{k\ell m} = \alpha_{k\ell m} f_k f_\ell f_m \ ,$$

$$f_k = (\sin kBt)/kB \ , \qquad k \neq 0 \ ,$$

$$f_k = t \ , \qquad k = 0 \ . \tag{100}$$

As in Eq. (71)

$$\sigma_R^2 = \sum_{k,\ell,m} \sum_{k',\ell',m} C_{k\ell m} C_{k'\ell'm'} \bar{r}_{k\ell m,k'\ell'm'} \ . \tag{101}$$

But one can show from Eqs. (50), (51), (97), and (98) that

$$r_{000,000} = 0 \ . \tag{102}$$

$$\bar{r}_{k\ell m,k'\ell'm'} = 0 \ , \quad \text{unless } k = k',\ \ell = \ell',\ m = m' \ . \tag{103}$$

Therefore

$$\sigma^2 = 3 \sum_{k=1}^{\infty} \sum_{\ell=0}^{\infty} \sum_{m=0}^{\infty} C_{k\ell m}^2 \bar{r}_{k\ell m,k\ell m} \ , \tag{104}$$

$$\sigma_R^2 = 3 \sum_{k=1}^{\infty} \sum_{\ell=0}^{\infty} \sum_{m=0}^{\infty} C_{k\ell m}^2 \frac{\bar{\varepsilon}_{k\ell m,k\ell m}}{1 - \rho_{k\ell m}^2} \ , \tag{105}$$

$$\sigma_R^2 = 3 \frac{\alpha_{000}}{N} \sum_{k,\ell,m} \frac{\beta_k f_k^2 \beta_\ell f_\ell^2 \beta_m f_m^2}{1 - \rho_{k\ell m}^2} \ , \tag{106}$$

$$\sigma_R^2 = \frac{3}{N} \sum_{k,\ell,m} \frac{\left[F_k F_\ell F_m\right]^2}{1 - \rho_{k\ell m}^2} \ , \tag{107}$$

since $\alpha_{000} = T^3$. Here

$$F_k \equiv \sqrt{\left(\beta_k/T\right)}\ f_k \ ,$$

$$= (t/T) \ , \qquad k = 0 \ ,$$

$$= (\sqrt{2}/\pi) \sin (kBt)/k \ , \qquad k \neq 0 \ . \tag{108}$$

Since high harmonics may contribute substantially to the variance, it seems worthwhile, in Eq. (107), to use a rigorous expression for $\rho_{k\ell m}$. For the sake of simplicity, we assume that the medium in the cube is a pure absorber. Then

$$\rho_{k\ell m} \equiv \lambda_{k\ell m}/\lambda_{000} , \tag{109}$$

$$\lambda_{k\ell m} = \frac{\Sigma_t}{B\sqrt{k^2 + \ell^2 + m^2}} \tan^{-1}\left(\frac{B\sqrt{k^2 + \ell^2 + m^2}}{\Sigma_t}\right) . \tag{110}$$

Via Eqs. (107), (108), and (110), it is possible to compute σ_R^2 numerically to any desired level of accuracy.

By methods now familiar, it can be shown that

$$E\left\{r_R^i r_R^{i-1}\right\} = \frac{3}{N} \sum_{k,\ell,m}^{\infty} \rho_{k\ell m} \frac{\left[F_k F_\ell F_m\right]^2}{1 - \rho_{k\ell m}^2} , \tag{111}$$

and that, for M sufficiently large

$$\sigma_{R,M}^2 = \frac{1}{NM} \sigma_R^2 \left\{1 + \frac{6 \sum_{k,\ell,m} \rho_{k\ell m}\left[F_k F_\ell F_m\right]^2 \Big/ \left(1 - \rho_{k\ell m}^2\right)\left(1 - \rho_{k\ell m}\right)}{\sigma_R^2}\right\} . \tag{112}$$

On the other hand, MacMillan's method gives

$$\sigma_{R,M}^2 \approx \sigma_{R,M}^2 \equiv \frac{1}{NM} \sigma_R^2 \left\{1 + \frac{6 \sum_{k,\ell,m} \rho_{k\ell m}\left[F_k F_\ell F_m\right]^2 \Big/ \left(1 - \rho_{k\ell m}^2\right)}{\sigma_R^2\left(1 - \rho_{100}\right)}\right\} . \tag{113}$$

We see, again, that MacMillan's approximation would be rigorous if all the eigenvalues were zero, except for λ_{000} and for $\lambda_{100} = \lambda_{010} = \lambda_{001}$.

In the absence of correlations we would have

$$\sigma_{R,M}^2 = \frac{1}{NM} \sigma_R^2 .$$

Thus, in Eqs. (112) and (113), the quantities in parentheses act as amplification factors embodying the correlation effects. Define

$$A = 1 + \frac{6 \sum_{k,\ell,m} \rho_{k\ell m} \left[F_k F_\ell F_m\right]^2 \Big/ \left(1 - \rho_{k\ell m}^2\right)\left(1 - \rho_{k\ell m}\right)}{\sigma_R^2}, \tag{114}$$

$$\hat{A} = 1 + \frac{6 \sum_{k,\ell,m} \rho_{k\ell m} \left[F_k F_\ell F_m\right]^2 \Big/ \left(1 - \rho_{k\ell m}^2\right)}{\sigma_R^2\left(1 - \rho_{100}\right)}. \tag{115}$$

Here A is the "true" amplification factor while $\hat{A}$ is the approximate "MacMillan amplification factor". It will be seen that A and $\hat{A}$ depend only on the absorption probability $p = (t/T)^3$, and on $\tau B^2 = B^2/3\Sigma_t^2$. The intergeneration correlation coefficient

$$C \equiv E\left\{r_R^i r_R^{i-1}\right\} \Big/ \sigma_R^2$$

is also determined by p and τB^2. Values of A, $\hat{A}$, and C are listed in Table I for various combinations of p and τB^2. In the last colume of Table I we have displayed the percentage errors in $\sqrt{\hat{A}}$ (in place of the errors in A) since $\sqrt{A}$ is a measure of the uncertainty in the fission rate within R. Note that when $\tau B^2 = 0.1$ the dominance ratio (ρ_{100}) is approximately equal to 0.9, and that this dominance ratio is not at all large for a thermal reactor. On the other hand, $\tau B^2 = 0.3$ corresponds to $\rho_{100} \approx 0.7$, a dominance ratio characteristic of a small fast reactor.

The entries in Table I display some fairly obvious regularities. First of all we see that:

(a) As one would expect, the amplification factors tend to be large when the dominance ratio is close to one. In fact it is clear from Eq. (114) that as $\rho_{100} \to 1$, $A \to \infty$.

(b) Secondly, A and C are large when p is large (i.e., when t/T is large), and both are small when p is small. For $k \ll (1/\pi)(T/t)$, we see that $F_k \approx p$: thus if p is small, F_k is almost constant for a wide range of k's. Many modes, including many for which $\rho_{k\ell m}$ is small, then contribute to the sums in Eqs. (111) and (112). The presence of $\rho_{k\ell m}$ as a factor in these sums then tends to make the correlation coefficient C, and the amplification factor, A, small. On the other hand, when p is large, F_k decreases as 1/k, enhancing the importance of modes near the main mode. A and C then become large.

Finally we see that when $\rho_{100} \approx 0.09$, the MacMillan method is reasonably accurate down to p = 0.005. In leakier reactors with $\rho_{100} \approx 0.07$, the MacMillan method is adequate for p's as low as 0.001. Eventually, as p becomes still smaller, C approaches zero and A approaches one. Significant errors in the MacMillan method do occur, but in a range where p is so small that the binomial estimator (the only absorptions estimator we can deal with here) starts to become ineffective.

TABLE I

Amplification Factors and Correlation Coefficients for Various Values of the Problem Parameters

τB^2	p	C	A	$\hat{A}$	Error in $\hat{A}$ (%)
0.1	0.25	0.86	18.1	18.2	0
0.1	0.1	0.84	16.5	17.8	4
0.1	0.01	0.69	8.9	14.8	29
0.1	0.005	0.62	7.0	13.4	38
0.1	0.001	0.45	3.8	10.0	160
0.3	0.1	0.67	6.4	7.7	10
0.3	0.01	0.49	3.7	5.9	26
0.3	0.001	0.28	2.0	3.8	38

On the other hand there are many other estimators which can be introduced in this range. Track length estimators, collision and line-of-sight estimators are often used, when p is small, to reduce the variance in reaction rate estimates. We cannot say what these estimators will do to correlation coefficients or amplification factors. It is possible, however, that such estimators reduce variance by suppressing the uncorrelated "noise" component of the statistical fluctuations, without strongly affecting correlated large-scale fluctuations. In this case, amplification factors would again become large and a more extended analysis of MacMillan's method would be necessary.

Part II

THEORETICAL STUDY OF THE MONTE CARLO EIGENVALUE COMPUTATION, BIASES

A. Calculation of the Eigenvalue Bias in the Mathematical Model

As discussed in Part I, the retention of terms nonlinear in statistical fluctuations when calculating deviations from the fundamental mode eigenvector will give rise to estimates of the bias in the calculated eigenvalue. Following the definitions of Section A of Part I, one obtains from Eq. (10)

$$\underline{\delta}^{i+1} = A\underline{\delta}^{i} - \frac{1}{N}\left(\hat{\underline{\tau}}^{T} \cdot \underline{\delta}^{i}\right) A\underline{\delta}^{i} + \underline{\varepsilon}^{i} , \tag{116}$$

where A and $\hat{\underline{\tau}}^{T}$ are defined by Eq. (14). The statistical error vector $\underline{\varepsilon}^{i}$ is defined in Eq. (6) and has the property $E\left\{\underline{\varepsilon}^{i}\right\} = 0$. However, due to the presence of nonlinear terms in Eq. (116), the expectation value of $\underline{\delta}^{i+1}$ does not generally vanish even in the limit $i \to \infty$. Since the deviation of the estimated eigenvalue of the i-*th* generation from the fundamental mode eigenvalue is given by

$$\Delta\lambda^i = \lambda^i - \lambda_1 = \frac{\lambda_1}{N} \hat{\underline{\tau}}^T \cdot \underline{\delta}^i , \tag{117}$$

the expected value of the bias may be obtained from

$$\Delta\lambda \equiv \lim_{i\to\infty} E\left\{\Delta\lambda^i\right\} = \frac{\lambda_1}{N} \lim_{i\to\infty} E\left\{\hat{\underline{\tau}}^t \cdot \underline{\delta}^i\right\}. \tag{118}$$

If we retain only terms quadratic in $\underline{\varepsilon}^i$ in Eq. (116), we have the equation

$$\underline{\delta}^{i+1} = A\underline{\delta}^i - \frac{1}{N}\left(\hat{\underline{\tau}}^T \cdot \tilde{\underline{\delta}}^i\right) A\tilde{\underline{\delta}}^i + \underline{\varepsilon}^i , \tag{119}$$

where

$$\tilde{\underline{\delta}}^{i+1} = \sum_{n=0}^{\infty} A^n \underline{\varepsilon}^{i-n} \tag{120}$$

was obtained in Eq. (15). Assuming again that an infinite number of generations have passed, Eq. (119) possesses the solution

$$\underline{\delta}^{i+1} = \tilde{\underline{\delta}}^{i+1} - \frac{1}{N}\sum_{n=0}^{\infty} \left(\hat{\underline{\tau}}^T \cdot \hat{\underline{\delta}}^{i-n}\right) A^{n+1}\tilde{\underline{\delta}}^{i-n} . \tag{121}$$

Defining

$$\alpha \equiv \Delta\lambda/\lambda_1$$

and noting that

$$E\left\{\tilde{\underline{\delta}}^{i+1}\right\} = 0 ,$$

we obtain

$$\alpha = -\frac{1}{N^2}\sum_{n=0}^{\infty} E\left\{\left(\hat{\underline{\tau}}^T \cdot \tilde{\underline{\delta}}^{i-n}\right)\left(\hat{\underline{\tau}}^T \cdot A^{n+1}\tilde{\underline{\delta}}^{i-n}\right)\right\} . \tag{122}$$

Upon substitution of Eq. (120),

$$\alpha = -\frac{1}{N^2}\sum_{n=0}^{\infty}\sum_{p=0}^{\infty}\sum_{q=0}^{\infty} E\left\{\left(\hat{\underline{\tau}}^T \cdot A^p \underline{\varepsilon}^{i-n-p-1}\right)\left(\hat{\underline{\tau}}^T \cdot A^{n+q+1} \underline{\varepsilon}^{i-n-q-1}\right)\right\}$$

or

$$\alpha = -\frac{1}{N^2}\sum_{n=0}^{\infty}\sum_{p=0}^{\infty} E\left\{\left(\hat{\underline{\tau}}^T \cdot A^p \underline{\varepsilon}^{i-n-p-1}\right)\left(\hat{\underline{\tau}}^T \cdot A^{n+p+1} \underline{\varepsilon}^{i-n-p-1}\right)\right\}, \tag{123}$$

since the statistical fluctuations in different generations are uncorrelated. Using Eqs. (18) and (23), we obtain

$$\alpha = -\frac{1}{N^2}\sum_{n=0}^{\infty}\sum_{p=0}^{\infty}\sum_{k}\sum_{\ell}\left(\hat{\underline{\tau}}^T \cdot A^p \underline{U}_k\right)\left(\hat{\underline{\tau}}^T \cdot A^{n+p+1} \underline{U}_\ell\right) \varepsilon_{k\ell} .$$

From Eq. (16), we obtain

$$\hat{\underline{\tau}}^T \cdot A^p \underline{U}_k = \left(\rho_k - 1\right) \rho_k^p ,$$

for $p \geq 1$ and for all modes k such that $\hat{\underline{\tau}}^T \cdot \underline{U}_k \neq 0$. It then follows that

$$\alpha = +\frac{1}{N^2}\sum_{k=1}^{\prime}\sum_{\ell=2}^{\prime} \rho_k\rho_\ell\left(1 - \rho_\ell\right)\sum_{n=0}^{\infty} \rho_\ell^n \varepsilon_{k\ell}$$

$$-\frac{1}{N^2}\sum_{k'=2}^{\prime}\sum_{\ell=2}^{\prime}\sum_{p=1}^{\infty}\left(1 - \rho_k\right)\left(1 - \rho_\ell\right) \rho_k^p\rho_\ell^p \sum_{n=0}^{\infty} \rho_\ell^{n+1}\varepsilon_{k\ell} ,$$

which reduces to

$$\alpha = \frac{1}{N^2}\sum_{k=1}^{\prime}\sum_{\ell=2}^{\prime} \frac{\rho_k\rho_\ell\left(1 - \rho_\ell\right)}{1 - \rho_k\rho_\ell} \varepsilon_{k\ell} . \tag{124}$$

As in Part I, $\rho_k = \lambda_k/\lambda_1$, and the summations designated with primes run over those modes with $\hat{\underline{\tau}}^T \cdot \underline{U}_k \neq 0$.

B. <u>Approximations for Estimating the Eigenvalue Bias</u>

Equation (124), for the fractional eigenvalue bias, may be written as

$$\alpha = \alpha_0 + \alpha_1 + \alpha_2 ,$$

where

$$\alpha_0 = \frac{1}{N^2} \sum_{\ell=2}^{\prime} \rho_\ell \varepsilon_{\ell 1} , \tag{125}$$

$$\alpha_1 = \frac{1}{N^2} \sum_{k=2}^{\prime} \sum_{\ell=2}^{\prime} \rho_k \rho_\ell \left(1 - \rho_\ell\right) \varepsilon_{k\ell}, \tag{126}$$

$$\alpha_2 = \frac{1}{N^2} \sum_{k=2}^{\prime} \sum_{\ell=2}^{\prime} \frac{\left(\rho_k \rho_\ell\right)^2 \left(1 - \rho_\ell\right)}{1 - \rho_k \rho_\ell} \varepsilon_{k\ell} . \tag{127}$$

In the above partition, α_0 contains all contributions which are weighted by the fundamental mode and can be expected to contain a large contribution to the bias; α_1 is the leading term in the series expansion of that part of α which contains contributions only from higher-order modes. The expressions for α_0 and α_1 can be converted into sums over regions rather than over modes by utilizing the following relations:

$$\varepsilon_{k\ell} = \sum_{\mu,\nu} U^*_{k\mu} U^*_{\ell\nu} \varepsilon_{\mu\nu} ,$$

$$\varepsilon_{\mu\nu} = N\left[\delta_{\mu\nu} W_\nu U_{1\nu} - U_{1\mu} U_{i\nu}\right] ,$$

and

$$\sum_{k=1}^{\prime} \rho_k \varepsilon_{k\ell} = \frac{1}{\lambda_1} \sum_{k=1} \sum_{\sigma} \sum_{\mu} \sum_{\nu} U_{k\sigma} U^*_{k\mu} U^*_{\ell\nu} \varepsilon_{\mu\nu}$$

$$= \frac{1}{\lambda_1} \sum_{\sigma} \sum_{\mu} \sum_{\nu} \delta_{\sigma\mu} U^*_{\ell\nu} \varepsilon_{\mu\nu}$$

$$= \frac{1}{\lambda_1} \sum_{\mu} \sum_{\nu} U^*_{\ell\nu} \varepsilon_{\mu\nu} ,$$

where completeness has been assumed. Applying the above relationships to Eq. (125), we obtain

$$\alpha_0 = \frac{1}{N^2}\left[\frac{1}{\lambda_1}\sum_{\mu}\sum_{\nu} U^*_{1\nu}\varepsilon_{\mu\nu} - \sum_{\mu,\nu} U^*_{1\mu}U^*_{1\nu}\varepsilon_{\mu\nu}\right]$$

$$= \frac{1}{N}\sum_{\nu}\left(\frac{1}{\lambda_1} - U^*_{1\nu}\right) U^*_{1\nu}W_\nu U_{1\nu} \ . \tag{128}$$

In Eq. (126), we obtain

$$\alpha_1 = \frac{1}{N^2}\left[\frac{1}{\lambda_1}\sum_{\ell=1}{}' \rho_\ell \left(1 - \rho_\ell\right) \sum_{\mu,\nu} U^*_{\ell\nu}\varepsilon_{\mu\nu}\right.$$

$$\left. - \sum_{\ell=1}{}' \rho_\ell\left(1 - \rho_\ell\right) \sum_{\mu,\nu} U^*_{1\mu}U^*_{\ell\nu}\varepsilon_{\mu\nu}\right] .$$

Noting that

$$\left(1 - \rho_\ell\right) U^*_{\ell\nu} = \sum_{\sigma} U^*_{\ell\sigma}\left[\delta_{\sigma\nu} - \hat{H}_{\sigma\nu}\right]$$

and

$$\sum_{\ell=1}{}' \rho_\ell U^*_{\ell\sigma} = \frac{1}{\lambda_1}\sum_{\ell=1}\sum_{\tau} U_{\ell\tau}U^*_{\ell\sigma} = \frac{1}{\lambda_1}\sum_{\tau}\delta_{\tau\sigma} = \frac{1}{\lambda_1} ,$$

we obtain

$$\alpha_1 = \frac{1}{N^2\lambda_1^2}\sum_{\sigma,\mu\nu}\left[\delta_{\sigma\nu} - \hat{H}_{\sigma\nu}\right]\varepsilon_{\mu\nu} - \frac{1}{N^2\lambda_1}\sum_{\sigma,\mu,\nu}\left[\delta_{\sigma\nu} - \hat{H}_{\sigma\nu}\right] U^*_{1\mu}\varepsilon_{\mu\nu} \ .$$

Inserting the expression for $\varepsilon_{\mu\nu}$ gives

$$\alpha_1 = \frac{1}{N\lambda_1}\sum_{\nu}\left(\frac{1}{\lambda_1} - U^*_{1\nu}\right)\left(1 - \sum_{\sigma}\hat{H}_{\sigma\nu}\right) W_\nu U_{1\nu} \ . \tag{129}$$

Define the quantity G_ν by

$$G_\nu = \sum_\mu H_{\mu\nu} = \lambda_1 \sum_\mu \hat{H}_{\mu\nu} \ . \tag{130}$$

Then

$$\alpha_1 = \frac{1}{N} \sum_\nu \left(\frac{1}{\lambda_1} - U^*_{1\nu}\right)\left(\frac{1}{\lambda_1} - \frac{G_\nu}{\lambda_1^2}\right) W_\nu U_{1\nu} \ .$$

If we now estimate α with the approximation

$$\alpha \approx \alpha_0 + \alpha_1 \ ,$$

we obtain

$$\alpha \approx \frac{1}{N} \sum_\nu \left(\frac{1}{\lambda_1} - U^*_{1\nu}\right)\left(\frac{1}{\lambda_1} + U^*_{1\nu} - \frac{1}{\lambda_1^2} G_\nu\right) W_\nu U_{1\nu} \ . \tag{131}$$

The approximation, Eq. (131), may be evaluated during a Monte Carlo run if the appropriately normalized importance function $U^*_{1\nu}$ is supplied. The other quantities appearing have the following interpretation: $U_{1\nu}/\lambda_1$ is the contribution of a unit source in region ν; W_ν is the weight of a site produced by the absorption of a unit-weight neutron in region ν; and G_ν is the weight of a site produced by a unit-weight neutron <u>born</u> in region ν.

Another approximation for the eigenvalue bias may be derived by noting that

$$\alpha_1 + \alpha_2 = \frac{1}{N^2} \sum_{k=2} \sum_{\ell=2} \sum_{n=0}^{\infty} \left(\underline{\hat{\tau}}^T \cdot \underline{U}_k\right)\left(\underline{\hat{\tau}}^T \cdot \underline{U}_\ell\right) \left(1 - \rho_\ell\right)\left(\rho_k \rho_\ell\right)^n \varepsilon_{k\ell}$$

$$= \frac{1}{N^2} \sum_{n=0}^{\infty} E\left\{\left(\underline{\hat{\tau}}^T \cdot [I - \underline{U}_1 \underline{U}_1^*] \hat{H}^n \underline{\varepsilon}^{i-n-1}\right)\right.$$

$$\left.\times \left(\underline{\hat{\tau}}^T \cdot [I - \underline{U}_1 \underline{U}_1^*][\hat{H}^n - \hat{H}^{n+1}] \ \underline{\varepsilon}^{i-n-1}\right)\right\} .$$

It follows that

$$\alpha_1 + \alpha_2 = \frac{1}{N^2} E\left\{\left[\hat{\underline{\tau}}^T \cdot \left(1 - \underline{U}_1\underline{U}_1^*\right) \tilde{\underline{\delta}}^i\right]\left[\hat{\underline{\tau}}^T \cdot \left(1 - \underline{U}_1\underline{U}_1\right)\left(\tilde{\underline{\delta}}^i - \tilde{\underline{\delta}}^{i+1}\right)\right]\right\} .$$

If, in the above equation, we make the approximation that

$$\tilde{\underline{\delta}}^i \approx \underline{\delta}^i \approx \underline{V}^i - \underline{V}_{ave} ,$$

then

$$\alpha_1 + \alpha_2 \approx \frac{1}{N^2} E\left\{\left[\left(\hat{\tau}^T \cdot \left[\underline{V}^i - \underline{V}_{ave}\right]\right) - \left(U_1^* \cdot \left[\underline{V}^i - \underline{V}_{ave}\right]\right)\right]\right.$$

$$\left. \times \left[\left(\hat{\tau}^T \cdot \left[\underline{V}^i - \underline{V}^{i+1}\right]\right) - \left(U_1^* \cdot \left[\underline{V}^i - \underline{V}^{i+1}\right]\right)\right]\right\} .$$

The approximation for $\alpha_1 + \alpha_2$ may be estimated from a Monte Carlo run if the importance functions are either supplied externally or generated by the run. When combined with α_0 from Eq. (128), an estimate of the eigenvalue bias will be obtained.

C. The Bias in the Eigenvector

The bias which occurs in the eigenvector is related to our previously defined quantities by

$$\Delta\underline{U} = \frac{1}{N} E\left\{\underline{\delta}^i\right\} , \tag{132}$$

where $\underline{\delta}^i$ is obtained from Eq. (116). The eigenvalue bias is related to $\Delta\underline{U}$ by

$$\alpha \equiv \frac{\Delta\lambda}{\lambda_1} = \hat{\underline{\tau}}^T \cdot \Delta\underline{U} .$$

By steps similar to those used in the derivation of Eq. (124), we obtain

$$\Delta\underline{U} = \frac{1}{N^2} \sum_{k=1} \sum_{\ell=2} \frac{\left(\hat{\underline{\tau}}^T \cdot \underline{U}_k\right) \varepsilon_{k\ell}}{1 - \rho_k\rho_\ell} \left[\left(\hat{\underline{\tau}}^T \cdot U_\ell\right) \underline{U}_1 - \rho_\ell \underline{U}_\ell\right] . \tag{133}$$

By examination of Eq. (133), two features may be noticed. First, contributions from modes may be present which do not contribute at all to the eigenvalue bias (i.e., those with $\hat{\underline{\tau}}^T \cdot \underline{U}_\ell = 0$). Second, modes with eigenvalues close to the fundamental, which contribute only slightly to the eigenvalue bias, are much

stronger contributors to the eigenvector bias. As a consequence, large local bias in source shape may be present even when eigenvalue bias is unimportant.

Some understanding of the potential difficulty may be gained by examining the leading term in Eq. (135) coming from the series expansion of $[1 - \rho_k \rho_\ell]^{-1}$:

$$\Delta \underline{U}^{(0)} = \frac{1}{N^2} \sum_{k=1} \sum_{\ell=2} \left(\hat{\tau}^T \cdot \underline{U}_k \right) \left[\left(\hat{\tau}^T \cdot \underline{U}_\ell \right) \underline{U}_1 - \rho_\ell \underline{U}_\ell \right] \varepsilon_{k\ell} \ .$$

Using the relationships of the previous section, we obtain

$$\Delta \underline{U}^{(0)} = H \frac{1}{N\lambda_1^2} \sum_\mu W_\mu U_{1\mu} \left[\frac{1}{\lambda_1} \underline{U}_1 - \sum_{\ell=1} U^*_{\mu\ell} \underline{U}_\ell \right] \ . \tag{134}$$

By using the assumed completeness of the eigenvectors, Eq. (134) may be reduced to the form

$$\Delta \underline{U}^{(0)} = H \underline{S} \ , \tag{135}$$

where

$$\underline{S} = \frac{1}{N\lambda_1^2} \left[\bar{W} - W \right] \underline{U}_1 \ . \tag{136}$$

In Eq. (136), W is the diagonal matrix with elements $W_\mu \delta_{\mu\nu}$ and

$$\bar{W} = \left(\frac{\sum_\mu W_\mu U_{1\mu}}{\sum_\mu U_{1\mu}} \right) I \ . \tag{137}$$

The lowest-order estimate for the eigenvector bias may thus be obtained from a source vector constructed by weighting the fundamental mode eigenvector by the deviation of the weight from its average over the fundamental mode. Applying the appropriate operator to $\underline{S}$ as defined in Eq. (136) will produce corresponding results for the bias in reaction rate estimates.

D. A Note on Reducing the Bias

It has been suggested by Lieberoth [1] that the bias may be reduced by modifying the Monte Carlo procedure. The analog of his suggestion which applies to our mathematical model may be described as follows.

(1) define a maximum weight $W_{max} \geq W_\nu$ for all ν;

(2) at each absorption, accept a source site with probability W_ν/W_{max}; and

(3) if the site is accepted, assign the weight W_{max}.

The effect of the procedure is to do a Monte Carlo calculation with only constant or zero weights allowed. A correlation matrix analogous to $\varepsilon_{\mu\nu}$ may be derived which is given by

$$\varepsilon^*_{\mu\nu} = N\left\{\delta_{\mu\nu}W_{max}U_{1\mu} - U_{1\mu}U_{1\nu}\right\}$$

and

$$\varepsilon^*_{k\ell} = N\left\{W_{max}\sum_{\mu}{}' U^*_{k\mu}U^*_{1\mu}U_{1\mu} - \delta_{k\ell}\delta_{\ell,1}\right\} . \tag{138}$$

The effect may be seen in Eq. (134) by the replacement $W_\nu \to W_{max}$, causing the lowest-order estimate of the eigenvector bias to vanish identically: $\Delta\underline{U}^{(0)} \to 0$ when $W_\nu \to W_{max}$. The increase in the magnitude of the correlation matrix, Eq. (138), will increase the magnitude of the nonvanishing terms, partly offsetting the gain obtained by eliminating the first term. Variances in calculated quantities will be increased as well.

If the presence of the bias is considered to be a potential difficulty, then the simplest and most direct approach is to try to use as large a number N of neutrons per generation as is feasible after sufficient source settling and to take advantage of the 1/N dependence of the bias.

Part III

MONTE CARLO CODE DEVELOPMENT AT ARGONNE NATIONAL LABORATORY, VIM DEVELOPMENT

Monte Carlo work at Argonne has, as its main goal, the construction of a Monte Carlo code for use in the analysis of Argonne critical experiments. The code presently available at Argonne is the VIM code originally writted by Levitt, et al, [4] at Atomics International. Responsibility for the maintenance of VIM, and for its future development, was transferred about a year ago from Atomics International to Argonne.

VIM is a continuous energy Monte Carlo code designed to treat a fast reactor lattice of rectangular subassemblies (or "drawers"). In the original VIM it was assumed that all interfaces within any drawer would be planar surfaces, and that each planar interface would be parallel to one of the three cartesian coordinate planes. Recently, however, combinatorial geometry routines written by MAGI [5] have been incorporated into VIM. The use of combinatorial geometry allows VIM calculations to be made in complex, irregular geometries as well as in lattice geometries. The combinatorial routines are being extended to the lattice environment to permit a full range of geometric options within repeating subassemblies. This capability is required for the

analysis of small-sample-worth experiments and for the analysis of pin-type fuel loadings, since small cylinders are used within the drawers.

The VIM cross-section library at present consists of 20 isotopes. Cross sections are specified pointwise from 10 eV to 10 MeV with different interpolation schemes possible in different energy ranges. ENDF/B data specifications are closely followed in the VIM library with respect to scattering angular distributions and individual isotopic fission spectra. Probability tables [6] are used for unresolved resonance cross sections. In the present library, all resonance data has been Doppler broadened to 300°K. Since the VIM cross-section library codes are now operational at Argonne, the number of isotopes in the library is being extended and it is now feasible to generate a multi-temperature library.

Among the modifications to VIM, which have been made at Argonne, has been the inclusion of variable dimensioning with dynamic allocation of storage. A small test problem may be executed in a total core requirement of about 60,000 words, with a realistic full-core eigenvalue problem requiring perhaps 400,000 words. The Monte Carlo calculation is completely core contained. The use of variable dimensioning has removed most limitations on problem complexity, with the exception of a temporary limit of 20 isotopes in a calculation; core size is the practical limitation on problem specification.

The computation of small sample worths is a primary goal in the VIM code development program. Such calculations are difficult simply because the samples are small. It is our intention, at present, to explore carrying out these calculations by a combination of three techniques, namely,

(1) roulette and splitting, to build up the density of particles around the sample;

(2) line-of-sight estimation, to make optimum use of collisions near the sample; and

(3) complete correlation.

By "complete correlation" we mean the simultaneous treatment of perturbed and unperturbed problems in a single random walk. The proposed computational method has the advantage that it is exact and that it does not require the simulation of the adjoint transport equation.

REFERENCES

1. J. LIEBEROTH, "A Monte Carlo Technique to Solve. The Static Eigenvalue Problem of the Boltzmann Transport Equation," *Nukleonik*, Bd. II, 213 (1968).

2. R. C. GAST, "Monte Carlo Eigenfunction Iteration Strategies That Are and Are Not Fair Game," WAPD-TM-878 (1969).

3. D. B. MAC MILLAN, "Monte Carlo Confidence Limits for Iterated-Source Calculations," *Nucl. Sci. Eng.*, 50, 73 (1973).

4. L. B. LEVITT and R. C. LEWIS, "VIM-1, A Nonmultigroup Monte Carlo Code for Analysis of Fast Critical Assemblies," AI-AEC-12951, Atomics International (1970).

5. M. O. COHEN et al, "SAM-CE: A Three-Dimensional Monte Carlo Code for the Solution of the Forward Neutron and Forward and Adjoint Gamma Ray Transport Equations," MR-7021, Mathematical Applications Group Inc. (1971).

6. L. B. LEVITT, "The Probabilty Table Method for Treating Unresolved Resonances in Monte Carlo Criticality Calculations," *Trans. Am. Nucl. Soc.*, 14, 648 (1974).

DISCUSSION

Gast: I think Don MacMillan recognized that his error estimates would be conservative for very small regions. His arguments are much more valid when applied to integrals over large subregions of the critical system than when they are applied to small regions. The modifications that we made at Bettis did recognize that, as the correlation coefficient of lag 1 approached zero, MacMillan's formula was overconservative and we did try to remedy this defect. We would be interested in the results of any tests you make on the Bettis scheme.

Gelbard: I would also like to look at a model problem that would have some leakage. We could put an anti-reflecting boundary on one side and then look at configurations of different sizes. At least the presence of leakage would tend to make the model problem a shade more realistic. The other possibility that seems to me to be interesting is to try to analyze the fluctuations that you observe in a practical Monte Carlo, to try to compute from those fluctuations, somehow, what the eigenvalues are. The question is, how reliably can you compute correlation coefficients of lag one, or lag two, or lag three, from a reasonable number of generations? If it is a hundred generations, that is one problem. If you can tolerate, let us say, going to a smaller number of histories per generation and then using 200 generations, it might be a lot easier. That is why I said, at the beginning of my talk, that there would be some advantage to being able to cut down the number of histories per generation. Incidentally, in our second paper, Prael has proposed ways to estimate the bias in the course of the running of a Monte Carlo, so that one might be able to tell in any given case whether it is permissible to use a relatively small number of histories per generation. If you can estimate the eigenvalues from the correlation coefficients of various lags you would get around the problem of having to use the dominance ratio, computed externally. In addition you might be able to get a much better fit than MacMillan gets by using only one eigenvalue, and taking that eigenvalue to be the dominance ratio. This is the sort of thing we would like to be able to do.

Gast: We at Bettis have leaned toward doing more generations with fewer histories because of the problem of estimating the correlation coefficient. But the extent to which we can do this is limited, because we operate in a manner similar to 05R. Between generations there is a good deal of what I call "dead time", used for reading nuclear data and so forth. To avoid building up this dead time, we have more or less decided to use no fewer than 500 histories per iteration. But if it were not for that problem we would like to go to a lower number.

Borgwaldt: I hope that, in the afternoon, I will have the occasion to say a little bit more about our ideas concerning eigenfunction strategies. At this point I have only one simple question. Your theoretical work is based on a simple matrix model and I think that this is a reasonable approach. But have you tried to verify, with a model Monte Carlo calculation, that the approximations which you have had to introduce, like linearization for instance, are justified? Have you done model calculations to test your approximations?

Gelbard: No, simply because we have not had the time to do it. But this is certainly something that we would like to do. We would like to do a calculation specifically for the model problem, first, and to see if the theory

makes sense. Then, if it does, we want to try to analyze realistic problems. But, we have not had time to do this yet.

Coveyou: I would just like to comment on what Bob Gast said. There is no reason at all why you cannot do an 05R calculation, let us say, with 400 neutrons per generation, and then analyze it as if it were ten separate calculations each with forty neutrons per generation.

Gelbard: I think that this suggestion is very similar to a proposal that Kalos made earlier, namely that you batch the neutrons in independent groups and try to look at their statistical behavior from group to group.

Coveyou: It seems to me that there are two sources of the errors and biases that you get. One is precisely the fact that you analyze histories in batches. That is apparent from your paper since, invariably, the estimates of the variances and covariances all had a factor of exactly N in them. In other words, the fact that you had N particles had absolutely no influence on your results.

Gelbard: It just lowered the fluctuations uniformly for all modes.

Coveyou: Really you are not doing a batch of N particles and averaging over them — you are doing N different experiments with one particle each.

Kalos: That is right.

Gelbard: You are not saying that we have omitted something in our treatment of the correlation between generations are you?

Coveyou: No, no, the point here is that, because you kept the number of particles fixed, it is clear that you are not really analyzing what happened to N particles in a generation. Well, in a way you are, but you are also equivalently analyzing N different experiments, each with one particle. And that is shown by the fact that you have invariably just one N-dependent factor. The only way your results depend upon N is that N divides the variances and covariances.

Gelbard: That is right. Now, do you think that is an unrealistic model?

Coveyou: No, but I am beginning to doubt that you get any advantage, from a statistical point of view, out of using a large number of particles in a generation. I would guess, just off the top of my head, that possibly if you analyze the progeny of one source neutron, and then get estimates from each source neutron, that your error may even be less.

Gelbard: You run into the bias problem.

Coveyou: No, I think that such a procedure removes a part of the bias problem. Part of the bias comes from this normalization.

Gelbard: Oh, I see, you are talking about not normalizing, you are talking again about the suggestion that one should treat each particle separately and make independent estimates from each history.

Coveyou: Right.

Gelbard: Yes, but I think you are coming to the same question as to whether there is an advantage to be gained by not normalizing each generation. That is really the question we were discussing earlier. I think it is a really crucial question; I think this idea of not normalizing is very similar to the idea of weighting each generation by the eigenvalue. That idea is, essentially, in Lieberoth's paper. Originally I thought it was a good idea, but then I found myself asking whether you really don't always have to normalize in one way or another before you <u>use</u> an eigenfunction.

DISCUSSION ON VIM CODE

Borgwaldt: Do you keep the data for the nuclear cross sections in storage and, if not, what is the relationship between your turnaround time and the computer time, the CPU time?

Prael: We keep them in-core at all times; and, turnaround is terrible anyway. Even the small job that runs only 15 minutes here does not get good turnaround time.

Grimstone: With this very large amount of data, is your main problem the storage and the resultant slow turnaround time, or do you feel you have any problem with the actual speed of the code itself in handling this very large amount of data?

Prael: I think we probably spend most of our running time sorting through all the cross-section data; and the running time is about linear in the number of isotopes in the problem.

Gast: I was just going to contrast your experience with what we have found at Bettis. I think that of all the run time in our Monte Carlo (with what we consider a very detailed cross-section representation), about only 40% typically is in the tracking. The remainder of the time is in handling data, reading in and out, etc. I think the question of the adequacy of the cross-section representation is a very difficult one when one is using Monte Carlo as a standard. The approach that we have frequently taken is, in effect, to base the cross-section structure on intuitive feelings as to what constitutes an adequate representation, and then to hold the cross sections fixed for perhaps several years. This is essential, we feel, if (for example), we are to be able to review and intercompare results of benchmark studies carried out at different times. Independently a cross-section evaluation group is constantly checking, through subsidiary calculation, the adequacy of the cross-section representation. If it is decided that more detail is required, then there has to be an overlap period in which old calculations are compared with new, and the process really becomes very complex. To avoid such complications, the cross sections and their representations in our Monte Carlo codes tend to be held fixed. For example, our current cross sections and library structures have been fixed for perhaps five years. We are, however, just now starting some subsidiary calculations to see if our current representation really need to be refined as the cross sections on ENDF/B themselves become more refined.

Prael: That is certainly the right way to go. At the present time we are taking our first close look at a detailed analysis of the cross sections that are available to us. I would hope that we will be able to fix our library for the duration of the use of ENDF/B-III very shortly; and, when we prepare for any changes which may be required by ENDF/B-IV, we hope to use our experience to help us generate any new sets we may need.

Cashwell: The one-fourth million of words was for how many elements?

Prael: About twenty. We are limited to twenty isotopes per problem.

Cashwell: That is about our limit. We were shooting for an even smaller number — about fifteen — but I think we can do better.

Prael: One change we did make in the VIM code, here, was to introduce the use of the Argonne BPOINTER package for dynamic allocation of storage. Except for the limit of twenty isotopes per problem, there are essentially no other fixed limits on problem specifications other than the size of the core.

Gelbard: Any other questions? I gather that our code is one of the few that does not do any packing. Is that correct? It seems to me that KAMCCO actually uses packing for cross-section data.

Borgwaldt: Yes.

Gelbard: And is that what the Los Alamos codes do?

Cashwell: Yes.

Borgwaldt: So you use packing, and the Bettis code tries to handle the cross-section problem by reading the cross sections in batches. Is that correct? In other words you process neutrons through an energy range, and then throw out the cross sections for that range?

Gast: Yes, but we use packing too.

Kalos: How much can you pack into an IBM/360 word?

Gast: On our CDC-7600 we pack two words into one.

Kalos: Yes, but you have a big word; the IBM/360 has a little word.

Borgwaldt: I mean by packing that we leave no gaps between data.

Kalos: Oh, you squeeze.

Borgwaldt: Yes, we squeeze. Perhaps I may make another remark. At our computer installation, when we use more than 480 K bytes, we get a very low priority. For this reason we have always been forced to use a data representation that does not use up too much storage.

Prael: One of the limitations that we are working under is to try to keep our code easily convertible to any other machine, and some of the features that one would commonly use in coding for IBM machines are not used. For instance, integer arrays of small magnitude are not written in half-words but in full-words, which is totally wasteful.

Kalos: Perhaps it is worth mentioning that the SAM code uses both packing (in the sense that you are using the term) and banding, or supergrouping of the cross sections. They are swapped in and out of core as needed, but they are not sparse, either.

Gelbard: I gather that the reason that we can get away without packing of this sort, or banding, is that the core of our IBM machine is very large.

Kalos: Your particular IBM machine.

Gelbard: We may eventually be forced to do something more drastic than that, but so far we have been very conservative. Dr. Borgwaldt, I think that in your talk you mentioned that you used an optimum-combined estimator?

Borgwaldt: Yes. We compute the variance and the covariance and, in the usual fashion, we find the optimum combination for the two estimates. But we have found that usually this combined estimation does not have a very great effect. We have only found a very few cases in which the optimally combined estimator has given us a real advantage.

Gelbard: I think that the method which you are using to get combined estimators is a fairly straightforward one compared to the one that has sometimes been used at Bettis. At Bettis and Knolls Atomic Power Laboratory the treatment of the optimum-combined estimator takes into account the fact that the combination coefficients are, themselves, random variables. The variance of this combined estimator was analyzed in a very nice piece of work, some time ago, by a statistician at KAPL.* To compute the variance of the combined estimator is a very difficult statistical problem.

Coveyou: In fact, are you sure that the estimates are unbiased? I have a feeling that they are not. If you could calculate the combination coefficients from the sample, then I think that you will find that you have the old problem. You are dividing by a random variable, and the mean value of the quotient is not the quotient of the mean values. I think that, unless something is specifically done about this, the estimates themselves are biased. If you use part of the sample to calculate the combination coefficients, and the other part of the sample uses these combination coefficients, then the method is unbiased. But, if you use the same sample to calculate the combination coefficient as you do to calculate the mean values, then your results will definitely be biased.

Kalos: If you don't gain much by making the optimum combination, then I think you might as well forget the whole thing.

Coveyou: I would think so, yes.

Prael: We ought to check the paper. I think the author claims to show that his estimator is not biased.

Gelbard: I would like to cut this short as I don't think we have enough information on this subject at this time, and it seems to me that it would be worthwhile to look at this paper very closely. But I would like to ask Bob Gast whether, in his experience, very much is gained through use of the combined estimator, assuming that it is unbiased.

Gast: We have found that we have not bought much by combining estimators. We do not do it in our program, and since we have some comparisons

---

**Editor's Note:* Max Halperin, "Almost Linearly-Optimum Combination of Unbiased Estimates," *Am. Statistical Assoc. J.*, 56, 36 (1961).

which indicate that they would not buy us much, we have never taken combined estimators very seriously.

Kalos: I am glad to hear it. I heard about combined estimators many years ago and great claims were made for them; so I would always feel guilty about not using them.

Borgwaldt: I might perhaps add that you cannot say, a priori, whether the track length estimator or the collision estimator is better. In some cases one is better, in some cases the other. In most cases we find that track length estimation is better. We normally use the combined estimator at least to take advantage of the best features of both. It does not cost much either — that is another point.

SESSION IV

FOCUS — A NON-MULTIGROUP ADJOINT MONTE CARLO CODE WITH IMPROVED VARIANCE REDUCTION

J. E. Hoogenboom

Delft Technological University
Delft, The Netherlands

FOCUS — A NON-MULTIGROUP ADJOINT MONTE CARLO CODE WITH IMPROVED VARIANCE REDUCTION

J. E. Hoogenboom

Delft Technological University
Delft, The Netherlands

ABSTRACT

A description is given of the selection mechanism in the adjoint Monte Carlo code FOCUS in which the energy is treated as a continuous variable. The method of Kalos who introduced the idea of adjoint cross sections is followed to derive a sampling scheme for the adjoint equation solved in FOCUS which is in most aspects analogous to the normal Monte Carlo game. The disadvantages of the use of these adjoint cross sections are removed to some extent by introduction of a new definition for the adjoint cross sections resulting in appreciable variance reduction. At the cost of introducing a weight factor slightly different from unity, the direction and energy are selected in a simple way without the need of two-dimensional probability tables.

Finally the handling of geometry and cross section in FOCUS is briefly discussed.

INTRODUCTION

The problem of deriving an adjoint equation which is analogous to the normal integral equation has been discussed by Eriksson et al and Kalos [1]. With some changes we folow essentially their derivation although other ways are possible [2].

First the normal integral equation is given and from that an adjoint equation is deduced. This equation, however, is not easy to handle in a Monte Carlo game. Therefore a transformation of this equation is necessary by which adjoint cross sections are introduced. Some properties of these adjoint cross sections, especially when dealing with inelastic scattering, are discussed in more detail.

THE INTEGRAL TRANSPORT EQUATION

We define the following quantities:

$\psi(\underline{r},E,\underline{\Omega})dVdEd\Omega$ = the number of neutrons entering a collision in the volume element dV at $\underline{r}$ with energies between E and E+dE and directions in the solid angle $d\Omega$ about $\underline{\Omega}$.

$\chi(\underline{r},E,\underline{\Omega})dVdEd\Omega$ = the number of neutrons leaving a collision or the source in the volume element dV at $\underline{r}$ with energies between E and E+dE and directions in the solid angle $d\Omega$ about $\underline{\Omega}$.

$S(\underline{r},E,\underline{\Omega})dVdEd\Omega$ = the number of source neutrons in the volume element dV at $\underline{r}$ with energies between E and E+dE and directions in the solid angle $d\Omega$ about $\underline{\Omega}$.

$T(\underline{r}'\to\underline{r},E,\underline{\Omega})dV$ = the probability for a neutron at $\underline{r}'$ with energy E and direction $\underline{\Omega}$ to have its next collision in the volume element dV at $\underline{r}$.

$P_s(\underline{r},E)$ = the non-absorption probability for neutrons with energy E which have a collision at $\underline{r}$.

$C_s(\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})dEd\Omega$ = the probability for a neutron with energy E' and direction $\underline{\Omega}'$ having a scattering collision at $\underline{r}$ to obtain an energy after collision between E and E+dE and a direction in the solid angle $d\Omega$ about $\underline{\Omega}$.

From these definitions we obtain the following equation between ψ and χ

$$\psi(\underline{r},E,\underline{\Omega}) = \int T(\underline{r}'\to\underline{r},E,\underline{\Omega})\chi(\underline{r}',E,\underline{\Omega})dV' \tag{1}$$

$$\chi(\underline{r},E,\underline{\Omega}) = \int C_s(\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})P_s(\underline{r},E')\psi(\underline{r},E',\underline{\Omega}')dE'd\Omega'+S(\underline{r},E,\underline{\Omega}) \tag{2}$$

Substitution of (1) in (2) leads to the integral equation for χ

$$\chi(\underline{r},E,\underline{\Omega}) = \int K(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})\chi(\underline{r}',E',\underline{\Omega}')dV'dE'd\Omega'+S(\underline{r},E,\underline{\Omega}) \tag{3}$$

with the kernel K given by

$$K(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega}) = T(\underline{r}'\to\underline{r},E',\underline{\Omega}')P_s(\underline{r},E')C_s(\underline{r},E'\to E,\underline{\Omega}'\to\Omega) \tag{4}$$

From the definitions of T, P_s and C_s we have

$$T(\underline{r}'\to\underline{r},E,\underline{\Omega})=\Sigma_t(\underline{r},E)\exp\{-\int_0^{|\underline{r}-\underline{r}'|}\Sigma_t(\underline{r}'+s\underline{\Omega},E)ds\}\delta\left(\Omega-\frac{\underline{r}-\underline{r}'}{|\underline{r}-\underline{r}'|}\right)/|\underline{r}-\underline{r}'|^2 \tag{5}$$

$$P_s(\underline{r},E) = \frac{\Sigma_s(\underline{r},E)}{\Sigma_t(\underline{r},E)} \tag{6}$$

$$C_s(\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega}) = \sum_A N_A(\underline{r}) \sum_j \sigma_{A,j}(E') f_{A,j}(E'\to E,\underline{\Omega}'\to\underline{\Omega})/\Sigma_s(\underline{r},E') \tag{7}$$

where

$\Sigma_t(\underline{r},E)$ is the total macroscopic cross section at $\underline{r}$ for neutrons with energy E

$\Sigma_s(\underline{r},E)$ the total scattering cross section

$N_A(\underline{r})$ the atom density for a nuclide of type A

$\sigma_{A,j}(E)$ the microscopic cross section for reaction type j of nuclide A

$f_{A,j}(E'\to E,\underline{\Omega}'\to\underline{\Omega})$ the scattering function for reaction type j of nuclide A.

With these definitions the kernels T and C_s are normalised probability density functions for $\underline{r}$ and $(E,\underline{\Omega})$ respectively and can be sampled in the familiar way. The non-absorption probability P_s has been used explicitly in the formulas because it is mostly taken into account as a weight factor.

Suppose the aim of our calculation is a quantity F that is some average over the collision density ψ, say

$$F = \int g(\underline{r},E,\underline{\Omega})\Psi(\underline{r},E,\underline{\Omega})dVdEd\Omega \tag{8}$$

with g some specified function. If we want to calculate for example the response of a point detector at $\underline{r}_d$ with an efficiency $\eta(E)$ with regard to the neutron flux we have

$$g(\underline{r},E,\underline{\Omega}) = \eta(E)\delta(\underline{r}-\underline{r}_d)/\Sigma_t(\underline{r}_d,E) \tag{9}$$

We can also write F as an average over the density of particles leaving collisions

$$F = \int h(\underline{r},E,\underline{\Omega})\chi(\underline{r},E,\underline{\Omega})dVdEd\Omega \tag{10}$$

From (1) and (2) we obtain the relation between h and g

$$h(\underline{r},E,\underline{\Omega}) = \int g(\underline{r}',E,\underline{\Omega})T(\underline{r}\to\underline{r}',E,\underline{\Omega})dV' \tag{11}$$

Monte Carlo calculations of F will break down if $g(\underline{r},E,\underline{\Omega})$ is only non-zero in a small part of the phase space, as in the case of a point detector. One of the possibilities to overcome such problems may be the use of adjoint Monte Carlo.

THE ADJOINT EQUATION

If we define a kernel $K^\dagger$ adjoint to the kernel K from (4) by

$$K^\dagger(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega}) = K(\underline{r}\to\underline{r}',E\to E',\underline{\Omega}\to\underline{\Omega}') \tag{12}$$

then the following integral equation determines a function $\chi^\dagger(\underline{r},E,\underline{\Omega})$

$$\chi^\dagger(\underline{r},E,\underline{\Omega}) = \int K^\dagger(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})\chi^\dagger(\underline{r}',E',\underline{\Omega}')dV'dE'd\Omega'+h(\underline{r},E,\underline{\Omega}) \tag{13}$$

with $h(\underline{r},E,\underline{\Omega})$ from (10). Multiplying (13) by χ and (3) by $\chi^\dagger$ and integrating over all variables we obtain the relation

$$\int h(\underline{r},E,\underline{\Omega})\chi(\underline{r},E,\underline{\Omega})dVdEd\Omega = \int S(\underline{r},E,\underline{\Omega})\chi^{\dagger}(\underline{r},E,\underline{\Omega})dVdEd\Omega \tag{14}$$

or

$$F = \int S(\underline{r},E,\underline{\Omega})\chi^{\dagger}(\underline{r},E,\underline{\Omega})dVdEd\Omega \tag{15}$$

So we can calculate F by sampling $\chi^{\dagger}(\underline{r},E,\underline{\Omega})$ from (13) and use the physical source $S(\underline{r},E,\underline{\Omega})$ as the scoring function. From (15) we see that $\chi^{\dagger}(\underline{r},E,\underline{\Omega})$ is the contribution to F of a source neutron at $(\underline{r},E,\underline{\Omega})$, so $\chi^{\dagger}$ may be interpreted as the importance of a neutron to F. It will be clear that calculating F by solving the adjoint equation (13) will only be successful if the physical source S is not a point source or only non-zero in a small volume.

Now our problem is to design a sampling scheme for the kernel $K^{\dagger}$ which can be written as

$$K^{\dagger}(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega}) = \Sigma_t(\underline{r}',E)\exp\{-\int_0^{|\underline{r}-\underline{r}'|}\Sigma_t(\underline{r}+s\underline{\Omega},E)ds\}\delta\left(\underline{\Omega}-\frac{\underline{r}'-\underline{r}}{|\underline{r}-\underline{r}'|}\right)/|\underline{r}'-\underline{r}|^2$$

$$\times\frac{\Sigma_s(r',E)}{\Sigma_t(\underline{r}',E)}\sum_A N_A(\underline{r}')\sum_j \sigma_{A,j}(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')/\Sigma_s(\underline{r}',E) \tag{16}$$

From the δ-function we see that

$$\underline{r}' = \underline{r} - |\underline{r}-\underline{r}'|\underline{\Omega} \tag{17}$$

By substitution of $s'=|\underline{r}-\underline{r}'|-s$ in the integral in (16) we can rewrite $K^{\dagger}$ as follows

$$K^{\dagger}(r'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega} = \Sigma_t(\underline{r},E)\exp\{-\int_0^{|\underline{r}-\underline{r}'|}\Sigma_t(\underline{r}'-s'\underline{\Omega},E)ds'\}\delta\left(\underline{\Omega}+\frac{\underline{r}-\underline{r}'}{|\underline{r}-\underline{r}'|}\right)/|\underline{r}-\underline{r}'|^2$$

$$\times\sum_A N_A(r')\sum_j \sigma_j(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')/\Sigma_t(\underline{r},E) \tag{18}$$

The first part of this expression is equal to $T(\underline{r}'\to\underline{r},E,-\underline{\Omega})$, so we can sample this part of $K^{\dagger}$ in the familiar way except that tracking is in the direction $-\underline{\Omega}$. However, E and $\underline{\Omega}$ must then be known, so they must follow from the second part of (18). But this is not a normalised probability density function (pdf) and besides it depends on $\underline{r}$ owing to the factor $\Sigma_t(\underline{r},E)$.

To remove this difficulty and to obtain a selection mechanism analogous to the normal Monte Carlo game, we apply the following transformation

$$\xi^{\dagger}(\underline{r},E,\underline{\Omega}) = \Sigma^{\dagger}(r,E)\chi^{\dagger}(\underline{r},E,\underline{\Omega}) \tag{19}$$

$$S^{\dagger}(\underline{r},E,\underline{\Omega}) = \Sigma^{\dagger}(\underline{r},E)h(\underline{r},E,\underline{\Omega}) \tag{20}$$

$$L^{\dagger}(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega}) = \Sigma^{\dagger}(\underline{r},E)K^{\dagger}(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})/\Sigma^{\dagger}(r',E') \tag{21}$$

with $\Sigma^{\dagger}(\underline{r},E)$ a function that we shall define later. Then (13) and (15) transform into

$$\xi^{\dagger}(\underline{r},E,\underline{\Omega})=\int L^{\dagger}(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})\xi^{\dagger}(\underline{r}',E',\underline{\Omega}')dV'dE'd\Omega'+S^{\dagger}(\underline{r},E,\underline{\Omega}) \tag{22}$$

$$F = \int\frac{S(r,E,\Omega)}{\Sigma^{\dagger}(\underline{r},E)}\xi^{\dagger}(\underline{r},E,\underline{\Omega})dVdEd\Omega \tag{23}$$

If we perform the same manipulations on $L^\dagger$ as we did on $K^\dagger$ we can rearrange $L^\dagger$ to

$$L^\dagger(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})=\Sigma_t(\underline{r},E)\exp\{-\int_0^{|\underline{r}-\underline{r}'|}\Sigma_t(\underline{r}'-s'\underline{\Omega},E)ds'\}\delta\left(\underline{\Omega}+\frac{\underline{r}-\underline{r}'}{|\underline{r}-\underline{r}'|}\right)/|\underline{r}-\underline{r}'|^2$$
$$\times\frac{\Sigma^\dagger(\underline{r},E)}{\Sigma_t(\underline{r},E)}\sum_A N_A(\underline{r}')\sum_j\sigma_{A,j}(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')/\Sigma^\dagger(\underline{r}',E') \tag{24}$$

We now define the adjoint kernels

$$T^\dagger(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})=\Sigma_t(\underline{r},E)\exp\{-\int_0^{|\underline{r}-\underline{r}'|}\Sigma_t(\underline{r}'-s\underline{\Omega},E)ds\}\delta\left(\underline{\Omega}+\frac{\underline{r}-\underline{r}'}{|\underline{r}-\underline{r}'|}\right)/|\underline{r}-\underline{r}'|^2 \tag{25}$$

$$=T(\underline{r}'\to\underline{r},E,-\underline{\Omega})$$

$$P_s^\dagger(\underline{r},E)=\frac{\Sigma^\dagger(\underline{r},E)}{\Sigma_t(\underline{r},E)} \tag{26}$$

$$C_s^\dagger(\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})=\sum_A N_A(\underline{r})\sum_j\sigma_{A,j}(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')/\Sigma^\dagger(\underline{r},E') \tag{27}$$

we can write $L^\dagger$ as

$$L^\dagger(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})=C_s^\dagger(\underline{r}',E'\to E,\underline{\Omega}'\to\underline{\Omega})T^\dagger(\underline{r}'\to\underline{r},E,\underline{\Omega})P_s^\dagger(\underline{r},E) \tag{28}$$

If we require $C_s^\dagger$ to be a normalised pdf for E and $\underline{\Omega}$ we must define the function $\Sigma^\dagger(\underline{r},E)$ as

$$\Sigma^\dagger(\underline{r},E)=\sum_A N_A(\underline{r})\sum_j\int\sigma_{A,j}(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')dEd\Omega \tag{29}$$

In order to satisfy this condition we can define a (macroscopic) adjoint cross section according to the definition of Kalos [1], that we shall call

definition I:
$$\sigma_{A,j}(E')=\int\sigma_{A,j}(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')dEd\Omega$$
$$=\int\sigma_{A,j}(E)f_{A,j}(E\to E')dE \tag{30}$$

with the last equality from the fact that $f(E\to E',\underline{\Omega}\to\underline{\Omega}')$ only depends on $\underline{\Omega}$ through $\underline{\Omega}.\underline{\Omega}'$. Handling the adjoint cross sections just like normal cross sections we have

$$\Sigma_{A,j}^\dagger(\underline{r},E')=N_A(\underline{r})\sigma_{A,j}^\dagger(E') \tag{31}$$

$$\Sigma_A^\dagger(\underline{r},E')=\sum_j\Sigma_{A,j}^\dagger(\underline{r},E') \tag{32}$$

$$\Sigma^\dagger(\underline{r},E')=\sum_A\Sigma_A^\dagger(\underline{r},E') \tag{33}$$

With these definitions we can now sample $C_s^\dagger$ by first selecting a nuclide and reaction type from the pdf

$$p(A,j)=\frac{N_A(\underline{r}')\sigma^{\dagger}_{A,j}(E')}{\Sigma^{\dagger}(\underline{r}',E')}=\frac{\Sigma^{\dagger}_A(\underline{r}',E')}{\Sigma^{\dagger}(\underline{r}',E')}\,\frac{\sigma^{\dagger}_{A,j}(E')}{\sigma^{\dagger}_A(E')} \tag{34}$$

that is we select a nuclide according to the microscopic adjoint cross section for that nuclide and a reaction type according to the partial microscopic adjoint cross section for that reaction. Then results the problem of selecting energy and direction from the pdf

$$p(E,\underline{\Omega})=\sigma_{A,j}(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')/\sigma^{\dagger}_{A,j}(E') \tag{35}$$

To sample the kernel $L^{\dagger}$ completely we use $P^{\dagger}_s$ as a weight factor just as P_s in the normal game and select $\underline{r}$ from $T^{\dagger}$ by tracking in the direction $-\underline{\Omega}$.

Some words may be said about sampling the function $S^{\dagger}(\underline{r},E,\underline{\Omega})$ that appears as a "source" in the adjoint problem. For the general case of a "detector" with volume V_d and "efficiency" $\eta(\underline{r},E)$ with regard to the neutron flux, we have

$$g(\underline{r},E,\underline{\Omega})=\eta(\underline{r},E)\rho(\underline{r})/\Sigma_t(\underline{r},E) \tag{36}$$

with

$$\rho(\underline{r})=\begin{cases}1 \text{ if } \underline{r}\in V_d\\ 0 \text{ else}\end{cases} \tag{37}$$

The factor $\Sigma_t(\underline{r},E)$ in η appears because η is defined with regard to the flux and the function g is averaged over the collision density. If we want to calculate for example some reaction rate in a volume V_d of a system we take η equal to the macroscopic cross section for that reaction. If we want to calculate the flux at a point for a given energy range, we set η to unity for energies within that range.

From (11) we obtain for h

$$h(\underline{r},E,\underline{\Omega})=\int\eta(\underline{r}',E)\rho(\underline{r}')\,\frac{T(\underline{r}\to\underline{r}',E,\Omega)}{\Sigma_t(\underline{r}',E)}dV' \tag{38}$$

and for $S^{\dagger}$

$$\begin{aligned}S^{\dagger}(\underline{r},E,\underline{\Omega})&=\Sigma^{\dagger}(\underline{r},E)h(\underline{r},E,\underline{\Omega})=\Sigma^{\dagger}(\underline{r},E)\int\eta(\underline{r}',E)\ \rho(r')\,\frac{T^{\dagger}(r'\to r,E,\Omega)}{\Sigma_t(\underline{r},E)}dV'\\&=P^{\dagger}_s(\underline{r},E)\int_{V_d}\eta(\underline{r}',E)T^{\dagger}(\underline{r}'\to\underline{r},E,\underline{\Omega})dV'\end{aligned} \tag{39}$$

If we select successively $\underline{r}'$ uniform in V_d, $\underline{\Omega}$ isotropically, E from the pdf $\eta(\underline{r}',E)/\int\eta(\underline{r}',E)dE$, then $\underline{r}$ from $T^{\dagger}(\underline{r}'\to\underline{r},E,\underline{\Omega})$ and finally apply a weight factor $P^{\dagger}_s(\underline{r},E)$ then we have selected $(\underline{r},E,\underline{\Omega})$ from the pdf

$$p(\underline{r},E,\underline{\Omega})=\frac{1}{4\pi V_d}\,P^{\dagger}_s(\underline{r},E)\int\frac{\eta(r',E)}{\int\eta(\underline{r}',E)dE}\,T^{\dagger}(\underline{r}'\to\underline{r},E\ \underline{\Omega})dV' \tag{40}$$

So if we want to sample $S^{\dagger}$ we must introduce an additional weight factor

$$w=V_d\int\eta(\underline{r}',E)dE/\iint_{V_d}\eta(\underline{r}',E)dV'dE \tag{41}$$

and multiply the final result of our Monte Carlo calculation by $4\pi\iint\eta(\underline{r}',E)dV'dE$. Note that if η is independent of space the additional weight factor is unity.

ANALYSIS OF THE ADJOINT CROSS SECTIONS

To get a better insight in the properties of the adjoint cross sections we take the case of a constant elastic scattering cross section σ_{el} and assume isotropic scattering in the center of mass system. We then have

$$f_{el}(E\to E')=\frac{1}{E(1-\alpha)} \qquad \alpha E<E'<E \tag{42}$$

with

$$\alpha=\left(\frac{A-1}{A+1}\right)^2 \tag{43}$$

with A the mass of the nuclide. Then

$$\sigma_{el}(E')=\int_{E'}^{E'/\alpha}\sigma_{el}\frac{1}{E(1-\alpha)}dE=\sigma_{el}\frac{\log 1/\alpha}{1-\alpha} \tag{44}$$

The factor $\log(1/\alpha)/(1-\alpha)$ is always greater than unity but tends to unity for heavy masses. However, for the very light nuclides this factor becomes large and even infinite in the case of hydrogen. It is therefore necessary to introduce an upper limit in the integration over energy in (30) equal to the highest energy in the problem at hand. This removes difficulties with an infinite adjoint cross section but makes the adjoint cross sections problem dependent. If we compute the adjoint cross section for inelastic scattering with a threshold energy ε, still assuming a constant cross section and isotropic scattering, we obtain fig.1. The scattering function $f_{in}(E\to E')$ and the integration limits are derived later. The adjoint inelastic cross section for this case is constant for $E>\varepsilon/(A+1)^2$ and equal to

$$\sigma_{in}(E)=\sigma_{in}\frac{\log 1/\alpha}{1-\alpha} \qquad E>\varepsilon/(A+1)^2 \tag{45}$$

and falls off for lower energies.

From this figure we see that the adjoint cross sections will generally be greater than the corresponding normal cross sections. In systems with small absorption the adjoint scattering factor $P_s^\dagger=\Sigma^\dagger/\Sigma_t$ will therefore in general be greater than unity. Because the weight of the "adjoint particle" is multiplied by this factor at every event, the weight of the particle may become very large, which may result in large variancies. This is especially the case when both elastic and inelastic scattering is important. Suppose we have a system with only one nuclide that does not absorb neutrons and scatters them elastically or inelastically by one level with threshold energy ε. Suppose further that the cross sections are constant and that the elastic cross section is twice the inelastic cross section. We can then calculate the factor P_s for this medium. Fig. 2 shows the result, where we have taken a threshold energy of 1.037 MeV and a mass A=27 for the nuclide.

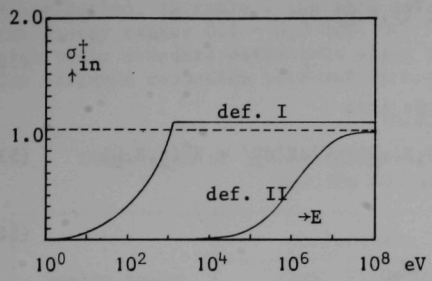

Fig.1. Adjoint inelastic cross section as a function of energy.

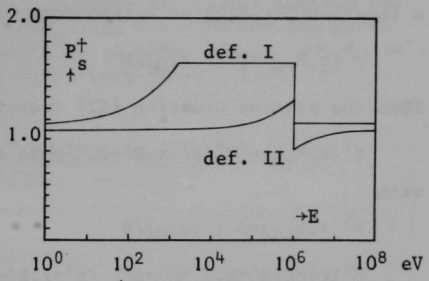

Fig.2. $P^{\dagger}_{s}$ as a function of energy.

We see that $P^{\dagger}_{s}$ is much greater than unity for a large energy range. If the nuclide had a very large mass, the factor $P^{\dagger}_{s}$ would have been 1.5 for $E<\varepsilon$ and unity for $E>\varepsilon$. So in the case of both elastic and inelastic scattering the factor $P^{\dagger}_{s}$ may even be much larger than in the case of elastic scattering only.

ANOTHER DEFINITION FOR THE ADJOINT CROSS SECTIONS

Up to now we always used the energy as a variable, but we could have equally well used the lethargy defined by

$$u=\log \frac{E_o}{E} \tag{46}$$

with E_o some reference energy. If we repeat the whole derivation we would have arrived at the following definition for the adjoint cross section

$$\sigma^{\dagger}_{A,j}(u')=\int\sigma_{A,j}(u)f_{A,j}(u\to u')du \tag{47}$$

On physical grounds we have

$$f_{A,j}(u\to u')\,|du'|=f_{A,j}(E\to E')\,|dE'| \tag{48}$$

But in (47) we integrate over u instead of u', so (47) is a different definition of the adjoint cross section. If we transform (47) back to energies with

$$f_{A,j}(u\to u')=f_{A,j}(E\to E')\left|\frac{dE'}{du'}\right|=f_{A,j}(E\to E').E' \tag{49}$$

and

$$du=\frac{dE}{E} \tag{50}$$

we arrive at our second definition for the adjoint cross sections

$$\text{definition II:}\qquad \sigma^{\dagger}_{A,j}(E')=\int\sigma_{A,j}(E)f_{A,j}(E\to E')\,\frac{E'}{E}\,dE \tag{51}$$

Instead of introducing the lethargy as a variable we could have introduced

a function $\zeta^\dagger(\underline{r},E,\underline{\Omega})$ as a transform of $\xi^\dagger(\underline{r},E,\underline{\Omega})$ by

$$\zeta^\dagger(\underline{r},E,\underline{\Omega})=\xi^\dagger(\underline{r},E,\underline{\Omega})/E \tag{52}$$

Then the adjoint equation (22) transforms into

$$\zeta^\dagger(\underline{r},E,\underline{\Omega})=\int M^\dagger(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})\zeta^\dagger(\underline{r}',E',\underline{\Omega}')dV'dE'd\underline{\Omega}' + R^\dagger(\underline{r},E,\underline{\Omega}) \tag{53}$$

with

$$R^\dagger(\underline{r},E,\underline{\Omega})=S^\dagger(\underline{r},E,\underline{\Omega})/E \tag{54}$$

$$\begin{aligned} M^\dagger(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})&=E'L^\dagger(\underline{r}'\to\underline{r},E'\to E,\underline{\Omega}'\to\underline{\Omega})/E \\ &=D^\dagger(\underline{r}',E'\to E,\underline{\Omega}'\to\underline{\Omega})T^\dagger(\underline{r}'\to\underline{r},E,\underline{\Omega})P_s^\dagger(\underline{r},E) \end{aligned} \tag{55}$$

$$D^\dagger(\underline{r}',E'\to E,\underline{\Omega}'\to\underline{\Omega})=\sum_A N_A(\underline{r}')\sum_j \sigma_{A,j}(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')(E'/E)/\Sigma^\dagger(\underline{r}',E') \tag{56}$$

For F we get from (23)

$$F=\int \frac{ES(\underline{r},E,\underline{\Omega})}{\Sigma^\dagger(\underline{r},E)}\ \zeta^\dagger(\underline{r},E,\underline{\Omega})dVdEd\Omega \tag{57}$$

When sampling $M^\dagger$ the only change compared with sampling $L^\dagger$ is that after selection of the nuclide and reaction type just as before, we have to select E and $\underline{\Omega}$ from the pdf

$$p(E,\underline{\Omega})=\sigma_{A,j}(E)f_{A,j}(E\to E',\underline{\Omega}\to\underline{\Omega}')\ (E'/E)/\sigma^\dagger_{A,j}(E') \tag{58}$$

which problem we shall discuss in the next section.

If we look again at the adjoint cross sections, now according to definition II, we get for the case of a constant elastic cross section

$$\sigma^\dagger_{el}(E')=\int_{E'}^{E'/\alpha}\sigma_{el}\frac{1}{E(1-\alpha)}\frac{E'}{E}dE=\sigma_{el} \tag{59}$$

The result for a constant inelastic cross section with threshold energy ε is also given in Fig. 1. Due to the factor E'/E in definition II for the adjoint cross section, it resembles now much more the normal cross section. The energy dependence appeared to be almost independent of the mass of the nuclide.

The weight factor $P_s^\dagger$ for the medium defined before is shown in Fig 2 when using definition II. From this figure we see not only that the maximum value for $P_s^\dagger$ is less than for definition I, but also that the energy range in which $P_s^\dagger$ is much greater than unity is rather small. Besides a value of $P_s^\dagger$ greater than unity at energies below the threshold will be compensated to some extent by a value of $P_s^\dagger$ less than unity at energies above the threshold.

To have an idea of the variance reduction obtained when using definition II, we calculated the flux at the center of a sphere containing the artificial nuclide mentioned before with mass A = 27. The threshold energy ε was again chosen at 1.037 MeV. The radius of the sphere was 5 mean free paths at energies below the threshold. A uniform and isotropic artificial neutron source was assumed in the sphere with a fission spectrum as the energy

distribution. In table I the flux per unit lethargy per source neutron for the energy ranges 0.1 - 0.2 MeV, 0.4 - 0.8 MeV and 1.4 -2.5 MeV are given with their standard deviations after processing of 1000 "adjoint particles". The variance reduction obtained by using definition II is clear.

TABLE I

Comparison of Standard Deviations Using Definitions I and II for the Adjoint Cross Sections

energy range (MeV)	definition I		definition II	
	flux per unit lethargy	standard deviation (%)	flux per unit lethargy	standard deviation (%)
0.1 - 0.2	$5.55\ 10^{-3}$	32	$6.89\ 10^{-3}$	6.2
0.4 - 0.8	$9.15\ 10^{-3}$	8.6	$9.26\ 10^{-3}$	2.6
1.4 - 2.5	$2.79\ 10^{-3}$	1.6	$2.79\ 10^{-3}$	0.8

SAMPLING THE ADJOINT COLLISION KERNEL

Now we have to look in more detail to the selection of the energy and direction from (58). We limit ourselves here to the case of inelastic level scattering. Elastic scattering is then automatically included by setting the level energy to zero. If the inelastic level has an energy Q above the ground state, the neutron must have at least the threshold energy

$$\varepsilon = \frac{A+1}{A} Q \tag{60}$$

to excite the level. From elementary collision mechanics one can derive the relation between the energy before and after the collision and the scattering angle. Throughout this section we shall adopt E' for the higher energy (that is the energy before collision in the normal case) and E for the lower energy We then have

$$\frac{E'}{E} = \frac{A^2(1-\varepsilon/E) + 2A\mu_C\sqrt{1-\varepsilon/E} + 1}{(A+1)^2} \tag{61}$$

$$\mu_o = \frac{A\mu_C\sqrt{1-\varepsilon/E} + 1}{\sqrt{A^2(1-\varepsilon/E) + 2A\mu_C\sqrt{1-\varepsilon/E} + 1}} \tag{62}$$

with μ_C and μ_o the cosines of the scattering angle in the center of mass system and the laboratory system respectively. For $\varepsilon=0$ we have the familiar expressions for elastic scattering.

Now we can write for the scattering function $f_{in}(E\to E', \underline{\Omega}\to\underline{\Omega}')$

$$f_{in}(E\to E', \underline{\Omega}\to\underline{\Omega}') = f_{in}(\underline{\Omega}\cdot\underline{\Omega})\,\delta\left(E' - \frac{A^2(1-\varepsilon/E) + 2A\mu_C\sqrt{1-\varepsilon/E} + 1}{(A+1)^2}\right) \tag{63}$$

We expand $f_{in}(\underline{\Omega}.\underline{\Omega}')=f_{in}(\mu_o)$ in a series of Legendre polynomials

$$f_{in}(\mu_o)= \frac{1}{2\pi}\sum_{1}\frac{21+1}{2}\, f_1(E)P_1(\mu_o) \tag{64}$$

with $f_1(E)$ denoting the 1-th Legendre component in the laboratory system. The factor $1/2\pi$ appears because the scattering will be azimuthally symmetric. If the angular dependence is needed in the center of mass system we can use

$$f^C_{in}(\mu_C)= \frac{1}{2\pi}\sum_{1}\frac{21+1}{2}\, f^C_1(E)P_1(\mu_C) \tag{65}$$

$f_{in}(\mu_o)$ and $f_{in}(\mu_C)$ are connected through the relation

$$f^C_{in}(\mu_C)=f_{in}(\mu_o)\,\frac{\partial\mu_o(\mu_C,E)}{\partial\mu_C} \tag{66}$$

By integrating (63) over $\underline{\Omega}$ we get the energy scattering function

$$\begin{aligned} f_{in}(E\to E')&=2\pi\int f_{in}(E\to E',\mu_o)d\mu_o=2\pi\int f^C_{in}(E\to E',\mu_C)d\mu_C \\ &= \frac{1}{E(1-\alpha)\sqrt{1-\varepsilon/E}}\sum_{1}(21+1)f^C_1(E)P_1(\mu_C) \end{aligned} \tag{67}$$

This function has to be used in the computation of the adjoint cross section as defined by (51). To determine the range of E and E' for which $f_{in}(E\to E')$ is non-zero we first examine the relation between μ_o and μ_C as given by (62) and shown graphically in fig. 3. If we set $A'=A\sqrt{1-\varepsilon/E}$ in (62) we have the same formula as for elastic scattering with a nuclide of mass A'. So for $E>\varepsilon A^2/(A^2-1)$ the functional dependence is the same as for elastic scattering and for $E=\varepsilon A^2/(A^2-1)$ it is the same as for elastic scattering on hydrogen. For $\varepsilon<E<\varepsilon A^2/(A^2-1)$ we see that the value of μ_o is restricted.

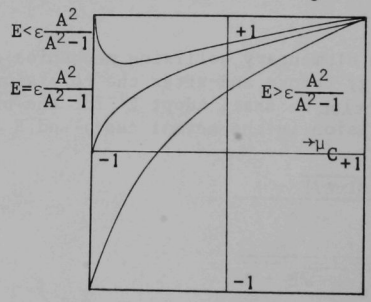

Fig. 3. μ_o as a function of μ_C for inelastic scattering.

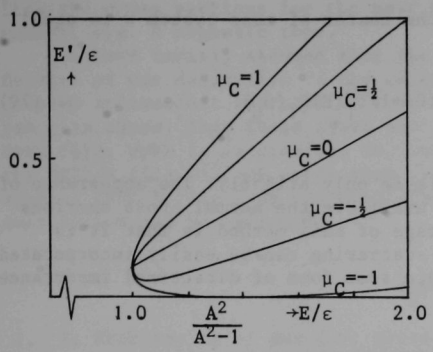

Fig. 4. Relation between E' and E with μ_C as a parameter.

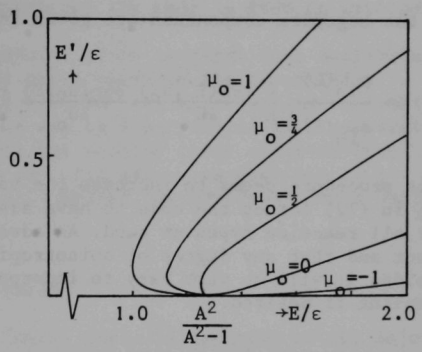

Fig. 5. Relation between E' and E with μ_o as a parameter.

In fig. 4 the relation between E' and E from (61) is shown with μ_C as a parameter. For $E>\varepsilon$ all values for μ_C between -1 and +1 are possible. For given E' as in the adjoint case this is also true for $E'>\varepsilon/(A+1)^2$ but not for $E'<\varepsilon/(A+1)^2$. However, if we use μ_o as a parameter we get fig. 5. Now for every E' all values for μ_o between -1 and +1 are allowed. So in the adjoint game it is easier to use μ_o. From (61) and (62) we can find E as a function of E' and μ_o.

$$E = \frac{(A+1)^2(A^2+2\mu_o^2-1)E'+A^2(A^2-1)\varepsilon-2\mu_o(A+1)\sqrt{(A+1)^2(A^2+\mu_o^2-1)E'^2+A^2(A^2-1)\varepsilon E'}}{(A^2-1)^2} \tag{68}$$

By setting $\mu_o=-1$ and +1 we obtain the minimum and maximum value for E given E'

The pdf for E irrespective of $\underline{\Omega}$ can be written as

$$p(E) = \frac{\sigma_{in}(E)f_{in}(E\to E')E'/E}{\sigma_{in}^{+}(E')} = \frac{\sigma_{in}(E)}{\sigma_{in}^{+}(E')}\frac{E'}{E^2(1-\alpha)\sqrt{1-\varepsilon/E}}\sum_l(2l+1)f_l^C(E)P_l[\mu_C(E,E')] \tag{69}$$

By numerical integration it is possible to construct in advance of the Monte Carlo calculation probability tables from which E can be selected by interpolation. This is the method suggested by Kalos and used in the ANTE code [3]. This requires the storing of two dimensional tables. In FOCUS we followed an other method. We select $\underline{\Omega}$ isotropically and apply a weight factor to account for not using the correct pdf. Than E is calculated from (68). By integrating (58) over E we obtain the true pdf for $\underline{\Omega}$

$$p(\underline{\Omega}) = \frac{1}{2\pi}\frac{\sigma_{in}(E)}{\sigma_{in}^{+}(E')}\frac{E'}{E}\frac{\partial E(E',\mu_o)}{\partial E'}\sum_l\frac{2l+1}{2}f_l(E)P_l(\mu_o) \tag{70}$$

where E is now determined by means of (68). The weight factor must be

$$w = \frac{\sigma_{in}(E)}{\sigma_{in}^{+}(E')}\frac{E'}{E}\frac{\partial E(E',\mu_o)}{\partial E'}\sum_l(2l+1)f_l(E)P_l(\mu_o) \tag{71}$$

If the Legendre components are given in the center of mass system w is equal to

$$w=\frac{\sigma_{in}(E)}{\sigma^{\dagger}_{in}(E')}\,\frac{E'}{E}\,\frac{\partial E(E',\mu_o)}{\partial E'}\,\frac{\partial\mu_C(\mu_o,E)}{\partial\mu_o}\sum_{l}(2l+1)f^{C}_{l}(E)P_l(\mu_C) \tag{72}$$

This procedure seems to increase the variance only slightly. The appearance of σ_{in} in (72) forces the code to have also available the normal cross sections for all reaction types at hand. An advantage of this method is that it is exact and that any degree of anisotropic scattering can be easily incorporated. Besides it will be very easy to incorporate some form of directional importance sampling if desired.

GEOMETRY AND CROSS SECTION HANDLING IN FOCUS

Geometry handling

Because the adjoint transport kernel $T^{\dagger}$ is the same as the normal transport kernel except for the flight direction, the geometry treatment may be same as in normal Monte Carlo codes. FOCUS uses the well known general geometry package from the O5R [4] and MORSE [5] codes. This geometry package allows any quadratic surface in the system and offers the possibility of detailed geometry for only a restricted part of the system without the need of testing all surfaces in the system for crossing.

Cross section handling

For sampling the adjoint transport kernel $T^{\dagger}$ we need the total cross section Σ_t for every medium in the system. Besides we need the weight factor $P_s=\Sigma^{\dagger}/\Sigma_t$ for every medium. To select a nuclide and a reaction type when sampling the adjoint scattering kernel $D^{\dagger}$ there are two ways. We can select directly a reaction type of a nuclide from the probabilities $\Sigma^{\dagger}_{A,j}/\Sigma^{\dagger}$ or select first the nuclide from the probabilities $\Sigma^{\dagger}_{A}/\Sigma^{\dagger}$ and then select for this nuclide the reaction type from the probabilities $\sigma^{\dagger}_{A,j}/\sigma^{\dagger}_{A}$. In FOCUS we have chosen the second method mainly because it saves computer storage when a nuclide with many inelastic levels appear in the more than one medium in the system.

Because of the applied section technique for the direction and energy, the partial cross sections for every nuclide are also needed together with the Legendre coefficients, if any. So the quantities $\Sigma_t,\Sigma^{\dagger}/\Sigma_t$ and $\Sigma^{\dagger}_{A}/\Sigma^{\dagger}$ are grouped by medium and the quantities $\sigma^{\dagger}_{A,j}/\sigma^{\dagger}_{A},\sigma_{A,j}$ and f^{C}_{l} are grouped by nuclide. All these cross sections are tabulated by energy. Because great detail in energy may be desired the cross sections can not be stored for all energies in computer memory. Therefore the energy range is devided in groups with energy boundaries that differ by a factor of 2 (called supergroups in the O5R code). Each supergroup is devided in a variable number of subgroups. Only the cross sections for one supergroup are held in computer memory. So if a particle in the adjoint game has received an energy above the upper boundary of the supergroup for which cross sections are held in the memory, its parameters are stored untill all particles with energy in the current supergroup are processed

Then the cross sections for the next supergroup are read in from an external memory, e.g. a magnetic tape.

We have tacitly assumed that the adjoint cross sections were available. Because of our definition II the adjoint cross sections are essentially problem independent and need be computed once for all nuclides and scattering reaction types. Then these cross sections may be incorporated in a existent data file. Work is in progress to compute the adjoint cross sections from the ENDF/B file [6] and store them with the original cross sections in a special ENDF/B file from which all the data needed for an adjoint Monte Carlo calculation can be selected.

REFERENCES

1. B. Eriksson et al and M.H. Kalos, "Monte Carlo Integration of the Adjoint Neutron Transport Equation", *Nucl. Sci. Eng.*, 37, 410 (1969).
2. W. Matthes, "Monte Carlo Simulation of the Adjoint Transport Equation", EUR-4678, Euratom (1971).
3. "ANTE-2, Adjoint Monte Carlo time-dependent Neutron transport code in Combinatorial Geometry", DASA-2396.
4. R.R. Coveyou, J.G. Sullivan and H.P. Carter, "05R - A General-purpose Monte Carlo Neutron Transport Code", ORNL-3622, Oak Ridge National Laboratory (1965).
5. E.A. Straker et al, "The MORSE Code - A Multigroup Neutron and Gammaray Monte Carlo Transport Code", ORNL-4585, Oak Ridge National Laboratory (1970).
6. M.K. Drake (ed.), "Data Formats and Procedures for the ENDF Neutron Cross Section Library", BNL-50274, Brookhaven National Laboratory (1970).

DISCUSSION

Kalos: I am gratified to find that we can gain so much by so simple a device. It is really very very nice/ but I would also like to call attention to the fact that the 1/E function, which seems to be a very good universal importance function, could be replaced by other functions. This paper then provides us with a model which shows how one must define reaction cross sections in order to carry out importance sampling in full detail, and it is not a trivial matter to do so. But, apparently just the 1/E factor makes an enormous difference, and I hope that we can get it into our code as soon as possible.

Hoogenboom: Yes, of course, there are other possibilities. One is led to this scheme in particular if, right from the start, one works consistently with lethargies instead of energies in the derivation. Incidentally it seems rather natural that, if the elastic scattering cross section is constant, the adjoint cross section should also be constant, and the two cross sections should be equal. That this should be true does not help much from a practical point of view, but it seems to be natural.

Borgwaldt: Can you recall the name of the code in which you have implemented your scheme? And can you say whether you have any trouble defining these adjoint cross sections in the statistical resonance region?

Hoogenboom: The code is called FOCUS, because we have tried to use it to compute fluxes at a point.

Borgwaldt: I thought you were talking about an NT code.

Hoogenboom: Yes, that was the code developed, as I understood, by Mal Kalos at New York; and it was in that code that the adjoint cross sections were first defined.

Borgwaldt: How is that spelled?

Kalos: ANTE.

Hoogenboom: That is the code which uses your original proposal?

Kalos: Which uses the original definition of the adjoint cross section.

Gelbard: Other comments? I might point out that the use of a transformation specifically chosen to eliminate weights is implemented in the MARC code, a code for use in thermal groups only. It turns out that if you assume that, in the thermal group, water is the only moderator (and this is assumed in MARC) then you can pick a transformation function which eliminates the need for weights completely. This sort of a choice of a transformation becomes impossible in general, but it is conceivable that you could come up with a transformation function that would be helpful. There were some attempts of this sort made by Spanier and Levitt but, I think, they were not really totally successful. But to come back to the method that you are using, it has always seemed to me that the most likely source of trouble in this kind of method would be resonances. If you are right below a scattering resonance, the integral, over higher energies, over the scattering resonance, will give you a very large weight.

Hoogenboom: Yes.

Gelbard: This would be true in the resolved range and would be true at higher energies if you are dealing with, let us say, resonances of a scatterer like iron. Do you find this to be a problem?

Hoogenboom: Well, this is correct, but we have not done enough computations so that I can give you a complete answer. I have calculated fluxes in a system that was composed of uranium with aluminum as a construction material and have not seen any difficulties.

Kalos: What about the narrow resonance range? Very narrow resonances?

Hoogenboom: The weight factor fluctuates in the neighborhood of a resonance, being higher than one at the bottom of the resonance and lower than one near the top. To some extent the high weight factors are compensated by the low weight factors. Not compensated fully, but to some extent.

Gelbard: Well, not compensated insofar as the variance is concerned.

Kalos: It depends on the problem. If the neutron makes a few scattering collisions within the resonance then the higher weights will be compensated by the lower weights. That is, if the resonance is wide enough so that neutrons travel through the resonances slowly. In all cases the right thing to do is to introduce the local energy-dependent adjoint. And that should not be difficult to do in specific cases, especially in the narrow resonance range, although I have not worked it out.

Gelbard: Well, that is really what was done in MARC where the use of the adjoint was straightforward. But do you mean, by the adjoint, the importance function that you would have in an infinite medium of the material you are in?

Kalos: No, I don't mean that.

Gelbard: You mean a space-dependent adjoint?

Kalos: It would have to be somewhat space dependent ...

Gelbard: Yes, then you have a problem with the computation and use of the adjoint. Incidentally in looking at the efficiency of the adjoint mode in MARC we used a device which we found to be helpful. Through this device one can estimate how much inherent loss in efficiency is introduced through the introduction of weights in the adjoint. Let us suppose that you have a slab of a homogeneous material. Divide the slab in half. Let us say the problem is to compute the resonance absorption in region 2 due to a source in region 1. You can do this in the direct mode, ahd you can do the same thing in the adjoint mode. From a geometric point of view the adjoint and direct mode calculations are identical. In one case you put the source in region 1 and get the absorptions in region 2 and in the other case you put the adjoint source in region 2 and tally in region 1. Because you have no geometric difference between the two problems you will generally find the variance in the adjoint calculation is greater. You have, then, some estimate of how much efficiency you are losing. In other words, you have an estimate of where the break-even point would be between the efficiency of the adjoint and the efficiency of the

direct mode. This is a sort of figure of the merit for whatever scheme you are trying to use. One more comment. It seems to be true that the adjoint is easiest to use in fixed source problems. In the eigenvalue problem, some sort of additional technique is required. In the eigenvalue problem the use of the adjoint becomes a lot harder. Suppose that I want the eigenvalue perturbation due to, let us say, a small sample or to some small perturbation that I make in a region. I hope that we will get to discuss how you might use the adjoint in this case, because here there are additional problems beyond those which are encountered in the fixed source case. I don't mean to imply that you cannot use the adjoint, but there does seem to be difficulties.

Hoogenboom: I agree that there is a problem in this case. Suppose that, in a multiplying system, I want to calculate a detector response at several points, or in several small volumes. The source distribution is the eigenfunction, which must be used as a scoring function. But you don't know the eigenfunction. This is a problem which I don't know how to solve. It should be possible to make a connection between the adjoint and the forward calculation in attacking the eigenvalue problem.

Gelbard: One thing you can do is this. You can keep a record of the location of the source sites in the direct mode problem. Then in the adjoint problem, on every collision, you get a line-of-sight estimate (a point estimate essentially) at the location of the fission source sites. In other words, you ask what is the adjoint flux at the various fission source sites. You do this by line-of-sight estimation.

Kalos: We have done things of this kind, you know, and it can be a big pain in the neck.

Gelbard: I think it is out of the question to sight to every fission source site on every collision. But you might weight the source sites with an importance function and, on every collision, pick <u>one</u> that you are going to contribute to. Suppose, also, that you draw a circle around the point where you want to estimate a reaction rate. If the site that you are contributing to is outside of that circle, then instead of actually drawing a line to the site and computing the contributions, you simply select a path length and lay it off. If you don't get to the selected site you make no contribution: if you do, you score. If the selected site lies inside the circle, then you use point estimation.

Kalos: I want to remind you that, in principle, there is a very simple and direct way to attack this problem.

Hoogenboom: I don't see how at the moment, specifically. Perhaps you can explain it to me later.

A MONTE CARLO TECHNIQUE FOR LOCAL PERTURBATIONS IN MULTIPLYING SYSTEMS

W. Bernnat

Institute of Kernergetik
University of Stuttgart, Germany

A MONTE CARLO TECHNIQUE FOR LOCAL PERTURBATIONS IN MULTIPLYING SYSTEMS

W. Bernnat

Institute of Kernenergetik
University of Stuttgart, Germany

ABSTRACT

The use of the Monte Carlo method for the calculation of reactivity perturbations in multiplying systems due to changes in geometry or composition requires a correlated sampling technique to make such calculations economical or in the case of very small perturbations even feasible. The technique discussed here is suitable for local perturbations. Very small perturbation regions will be treated by an adjoint mode. The perturbation of the source distribution due to the changed system and its reaction on the reactivity worth or other values of interest is taken into account by a fission matrix method. The formulation of the method and its application is discussed.

INTRODUCTION

Changes in the geometry or composition of a reactor system produce perturbations in reactivity and fluxes. To compute such perturbations by two independent Monte Carlo simulations is frequently very expensive because of the error in the differential value is as much as the sum of the errors of the two independent calculations. It is known that the error of differential values can be reduced drastically by means of a correlated sampling technique. Such a correlation can be achieved by estimating the perturbed and unperturbed system with an identical set of histories [1]. To realize a complete correlation between the perturbed and unperturbed system is difficult if the perturbation is complicated to describe. However, for the case of local perturbations a relative sufficient result can be achieved even if only a subset of histories in the respective systems is identical. Obvious if the perturbed problem will be estimated over the set of histories of the unperturbed problem, the corresponding estimator must be weighted. The weighting

accounts for the different source distribution and transport-kernel of the perturbed system with respect to the unperturbed system. In multiplying systems the source distribution will be calculated by a source iteration procedure. In practice such a source iteration will be carried out only for the unperturbed system. The source distribution of the perturbed system therefore has to be approximated. If the perturbation of the source distribution is only slight a first order approach will be sufficient [2]. Approximations of higher order can be achieved by means of a fission matrix method described in this paper. This method is also helpful to improve the convergence of the source iteration [3,4]. In the computation of local perturbations by the correlated sampling method one has to differentiate wether perturbation regions are large enough so that a reasonable number of neutrons will cross these regions. In the case of large regions it will be sufficient to calculate the differential value by correlated forward simulations. For small perturbation regions an additional adjoint run should be carried out to get an adequate number of neutrons which cross the small regions [1,5].

THE CORRELATED SAMPLING METHOD

Reaction rates or integrals of the form

$$R = \int_{\Gamma} g(\underline{P})\psi(\underline{P})\,d\underline{P} = \int_{\Gamma} \xi(C)\,\mu(C)\,d\underline{P} \tag{1}$$

where $g(\underline{P})$ is a response function
$\psi(\underline{P})$ the collision density
at the phase space point $\underline{P}$

will be estimated by an estimator $\xi(C)$, where C denotes a chain of collision points $C(\underline{P}_1, \underline{P}_2, \ldots, \underline{P}_k, \ldots)$. $\mu(C)$ is the probability that the chain C occures. A change of the regarded system in geometry or composition leads to the integral

$$\hat{R} = \int_{\Gamma} g(\underline{P})\hat{\psi}(\underline{P})\,d\underline{P} = \int_{\Gamma} \hat{\xi}(C)\,\hat{\mu}(C)\,d\underline{P} \tag{2}$$

We are interested in the difference of the integrals (2) and (1) $\Delta R = \hat{R} - R$. To reduce the variance of ΔR the neutron chains $C(\underline{P}_1, \underline{P}_2, \ldots, \underline{P}_k, \ldots)$ and $\hat{C}(\underline{P}_1, \underline{P}_2, \ldots, \underline{P}_k, \ldots)$ in the perturbed and the unperturbed system should be identical. That means $\hat{R}$ should be estimated over the chains C of the unperturbed system. The corresponding estimator $\tilde{\xi}(C)$ for the perturbed system then must be weighted by

$$\tilde{\xi}(C) = \xi(C)\,\frac{\hat{Q}(\underline{P}_1)\prod_{l=2}^{k}\hat{K}(\underline{P}_{l-1}, \underline{P}_l)}{Q(\underline{P}_1)\prod_{l=2}^{k}K(\underline{P}_{l-1}, \underline{P}_l)} \tag{3}$$

ΔR now can be estimated by

$$\Delta R = \int_{\Gamma} [\tilde{\xi}(C) - \xi(C)]\mu(C)\, d\underline{P}$$

Obviously the nominator in (3) must be equal to zero if the denominator is zero. However, this is not a severe restriction since the definition of perturbed or unperturbed system may be interchanged or an artifical system can be defined as the unperturbed system.

In the case of local perturbations there are many neutron chains which do not cross a perturbation region. The chains $C\ (\underline{P}_1, \underline{P}_2, \ldots, \underline{P}_k, \ldots)$ of such histories are identical. The corresponding estimators $\tilde{\xi}(C)$ and $\xi(C)$ differ only by the factor $\hat{S}(\underline{P}_1)/S(\underline{P}_1)$.

For histories which cross after some collisions a perturbation region, the weighting is more complicated. Frequently reasonable results can be achieved if only the histories from a source event up to their first crossing into a perturbation region are correlated. All corresponding contributions of such histories in the perturbed and unperturbed system differ only by a constant factor and therefore they are linear dependent so that these histories will not contribute to the statistical error of the difference value ΔR. For local perturbations it is useful to split the computation of ΔR into two parts as described by equation (4)

$$\Delta R = \int_{\Gamma_U} \xi(C)\left[\frac{\hat{S}(\underline{P}_1)}{S(\underline{P}_1)} - 1\right]\mu(C)\, d\underline{P} + \int_{\Gamma_P} [\tilde{\xi}(C) - \xi(C)]\mu(C)\, d\underline{P} \tag{4}$$

The first integral is taken over all histories which have not yet crossed a perturbation region. The second integral is taken over histories after crossing a perturbation region. The first integral in (4) is easy to calculate by Monte Carlo if the source distribution $\hat{S}(\underline{P})$ of the perturbed system is known. If we have only a slight perturbation in the source distribution $\hat{S}(\underline{P})$ may be calculated by the first order approach [2]

$$\hat{S}^{\nu}(\underline{P}) = \int_{\Gamma} \hat{K}_F(\underline{P}', \underline{P})\, \hat{S}^{\nu-1}(\underline{P}')\, d\underline{P}' \tag{5}$$

Here the source distribution $\hat{S}^{\nu}(\underline{P})$ at the ν-th iteration step is calculated from the unperturbed source at the ν-1st step. $K_F(\underline{P}', \underline{P})$ is the number of fissions in $\underline{P}$ due to a source neutron at $\underline{P}'$ for the perturbed system.

A second technique to calculate the perturbed source distribution is possible by a fission matrix method as described

below. For the calculation of the second integral in (4), which is the main contribution to ΔR, we consider only those neutron events starting from points $\underline{P}_R$ of the surface A_R of a perturbation region into the region itself. To get the source distribution $S_R(\underline{P}_R)$ on the surface A_R of a perturbation region we assume that the region is black for all neutrons /5/. This is an adequate assumption if the perturbation region contains no fissionable materials. However, by a modified procedure we can also treat regions containing a fission source. The source $S_R(\underline{P}_R)$ we can describe by the following expression

$$S_R(\underline{P}_R)=\int_\Gamma D(\underline{P},\underline{P}_R)\psi(\underline{P})\,d\underline{P} \tag{6}$$

where $D(\underline{P},\underline{P}_R)$ is a suitable response function for $S_R(\underline{P}_R)$ which is only different from zero if $\underline{P}$ lies on a surface point $\underline{P}_R$. If a perturbation region is small, then $S_R(\underline{P}_R)$ cannot be estimated sufficiently by a forward calculation. Here we make use from the duality

$$S_R(\underline{P}_R)=\int_\Gamma D(\underline{P},\underline{P}_R)\psi(\underline{P})\,d\underline{P} = \int_\Gamma S(\underline{P})\psi^+(\underline{P})\,d\underline{P} \tag{7}$$

where $\psi(\underline{P})$ and $\psi^+(\underline{P})$ are defined by the forward and adjoint transport equations

$$\psi(\underline{P})=\int_\Gamma K(\underline{P}',\underline{P})\psi(\underline{P}')\,d\underline{P}' + S(\underline{P})$$

$$\psi^+(\underline{P})=\int_\Gamma K(\underline{P},\underline{P}')\psi^+(\underline{P}')\,d\underline{P}' + D(\underline{P},\underline{P}_R)$$

The second integral of (4) now we can write by

$$\int_{\Gamma_P}[\tilde{\xi}(C)-\xi(C)]\mu(C)\,d\underline{P} \tag{8}$$

$$=\iint_\Gamma\int_\Gamma g(P)[\hat{K}_Q(\underline{P}_R,\underline{P})-K_Q(\underline{P}_R,\underline{P})]\,S_R(\underline{P}_R)\,d\underline{P}_R\,d\underline{P}$$

The kernel $K_Q(\underline{P},\underline{P})$ in (8) is defined by

$$\psi(\underline{P}) = \int_\Gamma K_Q(\underline{P}',\underline{P})S(\underline{P}')\,dP'.$$

Now if we know $S_R(\underline{P}_R)$ by a forward or adjoint Monte Carlo simulation we are able to calculate the integral (8) by a complete or not complete correlated set of histories drawn from $S_R(\underline{P}_R)$.

THE FISSION MATRIX METHOD

To get an acceleration of the source iteration in eigenvalue problems Kaplan /3/ has developed a fission matrix method. A refinement of this method has been made by Carter et al. /4/. The basic idea is the following.

We assume that the zones of the system containing fissionable material are subdivided into a number of spatial regions ΔVi, $i = 1,2,\ldots,N$. Then we compute the number of fission neutrons in ΔVi due to a source neutron in ΔVj by

$$m_{ij} = \frac{\int_{\Delta Vi}\int_{\Delta Vj} K_F(\underline{P}',\underline{P}) \quad S(\underline{P}')\, d\underline{P}'\, d\underline{P}}{\int_{\Delta Vi} S(\underline{P})\, d\underline{P}} \tag{9}$$

where $K_F(\underline{P}',\underline{P})$ is the number of fission neutrons in $\underline{P}$ due to a source neutron at $\underline{P}'$. The numbers m_{ij} form a matrix $M \equiv \{m_{ij}\}$ which is a discrete approximation of the kernel $K_F(\underline{P}',\underline{P})$. Now it is evident that the following equation holds.

$$\underline{S} = \frac{1}{k} M^T \underline{S} \tag{10}$$

where $\underline{S} \equiv \{si\}$ is the source vector and M^T is the transposed of the matrix M. The largest eigenvalue of M^T is an estimate of the criticality factor. If the source distribution $S(\underline{P})$ in (9) is identical with the equilibrium distribution then the estimate of k by equation (10) is unbiased. But for a sufficient number of regions and sufficiently long tracking a very good approximation of the elements m_{ij} can be achieved even if the source distribution $S(\underline{P})$ in (9) differs remarkably from the equilibrium distribution. This fact was the reason for us to use this method for perturbation problems. Under the assumption that $S(\underline{P})$ must not be necessarily identical to the equilibrium distribution we can compute the corresponding elements $\hat{m}_{ij}$ of the perturbed problem by means of the unperturbed source distribution $S(\underline{P})$

$$\hat{m}_{ij} \approx \frac{\int_{\Delta Vi}\int_{\Delta Vj} \hat{K}_F(\underline{P}',\underline{P}) \quad S(\underline{P}')\, d\underline{P}'\, d\underline{P}}{\int_{\Delta Vi} S(\underline{P})\, d\underline{P}} \tag{11}$$

As pointed out above, this is a good approximation of $\hat{m}_{ij}$ although the true equilibrium distribution $\hat{S}(\underline{P})$ has not been taken into account. However, the restriction $si = 0 \Rightarrow \hat{s}i = 0$ must hold. The criticality worth of the perturbed system is the largest eigenvalue of the equation

$$\hat{\underline{S}} = \frac{1}{\hat{k}} \hat{M}^T \hat{\underline{S}} \tag{12}$$

The reactivity change Δk is very easy to calculate now by $\Delta k = \hat{k} - k$. We get k, $\hat{k}$ and $\underline{S}$, $\hat{\underline{S}}$ by numerical solution of the equations (10) and (12).

For the calculation of M and $\hat{M}$ the correlated sampling method as described above must be used, of course. In practice we will compute m_{ij} by a forward source iteration. Then for local perturbations we compute $S_R(\underline{P}_R)$ on the surfaces of perturbation regions in the regarded system. $S_R(\underline{P}_R)$ may be determined during the forward source iteration or by an additional adjoint run /5/. If we have M and $S_R(\underline{P}_R)$ then we may compute starting from $S_R(\underline{P}_R)$ an arbitrary set of sequential runs for changed systems with complete correlated or not complete correlated sampling techniques to get Δm_{ij}. By writing

$$\hat{m}_{ij} = m_{ij} + \Delta m_{ij}$$

we get the elements of the matrix $\hat{M}$ so that we can solve equation (12). Changes in activation rates or fluxes we compute by equation (4). Here the difference between unperturbed and perturbed source distribution will be corrected zonewise by

$$\hat{s}i_{/si} \quad i = 1,2,\ldots,N$$

with si, $\hat{s}i$ from the solutions of (10) and (12).

In some problems it is more effective to compute the differential change $\partial R/\partial X$, not the absolute value ΔR. Examples were given by Olhoeft in the calculation of Doppler coefficients $\partial R/\partial T$ /6/, Takahashi for geometrical changes $\partial R/\partial X$ /7/ and others. As a consequence for such differential changes we have to compute $\partial m_{ij}/\partial x$ instead of Δm_{ij}. If such changes are very small we should no longer compute Δk as the difference of the eigenvalues of equation (10) and (12) because the relative numerical error in Δk increases.

Here it is advantageous to make use of the perturbation theory. We can show /3/, that the fundamental (right) eigenvector of the matrix M, defined by the equation

$$\underline{W} = \frac{1}{k} M \underline{W}$$

can be interpreted as an importance vector. The elements wi of $\underline{W}$ have the meaning of the importance of a fission neutron starting in Δ Vi (see also Matthes /2/). Now we can calculate small effects by the following

$$\frac{1}{k} \cdot \frac{\partial k}{\partial X} = \frac{\left[\underline{W} \frac{\partial M}{\partial X} \underline{S} \right]}{[\underline{W} M \underline{S}]}$$

where $[\underline{W} M \underline{S}]$ denotes a space integration. The computation of $\underline{W}, \hat{\underline{W}}$ is as easy as the computation of $\underline{S}, \hat{\underline{S}}$ if the matrixes M, $\hat{M}$ are known. The error in Δk or in other differential values of interest depends on the effectiveness of the correlated sampling

technique and the choice of the region subdivision and the number of followed histories, of course.

APPLICATION

The applicability of the here discussed method was studied on several problems. The program which we used for our purpose was the MORSE-K code /8/ wich is a modified version of the multigroup MORSE-code /9/. The calculation of strong local perturbations due to moving reflector elements or control rods in both compact reactors and power reactors (pebble-bed reactor) were calculated by means of the not complete correlated Monte Carlo technique, using the fission matrix method to take into account the source perturbation /10/. Reasonable results were obtained. At time the not complete correlated sampling technique and the fission matrix method is standard in MORSE-K. The complete correlated method is very problem dependent and requires often a special treatment by the user. The computation of differential effects for small perturbation regions by an additional adjoint run is possible by a special version of MORSE-K /5/. To demonstrate the method a simple one-dimensional problem was calculated by Monte Carlo and S_N. The description of the unperturbed system is given in tab. 1. The core is divided into 10 regions. In the basic problem all core regions contain the same material UZrH. The example is related to an experiment on a compact zirconium-hydride moderated and beryllium reflected zero power reactor. The thickness of the beryllium reflector is 6 cm in the unperturbed case. Now some changes of the system have been studied by Monte Carlo and S_N as described in Table 2. In this table also the k-values and the Δ k-values calculated by Monte Carlo and S_N are presented. It is clear that this changes are not very small, but nevertheless the saving of the computing time is remarkable if the correlated method is used. Further the assumption that the fission matrix elements may be calculated by the unperturbed source distribution can be proved. The fission source vectors (fission rates) for the unperturbed and three perturbed systems described by Table 2 are given in Tables 1 and 3. The results are quite sufficient. Though this example is for a small reactor only, the application of the method to power reactors has proved to be very efficient /10/.

CONCLUSIONS

For local perturbations in multiplying systems it is effective to calculate differential values by a correlated sampling technique. Though a complete correlated sampling is the best method the operation with the simpler not fully correlated method is sufficient for many problems. To get unbiased differential values the source perturbation must be taken into account. The fission matrix method here gives a reasonable approximation to the perturbed source. Large perturbation regions

may be treated by means of forward calculations for the unperturbed problem and the perturbed problem. In the case of small perturbation regions an additional adjoint run should be used. The method implemented in MORSE-K is not necessarily restricted to multigroup-Monte Carlo, however, running adjoint problems is simpler by a multigroup code.

REFERENCES

1. J. SPANIER and E. M. GELBARD, Monte Carlo Principles and Neutron Transport Problems, Addison-Wesley, Reading, Massachusetts (1969).

2. W. MATTHES, "Calculation of Reactivity Perturbations with the Monte Carlo Method", Nucl. Sci. Eng., 47, 234 (1972).

3. E. L. KAPLAN, "Monte Carlo Methods for Equilibrium Problems in Neutron Multiplication", UCLR, 5275-T (1958).

4. L. L. CARTER, N. S. Mc CORMICK, "Source Convergence in Monte Carlo Calculations", Nucl. Sci. Eng., 36, 438 (1969).

5. E. LINDAUER, "Monte Carlo Berechnung geometrischer Störungen bei Kernreaktoren unter Benutzung adjungierter Beziehungen", Thesis Technical University of Berlin (1973).

6. S. E. OLHOEFT, "The Doppler Effect for a Non-Uniform Temperature Distribution in Reactor Fuel Elements", WCAP-2048, Westinghouse Electric Corporation, Atomic Power Division, Pittsburgh (1962).

7. H. TAKAHASHI, "Monte Carlo for Geometrical Perturbation and its Application to the Pulsed Fast Reactor", Nucl. Sci. Eng. 41, 259 (1970).

8. F. A. R. SCHMIDT, W. BERNNAT, MORSE-K Manual, Rep. IKE 4-1, Institut für Kernenergetik, University of Stuttgart (1972).

9. E. A. STRAKER et al, "The MORSE Code-A Multigroup Neutron and Gamma-Ray Monte Carlo Transport Code", ORNL-4585, UC-80-Reactor Technology (1970).

10. W. BERNNAT, A. HUEBNER, "Monte Carlo Verfahren zur Berechnung der Regelstabwirksamkeit im oberen Hohlraum von Kugelhaufenreaktoren", in Proc. Reactor Meeting DAtF Berlin, (1974), p. 37.

TABLE I

Description of the Unperturbed Test Problem (Infinite Cylinder). Fission Rates Calculated by Monte Carlo and S_N

Zone	Material	Outer Radius (cm)	Fission Rate[a] Monte Carlo	Error[b] (%)	Fission Rate S_N
1	$UZrH_x$	1.5	0.0333	6.8	.0342
2	$UZrH_x$	2.0	0.0264	6.0	.0263
3	$UZrH_x$	2.5	0.0335	5.0	.0335
4	$UZrH_x$	3.0	0.0407	4.2	.0404
5	$UZrH_x$	4.0	0.1007	3.2	.1005
6	$UZrH_x$	5.0	0.1247	3.0	.1242
7	$UZrH_x$	6.0	0.1425	3.0	.1443
8	$UZrH_x$	7.0	0.1587	2.6	.1606
9	$UZrH_x$	8.0	0.1692	2.7	.1739
10	$UZrH_x$	9.0	0.1928	2.7	.1954
11	Be	15.0	0.	-	0.

[a] 90 batches with 200 neutrons are processed for the source iteration.

[b] the estimated error is given in $\pm 2\sigma$-bounds

TABLE II

Description of Changed Systems and Results Calculated by Monte Carlo and S_N

	Change	Monte Carlo k^a	Monte Carlo $\Delta k/k$	Error (%)^b	S_N $\Delta k/k$ ^{c}
0	No	1.023	-	-	-
1	Outer radius 14.14 cm	0.9894	-0.034	7.5	-0.035
2	Outer radius 13.28 cm	0.9525	-0.074	6.1	-0.078
3	Outer radius 11.14 cm	0.8259	-0.239	3.2	-0.242
4	Al in zone 1 Outer radius 15.00 cm	1.002	-0.021	9.5	-0.020
5	Al in zone 1 Outer radius 11.14 cm	0.7994	-0.280	3.2	-0.284

[a]The error in k is 1.1 % ($\pm$ 2 σ - bounds)

[b]the estimated error is given in $\pm$ 2σ - bounds

[c]the unperturbed value of k calculated by the S_N method is k = 1.0335.

TABLE III

Comparison of Fission Rates for Changed Systems Calculated by Monte Carlo and S_N

	Changed system 2		Changed system 3		Changed system 4	
Zone	Fission Rate M. C.[a]	Fission Rate S_N	Fission Rate M. C.	Fission Rate S_N	Fission Rate M. C.	Fission Rate S_N
1	0.0332	0.0338	0.0326	0.0332	0.0	0.0
2	0.0263	0.0260	0.0253	0.0254	0.0258	0.0256
3	0.0333	0.0330	0.0324	0.0321	0.0336	0.0329
4	0.0404	0.0397	0.0390	0.0384	0.0402	0.0400
5	0.0989	0.0980	0.0942	0.0934	0.0992	0.0998
6	0.1202	0.1196	0.1122	0.1115	0.1236	0.1239
7	0.1355	0.1367	0.1228	0.1234	0.1451	0.1459
8	0.1489	0.1487	0.1285	0.1285	0.1619	0.1614
9	0.1546	0.1553	0.1257	0.1257	0.1743	0.1749
10	0.1613	0.1613	0.1131	0.1143	0.1963	0.1967

[a]Calculated by Monte Carlo, the estimated errors are similar as in Table I.

DISCUSSION

Hoffman: Do I understand correctly that, if you extend the region around your local perturbation to include your entire system, you are doing essentially the same thing that Matthes did? Is that fair to say?

Bernnat: No. I don't understand your question.

Hoffman: You had two terms in your treatment. One was just an integral over the perturbed region. If you extended that boundary of the perturbed region to include your entire system, would you be doing essentially the same thing that Matthes did?

Bernnat: I would say that in this case the perturbation is not a local perturbation and this method would not be adequate. You would then have to use complete Monte Carlo correlation methods. If your whole system is changed, for example, due to a temperature change or a change in atomic density, and if the cross-section perturbations are not very large, then it is better to do the calculation by a complete correlation calculation. But, here in the case of strong local perturbation the complete correlation method seems to me to be not very efficient because the first collision in the perturbed region, in the perturbed system, may change the weight of one of the sample particles to zero. Then, on subsequent collisions, the weight of the particle representing the perturbed system will remain zero.

Gelbard: Are there other comments?

Borgwaldt: I have a question for Dr. Gelbard. In your paper you mentioned the method of complete correlation. Could you give some details about the method which you are going to use?

Gelbard: At this point, we are hoping that we can get away with something very much less sophisticated than Bernnat's method. We want to do a local perturbation calculation by introducing a composition halfway between the perturbed and unperturbed composition, tracking in this intermediate composition and simultaneously using weights for the perturbed and unperturbed problems. Now we realize that the most serious difficulty, if you try to do this, is just the difficulty that has been discussed already, namely that you have to get particles into the perturbed region. We would like to try the simplest approach we can think of, which is simply to build up the density in the neighborhood of the perturbation by splitting. We don't know whether this will really work. Before we would make a very large commitment to this method, we think that we can get an estimate as to whether it is feasible or not by simply asking ourselves, not whether we can compute the perturbation due to a small sample accurately, but simply whether we can compute the absorption rate in a small sample reasonably accurately. We feel that if we can get enough histories into the sample so that we could get a reasonable estimate of the absorption rate, then by later introducing the simultaneous tracking technique we would perhaps be able to get a reasonable estimate of the perturbation. This proposed method does not turn out to be adequate, then we will have to try something more complicated, and probably something that does involve an adjoint. We have considered many methods involving the adjoint and there always seem to be some difficulties involved. If there are no other questions about Bernnat's paper, I would like to bring up some of these problems that

seem awkward to us. In computing the surface source you said you would use an adjoint method in which, essentially you start the adjoint source from the surface. You then track in the adjoint mode and tally the product of the adjoint flux and the source density. If you try to get a detailed energy angle space distribution of particles entering the perturbed region, then for every energy, angle and position you need an adjoint. Presumably, you don't do that; but then how can you be sure that you have an adequate energy angle and space distribution of incoming particles? That is the first question.

Bernnat: It may be that if I have no knowledge of the spectrum in the surrounding region, this adjoint method will be very inefficient. But it may be possible to calculate, by simpler methods, the spectrum in the neighborhood of such a perturbed region. If you have some knowledge about the spectrum, then you can start your adjoint mode calculation very easily. We have tested this method only in a multigroup program. In such a case it is very easy to use.

Gelbard: Well, then the energy variation would be simplified, but you would still have to get a space-angle distribution at each energy, and that still seems like quite a computational problem. But you would have a difficult time knowing whether you had adequately sampled the incoming particles, which leads me to the other problem. If you are trying to analyze an experiment, and there is a 10% discrepancy between the observed perturbation and somebody's computed perturbation, you would like to be able to guarantee that your perturbation calculation is accurate to much better than 10%. If you could do a rigorous Monte Carlo and get a reliable uncertainty estimate, then you could argue that, if you disagree with the experiment by five standard deviations, there is probably something wrong. On the other hand, if you do a calculation in which you have made approximations which you cannot evaluate, then you will probably have to go back and spend a good deal of effort trying to understand what these approximations do. In the end you will not be able to say whether the 10% discrepancy between Monte Carlo and experiment is real, or is due to your approximations. This is another thing that has made us reluctant to go this way. If you use the Green's function approach and try to accumulate the Green's function over regions in the course of the Monte Carlo, then you have the following delemma. If you make the regions too coarse you are introducing an unevaluated approximation; but, if you make the regions too fine, the statistical fluctuations in the Green's function again introduce an unevaluated approximation. You may argue that the fluctuations in one region will cancel those in the other regions. But again, if somebody pins you down and says,"Look, there is a 10% discrepancy between the Monte Carlo and the experiment. Now, is there something wrong, or are you just not giving us the right answer?" What do you say?

Bernnat: Well, I say that I have tested the method against an S_n calculation and the agreement was rather good. Of course, this was a very simple problem. Further, we have compared computed results with measurements and we have found that our agreement was within about 10%. In addition, we have changed the role of our reference and perturbed problems. This method tends to rather over or underestimate the perturbed source distribution, depending on whether you start from control rods in or from control rods out, and we have found that if we interchange the roles of the perturbed and unperturbed systems the results remain about the same. We think we can say that substantial discrepancies with experiment cannot be due to the method.

Gelbard: In the control rod calculation that you showed, did you use the adjoint method?

Bernnat: No.

Gelbard: Now, in the slab calculation, the adjoint method was a particularly easy one because you had no spatial variation over the surface of the slab. You had an angle variation and an energy variation. If the incoming flux is more or less isotropic, the angular variation does not matter very much. Now, if you have a small sample perturbation due, let us say, to a finite cylinder that is inside some sort of a fairly complicated lattice array, then you have to worry about sampling the positions on the surface adequately, and you would have a more difficult problem. These are reasons why we are hoping to use an approach which has no approximations. In the method we were trying to use there are two difficulties. First, it is not clear that we can really get enough particles into the perturbed region to give us reliable estimates. The second problem, I think, has been mentioned by a number of people. When you go to such a biased sampling scheme, there is always a serious question as to whether you can trust the variances; so we could be driven to all kinds of other possibilities. But this is our motivation. We were trying a simple scheme first. We will go to all kind of lengths to build up the particle density in the perturbed region. For example, we might draw a boundary, let us say, a spherical boundary, around our cylindrical sample, and on every collision which is made within a certain distance of this sphere, we might do line-of-sight estimation into this sphere. So we build up the particle density in the neighborhood of the cylinder by splitting, and additionally do line-of-sight estimations. Finally, we plan to use a perturbation technique, taking as our base problem one intermediate between the perturbed and unperturbed problems. There is an interesting Russian paper in which the author tries to fix optimum parameters for an intermediate problem. In the case where you have material in or material out, you would like your base problem to be intermediate between these two extremes. The authors determine what the cross sections in this base problem should be in order to minimize the variance. We would use this sort of intermediate problem.

Whitesides: I should point out that Tom Hoffman developed a method while he was with us at Oak Ridge in which he modified the method of Matthes to be an exact treatment. We have done a number of test calculations with this modified method in situations where we could check Monte Carlo results against ANISN. We have calculated Δk's from 0.004 up to about 8% in k, and obtained quite good agreement with ANISN for various types of perturbations. The problem, as Gelbard pointed out, is getting enough histories into the region you are interested in. We have actually been applying Hoffman's method to practical problems, looking at control rod worths in ZPR-type experiments, and we think we are having better luck than most people in calculating the Δk's. Our largest error appears to be about 10% with most of the errors down about 3 or 4%. So we think this is a viable approach. We have it built into our code and are using it fairly extensively now.

Gelbard: So you're getting Δk's to within 3 or 4% of Δk?

Whitesides: Of the experimental value.

Gelbard: Oh, I see, you are getting Δk to within 3 or 4% of the experimental value. Are you analyzing central-worth experiments? Are you seeing the usual central-worth discrepancies?

Whitesides: For the particular experiments we have carried out we have not seen them. We have not analyzed some of the most difficult experiments, and we don't know whether we can. But, our agreement with the ANISN calculations is just spectacular. The technique is reasonably fast: it takes about twice as long to do a perturbation calculation as it does to do a normal calculation. The problems we are having now are ones that you would have expected. One is that the method works quite well for absorption changes, just because all of the changes are in the same direction. Scattering changes are more difficult and the errors, of course, are larger for those. At present we're using a collision estimator, so it is necessary to be careful in defining the unperturbed problem. The unperturbed problem should always be the problem in which more collisions occur. For example, if you have a void region, and perturb the configuration by putting material into the void, the configuration with the material present should be the base configuration. But we're putting a track length estimator into our code, and when we use the track length estimator it should be possible to use either problem as the base problem.

Gelbard: It might be interesting to hear a little more about Hoffman's method.

Hoffman: Well, Matthes' method involves calculating the adjoint simultaneously with the forward function, and this is done in the unperturbed system; Matthes ends up with a first-order approximation to his change in k. On the other hand, once we generate this perturbation source, we change the geometry system. We really keep track of two geometries and track in the perturbed system. By doing this we can make an adjustment in the importance function, and also an adjustment in one of the terms in the expression for Δk; and we end up with an exact method, except for the fact that (as in Matthes' method) the fissile material must be divided into sections where the importance function is assumed to be constant.

Gelbard: Now let me compare the Matthes method and the correlation method and see whether there is any disagreement about this. The main problem in the two methods is basically the same. You have to get enough particles into the sample.

Hoffman: Yes, but one thing that is nice about this method is that the change in the system is your perturbation source.

Gelbard: But you have to be able to get particles into the perturbed region.

Hoffman: You have to be able to calculate the perturbation source.

Kalos: You don't have to get them there. Once you calculate them, you put them there.

Gelbard: Well, don't you have to track in the direct mode first to get the particles into the region? Or you can use line-of-sight estimation; but in one way or another you have to get information about the flux entering the

perturbed region. Now it seems to me that in correlated tracking the problem is exactly the same. You have to build up a density of information in the perturbed region. In the Matthes method you have a problem in that some sources are positive and some sources are negative. In the direct correlation method you have the problem that some weights are bigger than one and some weights are smaller than one. All in all, I don't see that one method is likely to be any better than the other.

Hoffman: Well, all in all, I don't see why one method would be equivalent to the other either.

Gelbard: Oh, I am not saying that they are going to be the same all the time, but I see no a priori reason why I would go one way rather than the other. It is strictly a matter of convenience. I am saying that if Matthes' method really works, then it must be that you are able to get enough information into the perturbed region, that the perturbed source can be computed adequately. If you can do that there is a good likelihood that you have enough information in the perturbed region to estimate the difference between the perturbed and unperturbed weights adequately. So, in looking at the two methods, it seems to me that there is, on the face of it, no reason to prefer one over the other.

Bernnat: I find that the adjoint vector of the Matthes method is exactly the right eigenvector of the matrix I discussed earlier.

Gelbard: But, I need make no approximations at all in the complete correlation method. While (whether the approximation is a good one or a bad one) I do have to make an approximation if I am forced to assume that the adjoint is flat over blocks. I must worry about whether it is or it is not. So I don't see what I am burying for that approximation.

Bernnat: The problem is to get enough samples in the perturbation region.

Gelbard: Yes, you approach this problem by starting an adjoint from the surface.

Bernnat: Actually the control rod calculations I described were not performed by an adjoint calculation. In the control rod problem, the density of neutrons in the upper core was amplified by splitting.

Steinberg: I would like to comment on the direct correlation method. You can avoid the problem of weights greater than one by a procedure that we used in the SAMCEP code, where we set up our sampling based on the maximum of the density functions of the various perturbed and unperturbed distributions.

Prael: I don't think that any of the perturbations which have been considered there are small, when compared with the perturbations we have been asked to deal with at Argonne.

Gelbard: But the problems that were done at Oak Ridge were, as I understand it, small sample perturbations like those in some of the Argonne experiments, and they actually would be the kind of perturbation in which we would be interested.

Prael: How small? The samples in our experiments were cylinders a few centimeters high.

Whitesides: We have looked at several things, but obviously nothing that small. We have been dealing with several control rods moving inside the reactor.

Kalos: Ely, I think I have been getting lost — lost again! When you say you propose to do an exact calculation you are talking about a Δk? Do I understand correctly that, in effect, you are going to do two correlated problems and get an eigenvalue for each of them? Which means correlating weights in many generations?

Gelbard: That is right, we would. In other words, when a fission particle is produced — if it is produced in a region where in the other problem there would not be any fission — it carries a zero weight. You are saying they get uncorrelated as we go through generations?

Kalos: I'm afraid that sort of effect would tend to build up.

Whitesides: I would say that the method of Matthes works beautifully if the perturbation, no matter how small it is, extends over a large volume. But correlated sampling works best if you have a very small perturbed region with everything else in the unperturbed system. So it seems to me that they are complementary methods.

Gelbard: Do I understand correctly that, in doing the adjoint calculation, you pick the sample points uniformly on the surface? And that you then pick energies uniformly from an energy band? Do you pick angles uniformly?

Bernnat: Yes.

Gelbard: Once you have chosen an angle and a point, do you run just one history and then go on to another one? Or do you start more than one history from a given point in the given direction?

Bernnat: I don't understand quite ...

Gelbard: Well, there are two ways one might do it. On having chosen an energy, and a point on the surface, one approach would be to actually try, by using the adjoint method, to construct the angle-dependent source at that point. That would require running many adjoint histories from that point.

Bernnat: Yes.

Gelbard: I was thinking that another approach would not be necessarily to try to construct the incoming source at that point exactly. Once you pick a point you could randomly pick one energy and one angle and run one adjoint history from that point, and then switch to another point. It is not necessary really to try to construct the incoming flux with any accuracy at each point since you are going to take a sum anyhow over all contributions. Have you thought about this strategy, or exactly how do you do it?

Bernnat: No, I sample from a distribution. I don't pick one point and then sample all energies and all angles. I assume a distribution on the surface, say a constant spatial distribution, and perhaps a cosine distribution for the angle. Then I sample of lot of neutron histories from this distribution; but I don't repeatedly use the same sample point.

Gelbard: In other words, what you do is to pick a space point uniformly, pick an energy, pick an angle, send out one history, get its contribution, and then pick another starting set?

Bernnat: Yes.

Gelbard: You then have to evaluate the integral of the adjoint flux times the source.

Bernnat: Yes.

Gelbard: Presumably you have to assume then that the source is uniform over specified regions.

Bernnat: Yes.

Gelbard: Another approach would be the one that was suggested yesterday. That is to say that, when you are running the adjoint problem, then on each collision, in principle, you could use a line-of-sight estimator to one of the source points (if that were practical). You could weight the source points by the value of the importance function at each point. From the weights you could form a probability distribution. At each collision you could take a source point from this distribution and get a contribution to that source point.

Bernnat: I have never tried this.

Gelbard: I was thinking of what you might do to avoid any approximation. In this way, if it were practical, you could avoid the approximation that the fission source is uniform over boxes.

Bernnat: It may be practical but it is more difficult. I would have to change my process.

Gelbard: Could you not, in this way, completely avoid the Green's function method also? In other words, is it really necessary to separate out the perturbed source contribution and calculate it by a Green's function method? Couldn't you avoid the Green's function method completely.

Bernnat: Yes, I can. But in the case of very strong perturbations, I must take into account the perturbed source distributions, and I need the Green's function to calculate my perturbed source distribution. Of course, for my adjoint calculation I would not need a Green's function, really.

BOUNDED ESTIMATION OF FLUX-AT-A-POINT FOR ONE OR MORE DETECTORS

H. A. Steinberg

Mathematical Applications Group, Inc.
Elmsford, New York 10523, U.S.A.

BOUNDED ESTIMATION OF FLUX-AT-A-POINT FOR ONE OR MORE DETECTORS

Dr. Herbert A. Steinberg

Mathematical Applications Group, Inc.
3 Westchester Plaza
Elmsford, New York 10523

ABSTRACT

The procedure for biasing the collision density of Monte Carlo particles, so that the estimation of flux at a point detector is bounded, is described. The theory and program design of the solution of the problem of obtaining estimates for several detectors during a single Monte Carlo calculation is presented.

INTRODUCTION

The estimation of neutron or photon flux at a specified point in space is a major objective of radiation transport calculations. In carrying out calculations for problems in three dimensions, it is usually necessary to resort to Monte Carlo methods. The typical procedure that is invoked for estimating flux at a point involves the generation of particles with random positions and momenta, sampled from the collision density, and using a last flight estimator to obtain an estimate of the flux at the prescribed position. The standard last flight estimator W is given by:

$$W = g(\omega)\exp(-\lambda(x,D))/r^2 \tag{1}$$

(See Table I and Figure 1 for definitions.)

The major problem in the treatment of point detectors is that $W \to \infty$ as $r \to 0$, i.e. the estimator is unbounded. Furthermore, as an added complication for Monte Carlo calculations, $\langle W^2 \rangle$ is infinite when it is averaged with respect to the natural collision density. As a result convergence of a Monte Carlo average to the correct answer is slower than the usual $1/\sqrt{N}$, where N is the number of samples.

HISTORICAL BACKGROUND

The first significant attempt at solving this problem was made by Kalos[1], using the concept of a once more collided estimator. Here, instead of using the estimator described by Equation (1), a procedure was developed where an intermediate point was selected from a biased distribution. The estimation then was the product of two terms of the form of

Equation (1). Although the resultant estimation was still unbounded, the second moment, and therefore the variance, was finite, so that the convergence was $1/\sqrt{N}$.

However, the possibility of arbitrarily large estimates remained, thus leaving room for a large error in a given Monte Carlo calculation, albeit with a small probability. To mitigate this problem, Steinberg and Kalos[2] developed the concept of a bounded estimator. They invoked the principle that, as long as the mean is finite, it is always possible to establish a biased sampling scheme in which the singularity becomes part of the probability density.

SINGLE DETECTOR-THEORY

The problem of handling the $1/r^2$ singularity in Equation (1) was treated by using biased distributions for the selection of the direction out of collision (or source) and for the position of the next collision. This results in a biased collision density in the neighborhood of the detector proportional to the "natural" collision density divided by r^2. The resultant modified estimator Wr^2 is therefore bounded.

As an aid to the implementation of the biasing procedure, the reselection principle was developed. In general terms, this procedure involves first the selection of a point in phase space by means of some prescribed (e.g. natural) distribution. Then, if the point is in some critical domain (e.g. where W as given in Equation (1) is large), reselect the point within the domain from an appropriately biased distribution.

Reselection would be invoked when two conditions are met. First it is desirable to keep the biasing local, i.e. there is some critical domain in which biasing is necessary, while for the rest of phase space the "natural" distribution is preferred. Second, the "natural" probability, that a point selected in the first stage will fall in the critical domain, is not known a priori or is difficult to calculate. In [2] weight adjustments which do not use this probability are described.

To select the direction, it is convenient to define a coordinate system using the line between the particle position and the detector as a polar axis (direction ω_D). The differential solid angle $d\Omega$ can then be expressed as:

$$d\Omega = \sin\alpha d\alpha d\psi \tag{2}$$

where α is the polar angle and ψ is an azimuth around it. The biasing procedure here involves α only. Specifically the density function is set to be proportional to $1/\sin\alpha$ for $0<\alpha<\alpha_m$, where α_m is some preassigned limit. This means simply that within the range $(0,\alpha_m)$, α is chosen by means of a uniform distribution in angle, rather than uniform in the cosine of the angle. The biasing weight adjustment F_0 is then given by:

$$F_0 = K_0 g(\omega)\sin\alpha \tag{3}$$

where K_0 is determined by α_m.

To select the position of the subsequent collision, the density function is set proportional to $1/r^2$, and the weight adjustment factor F_1, proportional to r^2, is given by

$$F_1 = K_1((S-S_0)^2+r_0^2)/r_0 \tag{4}$$

where r_0, S_0, and S are given in Table I, K_1 is determined by the range over which S is biased, and includes the adjustments for the "natural" exponential attenuation. The product WF_0F_1, used for flux estimation, can then be shown to be bounded, solving the problem of the $1/r^2$ singularity.

SINGLE DETECTOR IMPLEMENTATION

An implementation of the theory was described in [2], using the reselection principle. Essentially if a collision was found to be within a prescribed volume around the detector, both the previously selected direction and the position along the track were reselected. This led to serious programming problems. Furthermore, no attempt was made to handle the "natural" exponential attenuation.

Steinberg and Lichtenstein[3] described an improved implementation procedure which took into account both the exponential attenuation and the $1/r^2$ term in carrying out the spatial sampling. In addition, to avoid the above mentioned programming difficulty, it was determined that the direction and distance biasing should be decoupled. The direction biasing, using the reselection principle, was carried out as soon as the direction was chosen (and not later as in [2]). Then the biasing of position along the track near a detector can be carried out without any further effect on the direction and also without the use of reselection.

Specifically, when selecting a direction ω_p, a test is made to determine if the initially selected direction ω_q lies within a specified cone, i.e. if the angle α_0 between ω_q and ω is less than α_m (see Figure 1). If so, a new α is chosen uniformly between 0 and α_m and ω_p is calculated from

$$\omega_p = a\omega_q + b\omega \tag{5}$$

where $a = \sin\alpha/\sin\alpha_0$ and $b = \sin(\alpha_0-\alpha)/\sin\alpha_0$. When $\alpha_0>\alpha_m$, ω_q becomes ω_p.

When reselection is used, the weight adjustment F_0 is then given by:

$$F_0 = \frac{\left(\sin\alpha\right)_{\alpha_m}}{1-\cos\alpha_m}\,\frac{g(\omega_p)}{g(\omega_q)} \tag{6}$$

For position along the track, biasing can be carried out without resorting to reselection. The procedure is kept local by defining a range of distances along the track, near the detector where biasing is used. The range (S_-,S_+) is determined from

$$S_{\pm} = S_0 \pm\sqrt{\rho^2 - r_0^2} \tag{7}$$

This interval defines the portion of the path within ρ of the detector. If $r_0 \geq \rho$, then no special biasing is needed.

For $S_- < S < S_+$, a sampling density $f(S) = g(S) / \int_{S_-}^{S_+} g(t)dt$ can be used where

$$g(S) = \mu(S)\exp(-\lambda(S_-,S))/((S-S_0)^2+r_0^2) \tag{8}$$

Because $\int g(S)dS$ is difficult to obtain analytically, $g(S)$ is replaced by an approximation, constructed by dividing up the range (S_-,S_+) into short intervals and replacing either the exponential or the inverse square factor by a constant. A full description of the algorithm is given in the SAM-CE Manual [4].

SEVERAL DETECTORS SOLUTION

The problem of treating several point detectors simultaneously although not a difficult problem in theory, was not satisfactorily solved in either [2] or [3].

In [2] the theoretical solution of sampling from the average distribution was described, but the implementation suggested proved impractical. In [3] a practical solution, involving the use of Russian Roulette procedure to reduce the problem to that of one detector, was described. When implemented, however, this approach results in a serious degradation in the statistical quality of the resultant flux estimates. A full description can be found in [4].

However, by combining the method of [3] directly with the theoretical solution of [2], it is feasible to define a procedure which would be satisfactory statistically, and still be tractable for implementation. The procedure is based on a simple but important theorem of Monte Carlo. Specifically:

Theorem 1:

Let $F_1(x), ---, F_n(x)$ be a collection of probability distribution functions. Let $a_1, ---, a_n$ be constants so that $\Sigma a_k = 1$ and all $a_k \geq 0$. Let $F(x) = \Sigma a_k F_k(x)$. Then $F(x)$ is a distribution function, and a random variable X with distribution function $F(x)$ can be generated by means of the following two-step procedure:

1: choose k at random with probability a_k.

2: choose a random variable X with distribution function $F_k(x)$.

Proof:

$$\text{Prob}(X \leq x) = \sum_k \text{Prob}\ (k \text{ selected})\ \text{Prob}\ (X \leq x \mid k \text{ selected})$$

$$= \sum_k a_k F_k(x)$$

$$= F(x) \qquad \text{Q.E.D.}$$

This theorem leads to the following estimation theorem.

Theorem 2:

Let $F_k(x)$ and a_k be defined as in Theorem 1, let X be chosen as described in Theorem 1, let $T(x)$ be integrable over the domain of x, and let $\{b_k(x)\}$ be any functions of x, provided that $\Sigma a_m b_m(x)$ is never 0. Then an unbiased estimate J of $\int T(x)dx$ is given by:

$$J(X) = \frac{T(X)b_k(X)}{f_k(X)\ \sum_m a_m b_m(X)} \tag{9}$$

where $f_k(x) = F_k'(x)$. Note that J depends on k.

Proof:

$$\langle J \rangle = \Sigma a_k \int J(x) f_k(x) dx$$

$$= \Sigma a_k \int \frac{T(x)b_k(x)}{\sum_m a_m b_m(x)} dx \tag{10}$$

$$= \int T(x)dx$$

Q.E.D.

Implementation of the solution described in [2] is then very straightforward. Before choosing a direction out of collision, detector number k is chosen with probability a_k (where the a_k are defined in advance to reflect the relative importance of the various detectors).

The direction and next position sampling are carried out as if the kth detector were the only detector present using essentially the same procedure as described above for one detector.

The estimation theorem can then be invoked using

$$b_m(x) = \max(1, \rho_m^2/r_m^2) \tag{11}$$

For each detector $W_j/\Sigma a_m b_m(x)$ is bounded. (In fact, $\Sigma a_j W_j/\Sigma a_m b_m(x)$ is bounded.) Furthermore $b_k(x)K_0K_1$, where k is the detector used for sampling and K_0K_1 are the biasing weight adjustments, is also bounded. Therefore the estimation for all detectors is bounded.

Since the bound on $W_j/\Sigma a_m b_m(x)$ can be $1/a_j$ larger than the one detector bound, it is recommended that this procedure be invoked only when necessary, i.e. if the direction out of collision has to be biased by more than one detector.

REFERENCES

1. M. H. Kalos, "On the Estimation of Flux-at-a-Point by Monte Carlo", Nucl. Sci. Eng. 16, 111 (1963).

2. H. A. Steinberg and M. H. Kalos, "Bounded Estimators for Flux-at-a-Point in Monte Carlo", Nucl. Sci. Eng. 44, 406 (1971).

3. H. Steinberg and H. Lichtenstein, "Implementation of Bounded Point Estimators in Point Cross-Section Monte Carlo Codes", Trans. Am. Nucl. Soc. 17, 259 (1973).

4. M. O. Cohen et al, "SAM-CE: A Three Dimensional Monte Carlo Code for the Solution of the Forward Neutron and Forward and Adjoint Gamma Ray Transport Equations - Revision B", DNA 2830F-B, Mathematical Applications Group, Inc. (1973).

TABLE I

SYMBOLS

VECTORS

D, D_k	position of (k^{th}) detector
x	collision position
y	position of previous event (source or collision)
ω, ω_k	direction from y to D or D_k
ω_p	direction from y to x
ω_0, ω_{0k}	direction from x to D or D_k

SCALARS

s	$\lvert x-y \rvert$
r, r_k	$\lvert x-D \rvert$ (or $\lvert x-D_k \rvert$)
R, R_k	$\lvert y-D \rvert$ (or $\lvert y-D_k \rvert$)
ρ, ρ_k	"critical" radius (always $\leq R$ (or R_k))
α_m, α_{mk}	$\sin^{-1}(\rho/R)$ (or $\sin^{-1}(\rho_k/R_k)$)
α, α_k	angle between ω (or ω_k) and ω_p
r_0, r_{0k}	$R\sin\alpha$ (or $R_k \sin\alpha_k$)
S_0, S_{0k}	$R\cos\alpha$ (or $R_k \cos\alpha_k$)
θ, θ_k	angle between ω_0 (or ω_{0k}) and ω_p

FUNCTIONS

$g(\omega_0)$	density function of angular distribution: scattering - $g(\omega_0) = \sigma(\theta)$, where $\sigma(\theta)$ is differential cross-section; source - prescribed, usually $1/4\pi$ (isotropic)
$\lambda(u,v)$	distance from u to v in mean free paths
μ	total macroscopic cross-section

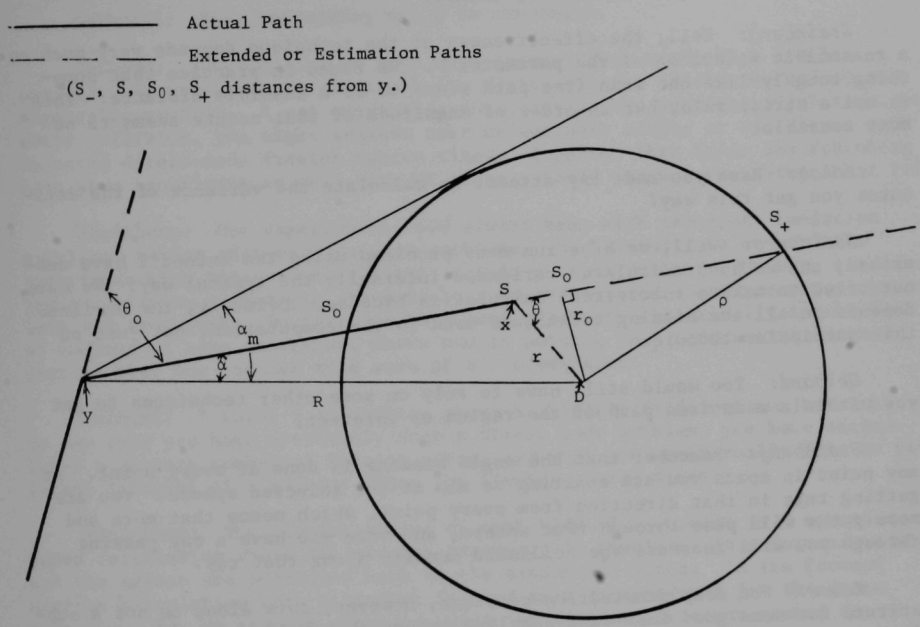

Fig. 1. Geometry for point detector.

DISCUSSION

Gelbard: What you are talking about seems to be, in some sense, equivalent to splitting continuously, with some distribution of splitting boundaries.

Kalos: No, it is better than splitting because the biasing of the angle makes you able to find the neighborhood of the point you are interested in without just having to reach it by chance.

Steinberg: Well, the effectiveness of the technique depends very much on a reasonable selection of the parameter ρ. We found in practice that something roughly like one mean free path seems to be a sensible distance. This is not a strict rule, but an order of magnitude of that nature seems to be most sensible.

Kalos: Have you made any attempt to calculate the variance of the estimates you get this way?

Steinberg: Well, we have run many problems using the method I have described, and we have calculated variances internally the typical way. We have not tried to make a theoretical calculation because, obviously, the variance depends on all the biasing techniques used in the computation, not just on this particular technique.

Gelbard: You would still have to rely on some other techniques to get you within a mean free path of the region of interest?

Steinberg: Remember that the angle biasing is done at every point. At any point in space you are starting to aim at the selected sphere. You are getting rays in that direction from every point, which means that more and more paths will pass through that sphere, and once you have a ray passing through you will increase the collision density along that ray.

Kalos: For deep penetration problems, however, this alone is not a substitute for some good important sampling scheme that brings you to within a few mean free paths of the point in question.

Steinberg: Well, I was thinking again of the small sample problem not ...

Kalos: In the small sample problem, if you are within a few mean free paths of the sample, this will do you a lot of good, I believe.

Cashwell: In the case of the point detector you turn this on when you get reasonably close to the detector?

Steinberg: The angle biasing takes place throughout the entire space. The cones get smaller and smaller as you get further away, but the biasing is always present. The position biasing is only put into effect when you are within that one mean path sphere.

Cashwell: One other question. Using the latest scheme, the Steinberg-Lichtenstein scheme... I gather you pick up some speed in implementation over the earlier methods?

Steinberg: There are practical programing problems in the earlier methods. The first problem by which we were hit was the almost impossible problem of retrieving the cross section for the previous collision. That cross section might no longer be in the fast memory and this was potentially a horrendous mess. Further, in our codes, the tracking portion is the most expensive and time-consuming portion and using the older scheme we have to go back and retrack. This way we avoid those two problems.

Cashwell: Yes, retracking would be expensive.

Gelbard: I would like to pick up a comment that was made a little earlier. We were talking about how one might join the adjoint method to an eigenvalue problem. I was suggesting that, during the adjoint calculation, at every collision, you might somehow make an estimate either at all, or at some selected direct-mode fission source sites. I gather that Kalos and Steinberg have some experience in this sort of thing. Could you say something about it?

Steinberg: Our experiences have always been with transport radiation problems, rather than with reactor problems. As a result, we do not have experience with eigenvalue problems in this sort of context. We have been using our biasing methods as concentrators in various and sundry transport problems and they work. In fact, we have used them very specifically in the kind of context we have described, where one is putting in a nonexisting detector just so that one can use this sort of a concentrator.

Gelbard: I think I am talking about a slightly different subject. Let us say that you have previously done a direct mode problem, you have accumulated a certain number of fission sites, and what you would really like to get is the integral of ψ adjoint times the fission source density.

Steinberg: The closest analog of that kind of a problem is a secondary gamma calculation where the neutron transport is simulated in a forward mode and the gammas are adjointed back to the production sites. In the forward neutron calculation this procedure is used as a concentrator for the gamma detector. So, we have an enhanced collision density, or gamma source density, in the neighborhood of the detector.

Gelbard: But, if you were trying to get the value of the adjoint flux at every one of 10,000 source sites, that would clearly be unfeasible. So that if you want to integrate the adjoint flux times the direct mode source, you could not on every collision get a contribution at every point.

Steinberg: No, what we do is to accumulate what we call a "forward source tape" and each adjoint gamma history contributes to just one source point.

Kalos: If you had at your disposal a large number of fission positions, I think that a sensible procedure might be to select one at random from some distribution which represented the importance function of the sites. That would then be the point at which you make an estimate. But you have the problem of organizing the collection of fision sites so that you can efficiently pick one that is worthwhile. I think that could probably be done. The procedure might well work. It would depend a lot on how well you could do the point estimation.

Gelbard: You have no direct experience with such a technique?

Kalos: Well, the way that we did it was rather clumsy. It worked. But how it might compare with some other hypothetical scheme is very hard to predict.

SESSION V

DIRECTION-DEPENDENT EXPONENTIAL BIASSING

R. C. Bending

Berkeley Nuclear Laboratories
U. K. Central Electricity Generating Board
Berkeley, Gloucestershire, England

DIRECTION-DEPENDENT EXPONENTIAL BIASSING

R. C. Bending

Berkeley Nuclear Laboratories
Berkeley, Gloucestershire, England

ABSTRACT

When Monte Carlo methods are applied to penetration problems, the use of variance reduction techniques is essential if realistic computing times are to be achieved. This paper describes a technique known as direction-dependent exponential biassing, which is simple to apply and therefore suitable for problems with difficult geometry. The material cross section in any region is multiplied by a factor which depends on the particle direction, so that particles travelling in a preferred direction 'see' a smaller cross section than those travelling in the opposite direction. A theoretical study shows that substantial gains may be obtained, and that the choice of biassing parameter is not critical.

The method has been implemented alongside other importance sampling techniques in the general Monte Carlo code SPARTAN, and results obtained for simple problems using this code are included in this paper.

INTRODUCTION

Importance sampling techniques enable considerable savings to be made in Monte Carlo calculations. In tractable geometries, an adjoint transport or diffusion calculation may be performed to obtain a near-optimum importance function. In more complicated geometries, however, it is often impossible to calculate an importance function, and a simpler approach is required. This paper gives a theoretical analysis of a biassing system which is suitable for shielding calculations, and which has been found very simple to apply.

In a shielding problem, the aim of a biassing system should be to encourage particle flights towards regions of low flux or of particular interest and to discourage flights towards the source. In the system described in this paper, which was developed for use in the Monte Carlo code SPARTAN, this aim is achieved by introducing a factor on the

material cross-section which depends on the flight direction relative to a preferred direction. As a result of this, particles travelling in the preferred direction see a smaller cross-section than those travelling in the opposite direction. A theoretical study based on a notional slab system shows that considerable gains may be achieved by this method, and that the choice of biassing parameters is uncritical.

The method has been implemented alongside other variance reduction method in the Monte Carlo code SPARTAN. A number of methods are available for the specification of preferred directions, allowing a fixed direction, a direction which is a function of position, or a direction which encourages flights towards a small target region. The paper includes a brief description of the results of some test cases using direction-dependent exponential biassing in SPARTAN.

EXPONENTIAL BIASSING

Consider a particle travelling in a material of macroscopic cross-section Σ. The distance to first event is obtained by sampling from the frequency function

$$p(x)dx = \Sigma\exp(-\Sigma x)dx \tag{1}$$

In the biassed case, we sample instead from the frequency function

$$q(x)dx = a\Sigma\exp(-a\Sigma x)dx \tag{2}$$

where 'a' is a positive biassing parameter. In order to compensate for the physical distortion caused by the introduction of 'a', we assign to a particle suffering an event at distance x a weight

$$w(x) = \frac{1}{a}\exp\left(-\Sigma x(1-a)\right), \tag{3}$$

so that

$$q(x)w(x) = p(x) \tag{4}$$

Suppose that the particle flight path emerges from the material at distance X. On reaching X in the unbiassed case, the neutron will have statistical weight 1.0 (i.e. will not have suffered an event) with probability

$$\exp(-\Sigma X) \tag{5}$$

and zero weight otherwise. In the biassed case, the particle reaches X with probability

$$\exp(-\Sigma aX) \tag{6}$$

and it is assigned at X a weight

$$\exp(-\Sigma(1-a)X) \tag{7}$$

so that the expected weight at X is preserved.

DIRECTION-DEPENDENT CHOICE OF BIASSING PARAMETER

In principle, the biassing parameter 'a', introduced in the previous section, may have any positive value, and its value may vary from flight to flight. A number of methods for the optimisation of 'a' have been proposed [1,2,3]. In this paper a method is examined in which the value of 'a' depends on the particle flight direction. This approach seems promising in view of the fact that the aim of importance sampling techniques in shielding calculations is to encourage flights towards regions of low flux.

The value of 'a' is calculated as

$$a = 1 - bu \tag{8}$$

where u is the cosine of the angle between the flight path and some preferred direction and b is a constant. This formula is chosen for its simplicity; any expression which is monotonic decreasing in u and positive in the range of u could be used, and no attempt is made in this paper to establish a 'best' formula.

In the appendix, random walk theory is applied to a slab problem in order to find the optimal choice of b (i.e. the value of b which minimises the variance of the estimated penetration through the shield for a given number of particle flights). It is found that the optimal value is only weakly dependent on shield thickness and scatter probability, so that the method is easy to apply.

CHOICE OF PREFERRED DIRECTION

The Monte Carlo code SPARTAN, in which the biassing system described in this paper has been implemented, provides three options for the choice of preferred direction in any region of the system being studied.

Fixed

The preferred direction is defined by a fixed vector (u_o, v_o, w_o). This option is used in slab-like geometry.

Variable

The preferred direction depends on position within the region as follows. If the particle is at (x, y, z) in Cartesian co-ordinates, then a vector in the preferred direction is given by

$$(u_o, v_o, w_o) = (\alpha x, \beta y, \gamma z)$$

for constants (α, β, γ) specified for that region. For example, $(\alpha, \beta, \gamma) = (1, 1, 0)$ gives a preferred direction along an outgoing radius of a cylinder.

Towards a Target Area

A point (x_o, y_o, z_o) is specified, and if the particle is at (x, y, z) then a vector in the preferred direction is given by

$$(u_o, v_o, w_o) = (x_o - x, y_o - y, z_o - z)$$

This option is used to encourage flights towards a particular region, such as that occupied by a particle detector.

USE OF RUSSIAN ROULETTE

Like many Monte Carlo codes, SPARTAN uses weighting in lieu of absorption as a variance reduction device. If a particle has an event for which the non-absorption probability is p_s, then, instead of continuing with probability p_s, the particle continues with probability 1.0, carrying a statistical weight of p_s.

When weighting is used in lieu of absorption, the only way in which a particle history can be terminated is for the particle to have an event in a purely absorbing region or to escape through the boundary of the system. In many situations, this leads to the uneconomic use of computing time, and Russian Roulette is introduced to provide an additional means of terminating particle histories. If a particle has a statistical weight w which is less than a specified limit w_o, then with probability w/w_o the particle is allowed to continue, with weight w_o, or else it is stopped.

When direction-dependent biassing is used, particle weights vary as a result of the application of equations 3 and 7, and the effect of the biassing may be reduced by Russian Roulette. For this reason, the use of Russian Roulette is delayed until after a specified number of events for each particle. It has been found in practice that the value of this number is important for the successful application of the technique to shielding problems. An appropriate value is the mean number of flights to escape, given by equation A15 of the appendix. Figure 1 gives the number of flights to escape per unit shield thickness (in mean free paths) as a function of the biassing parameter b. This value, is, of course, specific to the semi-infinite slab geometry used in the appendix, but is a useful approximation in more general situations.

RESULTS

Figures 2 and 3 give results obtained by the use of direction-dependent exponential biassing in SPARTAN. Both calculations are of neutron flux in one energy group in a symmetric slab, with data approprate for thermal neutrons in sodium (figure 2) and thermal neutrons in water (figure 3). The results of similar calculations without using biassing are also shown, together with the exact analytic results for these simple problems. For reasons of clarity, error estimates are included only near the edge of the shield. Each case was run for a computing time of about 10 minutes.

For transport through sodium (figure 2) in which the probability of absorption at each event is about 0.1, biassing improves the results throughout the shield, particularly beyond 180 cms, where the unbiassed solution collapses rapidly. In water, where the probability of absorption at each event is about 0.006, a shield with the same total attenuation must have a far greater optical thickness, so that good results are more difficult to obtain than for sodium. Nevertheless, a substantial improvement is obtained in the biassed case, and the collapse of the solution, characteristic of the unbiassed calculation, is largely averted.

CONCLUSIONS

The method of direction-dependent exponential biassing has been shown to lead to a substantial improvement in results for simple penetration calculations. It appears that the improvement is brought about principally through a reduction in the number of flights required for a particle to traverse a given shield thickness. The method is used in conjunction with weighting in lieu of absorption and the delayed application of Russian roulette. Approximate theoretical studies indicate that the choice of biassing parameter is not critical, so that the application of the technique to more general problems is straightforward.

This technique, which has been implemented in the Monte Carlo code SPARTAN, is expected to be of value in geometries where a detailed importance function cannot readily be evaluated.

ACKNOWLEDGEMENT

This paper is published by permission of the United Kingdom Central Electricity Generating Board.

REFERENCES

1. L.B. LEVITT, "A Simplified Monte Carlo Approach to Deep-Penetration Problems", Trans. Am. Nucl. Soc., 7, 44 (1964).

2. L.B. LEVITT, "The Use of Self-Optimised Exponential Biassing in Obtaining Monte Carlo Estimates of Transmission Probabilities", Nucl. Sci. Eng., 31, 500 (1968).

3. R.H. KARCHER, R.C. ERDMANN and O.C. BALDONADO, "The Application of Track-Length Distribution Biassing in Monte Carlo Deep-Penetration Calculations", Nucl. Sci. Eng., 31, 492 (1968).

4. D.R. COX and H.D. MILLER, "The Theory of Stochastic Processes", Methuen, London (1965).

APPENDIX - OPTIMISATION OF THE BIASSING PARAMETER b

In order to study the effect of varying the biassing parameter b, we shall consider the special case of a slab shield which is infinite in the -x direction and of thickness d in the +x direction. An isotropic source of particles is situated at the origin. Particles are scattered

isotropically at each event, and the statistical weight of a particle (which is 1.0 at birth) is multiplied by the non-absorption probability p_s at each event, as well as by any factor arising from the biassing process (that is, we use weighting in lieu of absorption). The physical macroscopic cross-section of the shield material is Σ, and a biassing parameter 'a' is used to give a biassed cross-section $a\Sigma$.

If u is the cosine of the angle between the flight path and the x-axis, then the biassing parameter 'a' is given by

$$a = 1 - bu \tag{A1}$$

We shall treat the problem as a random walk defined in one dimension by

$$x_{n+1} = x_n + x \tag{A2}$$

where x_n is the distance of the particle from the x = 0 plane after n events, and where x has a frequency function f(x) with positive mean μ and variance σ^2.

Values of μ and σ^2 may be obtained as follows. Referring to figure 4 the frequency function of ℓ for a given value of u is

$$p(\ell|u)d\ell = (1-bu)\Sigma\exp(-\Sigma(1-bu)\ell)d\ell \tag{A3}$$

$$p(x|u)dx = p(\ell|u)\ \left|\frac{d\ell}{dx}\right|dx$$

$$= \frac{1}{|u|}\ \Sigma(1-bu)\exp(-\Sigma(1-bu)x/u)dx, \tag{A4}$$

where x and u have the same sign. Integrating over the frequency function of u,

$$f(x)dx = \frac{\Sigma}{2}\exp\ (b\Sigma x)\ \{E_1(\Sigma|x|)\mp bE_2(\Sigma|x|)\}dx \quad \text{for } \begin{cases} x > 0 \\ x < 0 \end{cases} \tag{A5}$$

where $E_i(q)$ is the i-th exponential integral
Direct integration yields the moment generating function of x,

$$f^*(\theta) = \int_{-\infty}^{+\infty} \exp(\theta x)f(x)dx \tag{A6}$$

$$= \frac{\Sigma}{2}\left\{\frac{2b}{b\Sigma+\theta} - \frac{\theta}{(b\Sigma+\theta)^2}\ \log\left\{\frac{\Sigma-b\Sigma-\theta}{\Sigma+b\Sigma+\theta}\right\}\right\} \tag{A7}$$

Expanding in powers of θ gives for the mean and variance of x

$$\mu = \frac{1}{\Sigma b}\left\{\frac{1}{2b}\ \log\left\{\frac{1+b}{1-b}\right\} - 1\right\} \tag{A8}$$

$$\sigma^2 = \frac{2}{\Sigma^2 b^2}\left\{\frac{2-b^2}{1-b^2} - \frac{1}{b}\ \log\left\{\frac{1+b}{1-b}\right\}\right\} - \mu^2 \tag{A9}$$

The variable x_n has mean $n\mu$ and variance $n\sigma^2$. As we are mainly concerned with large values of n, we may invoke the central limit theorem to say that x_n has a normal distribution, and we shall where necessary regard f(x) as a normal frequency function.

Now let N be the number of flights to escape. We may apply Wald's identity [e.g. reference 4, Ch.2, equ.74]:

$$E(\exp(-\theta x_N)\left\{f^*(\theta)\right\}^{-N}) = 1 \tag{A10}$$

where $f^*(\theta)$ is the moment generating function of x, as above, and E stands for 'expected value of'.

$$\text{Thus } E(\exp(-\theta x_N - N \log f^*(\theta))) = 1 \tag{A11}$$

Now, $\log f^*(\theta)$, which is simply related to the cumulant generating function of x, is equal to (treating f(x) as normal)

$$-\mu\theta + \tfrac{1}{2}\sigma^2\theta^2, \tag{A12}$$

so that

$$E(\exp(-\theta x_N + N\mu\theta - \frac{N}{2}\sigma^2\theta^2)) = 1 \tag{A13}$$

Expanding the exponential and collecting powers of θ, (noting that since μ is positive eventual escape is certain and

$$E(x_N^r) \simeq d^r), \tag{A14}$$

we have:

$$E(N) = d/\mu \tag{A15}$$

$$E(N^2) = \frac{2}{\mu} E(Nx_N) - \frac{d^2}{\mu^2} + \frac{d\sigma^2}{\mu^3} \tag{A16}$$

Now, since escape is certain, and always occurs at $x_N \simeq d$, N and x_N are independent, and

$$E(Nx_N) = E(x_N)E(N) = d^2/\mu, \text{ so that} \tag{A17}$$

$$E(N^2) = \frac{d^2}{\mu^2} + \frac{d\sigma^2}{\mu^3} \tag{A18}$$

Now, consider a particle which escapes on its Nth flight. At each collision, its weight is multiplied by p_s and by a factor given by equation 3, which depends on the flight length and on the value of the biassing parameter 'a'. This factor, which is not independent of N, is not easy to calculate. We shall assume that the score arising from this particle has the functional form

$$S_N = r^{N-1} \tag{A19}$$

and we shall treat r as independent of N for the purpose of evaluating the mean and variance of the score. Since we know that the mean score should

be independent of b, and since $r = p_s$ for zero b, we can find an appropriate value of r for each value of b. The calculation of the score for zero b is described later.

In order to find the mean and variance of S_N, we need to know the distribution of N. The mean and variance of N are given by equations A15 and A18, and we regard N as a continuous variable having a gamma distribution defined on $1 \leq N < \infty$. This is somewhat arbitrary choice, and is made only for computational convenience and because the gamma distribution has the right type of domain. Thus the frequency function of N is

$$f(N) = \frac{1}{\beta\Gamma(t)} \left(\frac{N-1}{\beta}\right)^{t-1} \exp\left(-\frac{N-1}{\beta}\right), \tag{A20}$$

with mean $1 + \beta t$ and variance $t\beta^2$.

In this case, using equations A15 and A18,

$$\mu_N = d/\mu = t\beta + 1 \tag{A21}$$

$$\sigma_N^2 = d\sigma^2/\mu^3 = t\beta^2 \tag{A22}$$

The mean score is thus

$$\overline{S} = \int_1^\infty f(N) r^{N-1} dN \tag{A23}$$

$$= (1 - \beta \log r)^{-t}, \tag{A24}$$

after a little algebra. Similarly,

$$\overline{S^2} = (1 - 2\beta \log r)^{-t} \tag{A25}$$

It is desirable to minimise the relative variance of the score, $(S^2 - \overline{S^2})/\overline{S^2}$, and also the mean number of flights μ_N, by varying b. We therefore seek a minimum of

$$C = \mu_N \left\{ \frac{(1-\beta\log r)^{2t}}{(1-2\beta\log r)^t} - 1 \right\} \tag{A26}$$

There remains the problem of finding the score for b = 0. In this case $\mu = 0$ and the moments of N are not defined. In order to ensure consistency in the mathematics, we introduce a second shield boundary at $x = -d$, and argue on physical grounds that the score for the one-sided shield will, for sufficiently large d, be half that for the two-sided shield.

In this case, we obtain

$$\mu = 0 \tag{A27}$$

$$\sigma^2 = 2/(3\Sigma^2), \tag{A28}$$

and a calculation analogous to that given in equations A10 to A18 yields the results

$$\mu_{N(2)} = d^2/\sigma^2 \tag{A29}$$

$$\sigma^2_{N(2)} = 2d^4/3\sigma^4 \tag{A30}$$

where the suffix 2 indicates the two-sided shield. In this case $r = p_s$, so the score for the one-sided shield is approximately

$$\overline{S} = \tfrac{1}{2}(1 - \beta \log p_s)^{-t} \tag{A31}$$

where β and t are given by

$$1 + \beta t = \mu_{N(2)} \tag{A32}$$

$$\beta^2 t = \sigma^2_{N(2)} \tag{A33}$$

This value of the score should apply for all b, so that from equation A24 we may write

$$r = \exp\left\{\frac{1}{\beta}\left(1 - \exp\left(-\frac{\log \overline{S}}{t}\right)\right)\right\} \tag{A34}$$

This value of r may then be used in equation A26 to estimate C.

The value of C (the parameter which is to be minimised) has been computed for various values of the other parameters. Figure 5 gives an example of the results obtained. These show that any value of b larger than about 0.2 leads to a substantial reduction in C, and that the position of the minimum is not closely dependent on p_s. Similar results apply for other values of Σd.

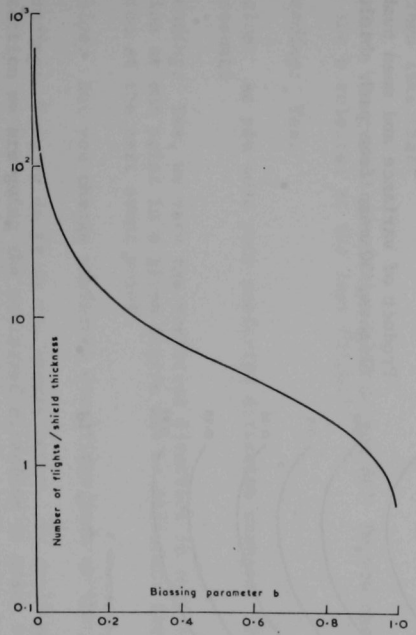

Fig. 1. No. of flights to escape.

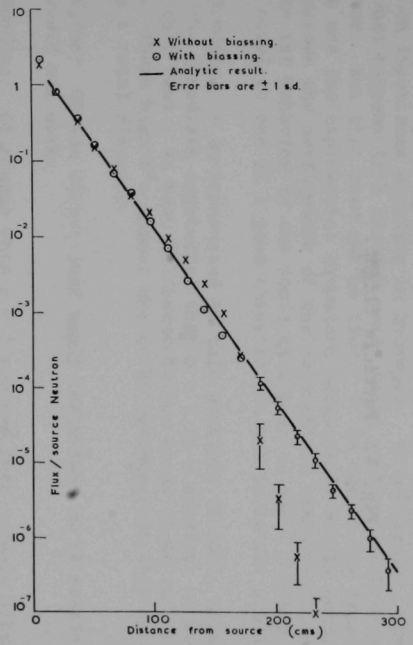

Fig. 2. Thermal neutrons in sodium: 1-group slab calculation.

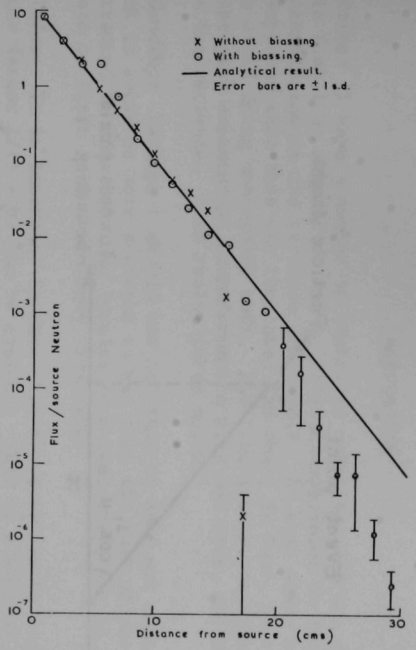

Fig. 3. Thermal neutrons in water: 1-group slab calculation.

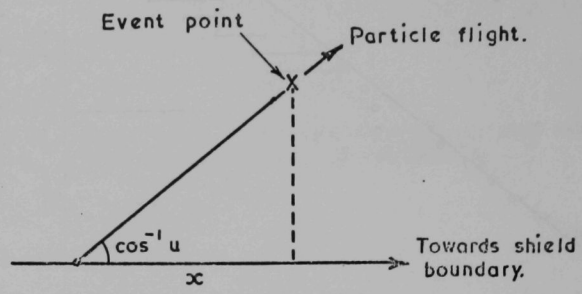

Material cross-section regarded as $\Sigma\,(1 - bu)$

Fig. 4. Particle flight.

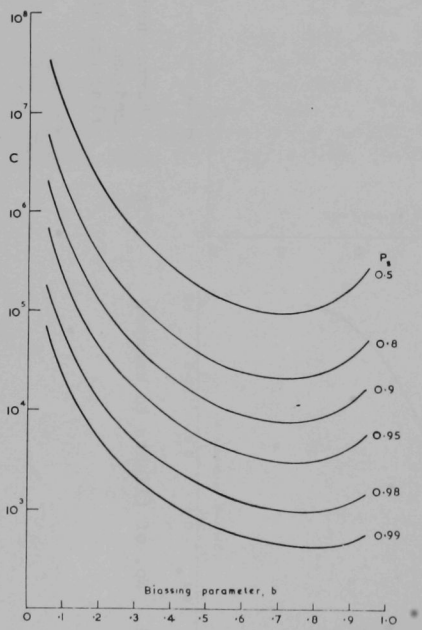

Fig. 5

Product of variance and mean number of flights, 30 mean free path shield.

DISCUSSION

Kalos: I have a number of technical questions and one comment. The comment is that I am interested and a little disappointed to discover that you do not actually reduce the variance per sample. In some naive schemes I had always supposed that this was the case, and that, in fact, what you did was simply to avoid doing unncessary work. But I am really surprised to discover that the pure exponential transformation does not work, and that it does not reduce the variance. I am really taken aback.

Bending: Perhaps I should not say that it does not improve the variance. I think it would be more accurate to say that the bulk of the improvement is not attained in this way. My results show that there was an improvement in variance at first as b moved away from zero.

Kalos: It is important to note that the increase in variance that comes from the factor $1/(1 - b)$ can be avoided. That is a kind of a spurious effect that really does not have to occur. I discovered this problem in a very old paper on importance sampling in hydrogen. In particular, you might ask yourself what happens if b becomes one? Using your approach that would be a total disaster. But that need not be true. It is not in the least true that when b equals one the exponential transform necessarily breaks down. So I think you can improve the performance of the exponential transformation by paying attention to the behavior in the limit as b goes to one. I don't suggest that using b equal to one is a good thing to do; but it is not a disaster area.

Bending: I am interested in this problem; in fact, because the equation gives the immediate impression that $b = 1$ is a diaster, we have simply eliminated that case. We have deliberately run cases close to one to see what happened. These high b's appear not to be optimum, but they do not appear to indicate a total disaster.

Kalos: Oh you cannot just apply the method as it stands with $b = 1$ — that would not work.

Bending: Of course when $b = 1$ you run into trouble like dividing by zero, and things like that.

Kalos: In curved geometries what do you do? Do you choose a path according to the b selected at the last flight?

Bending: Yes.

Kalos: As you move your preferred direction changes. Do you take that into account?

Bending: Yes, we vary the preferred direction in that the preferred direction at one point in a given region may be different from the preferred direction at the next event point.

Kalos: But you change preferred directions just as the events occur?

Bending: Yes, we do it as the events occur. We don't anticipate that in the direction we are going the preferred direction is going to be changing. We only consider the preferred direction at the last event point.

Kalos: Thank you.

Steinberg: I would like to comment just a little bit on some historical experience. About 15 years ago I was doing some one-dimensional gamma-ray calculations, and this is how I first got into Monte Carlo. At that time we started using the exponential transform and we quickly discovered that it would be quite valuable to make b position and energy-dependent. In fact, when we were doing one deep penetration problem we found that the closer you got to the escape surface the closer you wanted b to be to one. Soon afterwards we decided that the "bu" term, specifically, could be replaced by an arbitrary function. More or less by trial and error we tried to decide on some kind of reasonable optimum function of direction, energy, and position; and we found we got very successful results in the deep penetration problem.

Bending: We haven't explored the use of formulae other than 1 - bu and it has been obvious to us throughout that this was a totally arbitrary choice. It was justified by the moderate gains which it produced. But we would completely accept that there is undoubtedly a lot of room for improvement here.

Steinberg: Actually once you put in a general form for the "bu", you can show that this is equivalent, mathematically, to some biasing function in which the exponential form happends to be the form you are using. But you can use a different biasing function which will imply a more general form of the exponent in the exponential transformation.

Bending: We are principally interested in a method which can be applied with very little insight into the problem being solved, and without doing anything like an adjoint calculation. Now, if the more general exponential transformation has these properties, and can give better results without penalizing the user, without forcing the user to really understand the method, then this would be something in which we would be very interested. Sometimes separately, and sometimes along with the exponential transform, we use importance sampling based on splitting and Russian Roulette in the normal way. But we find that splitting and Russian Roulette can give completely wrong results, through misuse, more easily than our method. So, to some extent, political questions are involved in the choice of a method.

Steinberg: Yes, I think what you are saying is that the code designer, if he invests enough time and effort, can build in reasonably optimized importance functions for a whole wide class of problems. But unfortunately, in really complex codes, the code designer does not have complete control over what the code will ultimately be used for, and has to allow the user to put in importance function for untested situations.

Bending: Yes.

Kalos: I want to point out that Steinberg's experience, although it is very valuable, was for a one-dimensional deep gamma-ray penetration problem and he spent a year experimentally optimizing the importance functions. I would say that that was a discouraging report, rather than an encouraging report, from Dr. Bending's point of view. I would like to make one technical remark. In the criticality problem I believe that importance sampling is a valuable technique; but the importance function that one must use is not the probability of getting back to an important region in the current generation,

but the probability of getting back asymptotically, after a lot of generations. This should tend to influence your use of importance sampling to gather the nuetrons back from the moderator.

Bending: We have found experimentally that one cannot try to push this kind of approach too hard in criticality calculations.

Cashwell: Now, in slab geometry presumably your preferred direction does not change very much as you penetrate the slab. Do you have any experience in complicated geometry where the preferred directions change rather drastically as you go, in cases where you try to lead the particles down some tortuous path?

Bending: We have no experience of this sort. In fact in all the cases we have done in any geometry the preferred direction has not changed very much within several mean-free paths; and, I think, if one were in a situation in which the preferred direction changed radiacally, changed by a right angle, within roughly one or two mean-free paths, then I think that one would find this approach much less effective. But I am guessing. We have had no experience in that situation.

Bernat: I have a question with respect to criticality problems. It may be that in the case of a water reflector this importance sampling is a very good method. But in the case of beryllium or graphite or D_2O reflectors, I am not sure that this is a reasonable method.

Kalos: Did you say D_2O? Well, I have absolutely no experience with such a case, and I hesitate to speak to on this question. But I am sure that, in principle, savings can be obtained. I can only speculate that perhaps your negative experience is based on trying to do too much, which is easy to do. Also, I would strongly recommend that anybody who uses biasing, splitting, and Russian Roulette do it the way we do it at MAGI in the SAM codes, and this is described in an appendix to my paper called, "On the Integration of the Adjoint Gamma Ray Transport Equation." There is an algorithm there which, for the flight, is the minimum variance splitting and Russian Roulette scheme that one can devise.

Gelbard: Do you have another kind of identification for that paper? Is there a report number?

Kalos: *Nuclear Science and Engineering*, 33, 284 (1968). There is a little appendix in that paper that happens to be rather clear, for a change, and which describes the minimum variance splitting and Russian Roulette technique.

Gelbard: Let me get onto a somewhat different subject. I noticed in your last slide that the truth was pretty far outside the error bars, and I wonder how serious a problem this is. I refer again to the problem mentioned many times before, namely, that when you use a very severe kind of biasing it is very difficult to estimate the error bars. I remember a case where somebody I know was doing a shielding calculation, got very nice standard deviations, but later found out that he was a factor of 5 away from the truth. Now, you were specifically talking about the problem of constructing a foolproof system, and isn't this a difficulty in constructing a foolproof system? The user has, at this point, no warning of any sort of trouble. I want to ask, in

this connection, whether a test of normality has been tried, and whether it has been found to be useful in giving the user a warning as to the validity of the error estimates.

Moore: It is our experience, in shielding calculations in general, that it is very easy to generate error bars which are not in any physical sense a realistic estimate of the error. We don't find that this is necessarily associated with a particular biasing scheme. It seems to be a property of deep penetration calculations in general. We have put some effort into this problem, but have not succeeded in finding a method of interpreting results which will guarantee that the error bars are reasonable. In fact, we have done what you suggested. We have put in some forms of normality testing and what we find is that, in shielding in particular, there are quite often situations in which the normality test indicates that the results are not good, when in fact they are quite satisfactory. But, when the results are not satisfactory, then almost universally the normality test tells us that they are not satisfactory.

Gelbard: What test do you use?

Bending: It is a Shapiro-Wilkes test based on the order statistics.

Coveyou: I am a little puzzled by all this because I have always been of the opinion that these distributions, except in very rare cases, are just not normal, and that the whole theory that is used to calculate confidence intervals and so on just simply does not apply here.

Kalos: That is a different question.

Coveyou: I know, but the point here is that there is a simple way to get error estimates which are rigorous.

Kalos: But it is conservative.

Coveyou: Yes, very conservative, just Chebycheff's inequality. If the error bars calculated from Chebycheff's inequality do not lie near a known truth, then you know that something is wrong.

Gelbard: But I think you are liable to have a very bad psychological effect on the user if you quote such conservative error bars.

Coveyou: I am not entirely sure that the effect is bad if you warn the user his bad results are bad.

Gelbard: I mean that it may have a bad effect, from the public relations point of view, if you quote errors which are so conservative that they unnecessarily discourage the user.

Coveyou: Well, you can explain what they mean. These are not error bars. These are not 95% confidence intervals for normally distributed samples. You can say what the error bars mean in terms of probability.

Gelbard: Incidentally a remark was made to the effect that this behavior is characteristic of shielding problems and I wanted to say that I agree. But I think that the reason this is characteristic of shielding problems is that,

if you do no biasing, practically no particles get through. Then you are certainly in trouble. And if you do biasing then the biasing gets particles through, but your sample distributions are non-normal.

Kalos: I cannot agree with the thesis that pathological confidence intervals are characteristic of biasing schemes. No biasing is also a biasing scheme. When a biasing scheme gives trouble it means that you need better biasing. I think normality is almost a sufficient but not a necessary test. Anyway, I would like to suggest a more practical approach. That is, that within practical limits, you vary the b's and see what happens in a practical problem. You instruct the user to run his problem twice with different b's and see what the differences are. This would be very illuminating.

Bending: There is one thing we try to do to avoid trouble. We use Russian Roulette to terminate histories when we are not interested in them any longer. But we delay the application of Russian Roulette until the particle has had at least μ_n flights, where μ_n is the mean number of flights to escape; and preferably we allow it to have rather more than this many flights before we apply the Russian Roulette. We do this when the situation is such that one can calculate μ_n, and then add something as a safety margin. We find that under such circumstances we normally do not run into serious problems with error bars being unrealistic. But this tends to be, to some extent, a slightly handwaving exercise in that it is only easy to use such an approach in a one-group slab shield. On the other hand, in a problem configuration with any degree of complexity, it becomes difficult to insure that one actually has an appropriate value for μ_n, unless one puts in a lot of conservatism in which case the program runs more slowly. But we have found that if one's μ_n's are sufficiently high, if you force the neutron to continue so that, virtually, its only way to terminate its history is to physically escape from the system, then in general, this avoids the problem of uninterpretable errors.

Gelbard: Is it correct, as Coveyou suggested, that you almost never find that the normality test is satisfied?

Bending: This is to a great extent true, and this is why I said that very often the results which the normality test throws out are, in fact, physically acceptable. I have toyed with various ideas which should enable one to draw conclusions from a series of results which are manifestly not drawn from a normal distribution, but I have not succeeded in finding a scheme which does better than simply to give you a warning and tell you that something is wrong. I would value someone else's experience in interpreting results which are obviously drawn from a very strongly skewed distribution.

Gelbard: In cases where you don't use biasing, but are examining an unlikely event, characteristically you tend to underestimate the probability of that event. Now, there was a suggestion here that by varying b you would see whether you are in trouble or not. I am wondering whether, with exponential biasing, you have any such a characteristic behavior. In the cases which you described, when the truth was outside your error bar, you tended to be below the truth. If that is true characteristically, then varying b will still, certainly, give you some information. But you might vary b and find what looks like a satisfactory situation, when in fact, with all the b's you have tried, your answers are systematically low. So there is still some danger. There does not seem to be any really decent way around this problem. I mention this

again in connection with the desire to make Monte Carlo foolproof. It seems to me that this is still a nonfoolproof aspect of Bending's method, and of just about any method designed to do this sort of difficult problem.

Kalos: It is easy to construct hypothetical problems in which this scheme will make unlikely the most important events. Consider any situation with a strong and important minimum in the cross sections at an energy well below the energy you are starting at. By biasing forward you are not giving the neutron the time it needs to slow down into the minimum and then stream through the window. You need an energy biasing but, in fact I think that this exponential biasing scheme even goes slightly in the wrong direction in such a situation. It is not, necessarily, absolutely worse than no biasing at all but I think that it goes in the wrong direction.

Gelbard: In some cases you may know that there is some special reason why you should not use a particular biasing scheme. On the other hand, you may not be aware that there is a very important cross-section window in your problem, in which case you might systematically get the wrong answer by using the biasing, and not realize that you are doing anything wrong.

Bending: In the kind of situation that you are speaking of, where the systematic error is induced by a particular form of the cross sections, I suspect that there would be no statistical effect which could be picked up by any statistical test, whether a normality test or any other kind. This would be a particularly difficult problem to detect.

Gelbard: Now I want to turn to a different subject. Yesterday I was arguing that Monte Carlo methods ought to be exact. But, if one is willing to accept approximations, then how useful are albedo methods to get rid of the problem of the reflector? If you have a regularly shaped outer boundary, and the core is very large, it would seem that you could just replace the reflector by an albedo, and this has in fact been done. Has anybody any feeling about the utility of this method? In a thermal reactor the Monte Carlo treatment of the reflector does take an awfully long time.

Whitesides: We have used this approach quite extensively in the KENO program, generally for rectangular reflectors. We assume that the particle comes back in at the point that it leaves the system, and to make this approximation reasonable the reflector boundary must be large. In the simplest form of the albedo method we do not account for corners, although we have actually treated corners, using two-dimensional DOT calculations to generate the albedo information.

Gelbard: Normally, you would get the albedo information from ANISN?

Whitesides: That is correct. We run a series of ANISN calculations. For every energy group we compute the albedo for various incoming angles. We assume that the incoming and outgoing flux can be considered azimuthally symmetric, and find that the rectangular boundary of the core does not have to be very large to make this approximation almost exact. But I want to point out that this method is only valid in certain special cases. To use it to treat reflectors with complicated shapes is not very practical.

Gelbard: Well, how about moving the albedo surface out into the reflector? Perhaps a mean free path or so out?

Whitesides: That helps, but also limits the savings.

Kalos: Do you think half a mean-free path might work?

Whitesides: I guess we actually used a mean-free path or two and felt that we had to move the boundary into the reflector by at least that much. Do you think that is too much?

Coveyou: Isn't that just a geometrical question? The point is that you ought to use as little of the reflector as you can to get a regular boundary around the pseudo-core. How far you move the boundary cannot have any physical significance.

Kalos: It can, because of the approximations you are making. You are returning each reflected neutron exactly at the point where it entered the albedo boundary, instead of smearing the reflected neutrons over a mean-free path or so.

Whitesides: That is right.

Gelbard: Let us say that the geometry of a boundary is roughly cylindrical, but has corners. You can ask how far out in the reflector you have to get before the flux distribution in the reflector really has cylindrical symmetry. If you can go out a mean free path, and if the effect of corners has worn off in that the flux is azimuthally symmetric, then at that point you could safely put an albedo boundary.

Whitesides: Actually what we prefer to do in the reflector is to use Russian Roulette and splitting, with importance functions generated by ANISN slab calculations. We feel that this approach works much better than the albedo method if you have got to move the boundary into the reflector.

MONTE CARLO METHODS FOR SHIELD DESIGN CALCULATIONS

M. J. Grimstone

UKAEA, Atomic Energy Establishment
Winfrith, Dorchester, Dorset, England

MONTE CARLO METHODS FOR SHIELD DESIGN CALCULATIONS

M. J. Grimstone

UKAEA, Atomic Energy Establishment
Winfrith, Dorchester, Dorset, England

ABSTRACT

A suite of Monte Carlo codes is being developed for use on a routine basis in commercial reactor shield design. The methods adopted for this purpose include the modular construction of codes, simplified geometries, automatic variance reduction techniques, continuous energy treatment of cross section data, and albedo methods for streaming. Descriptions are given of the implementation of these methods and of their use in practical calculations.

1 INTRODUCTION

The objectives of the work on Monte Carlo at Winfrith have been to develop methods which can be used by engineers for commercial reactor shield design. The central problem in practical shielding calculations is the determination of radiation attenuation in systems with complicated geometries. Whilst Monte Carlo is the obvious choice of method for treating problems with strong geometric effects, it is not well suited to deep penetration calculations. Two general lines of approach have been adopted to overcome these difficulties: semi-analytic techniques exploiting the well-tried removal and albedo approximations are used for the treatment of streaming phenomena in ducts and voids; and adjoint diffusion theory solutions furnish approximate importance functions for acceleration in bulk penetration problems.

This marriage of the conventional methods which have long been established for shield design with Monte Carlo in "hybrid" calculations has proved to be the deciding factor which has led to the extensive use of Monte Carlo as a shield design tool in the United Kingdom. Thus Dutton, for example, used line-of-sight formulae to calculate neutron sources from AGR fuel clusters for the Monte Carlo treatment of the PCRV penetration-head region [1]. Similarly, Avery and Pugh used two-dimensional Monte Carlo to investigate streaming phenomena in the refuelling machine of the Magnox station at Trawsfynydd [2]. Avery and Warman have employed an albedo-Monte Carlo method in calculating the streaming of neutrons and gamma-rays along air gaps and passageways between the vessel and primary shield of a PWR [3]. The use of diffusion theory in adjoint mode to provide approximate importance functions for accelerating Monte Carlo calculations has been demonstrated by Bendall and McCracken [4].

This paper describes how the methods of coding, data handling and variance reduction have evolved as a result of this emphasis on design applications. The calculational scheme is under continual development to meet the changing needs of reactor programmes; the current status of codes is reviewed, and some indication is given of future trends.

2 GENERAL DESCRIPTION OF THE METHODS

The capabilities outlined above are being provided for shield designers in the form of codes of modular construction [5]. There are two levels of approximation in the basic suite of modules, the Monte Carlo options being supplemented by the fast and approximate methods based on line-of-sight, kernel albedo, and diffusion theory in the various forms used for shielding applications (eg the COMPRASH removal-diffusion method [6] and the Adjusted Diffusion Coefficients (ADC) Method [7]). Apart from their role in the hybrid Monte Carlo calculations these approximate methods furnish a powerful range of tools in their own right for rapid survey calculations and for repetitive runs in shield optimisation studies. The codes and methods developed at Winfrith are complementary to the generalised geometry code MONK which is the subject of another paper at this Conference [8]. This code is used primarily for criticality applications, but it can be utilised for shielding studies where the geometric capabilities are essential for an adequate representation of the problem.

There are a number of special features required for Monte Carlo methods used for shield design. The code system must be flexible so that it is easy to introduce new capabilities as different problems arise. Speed of execution and economical storage requirements are important, and these may even have to be pursued at the expense of some accuracy. It is also essential to reduce as far as possible the degree to which the efficiency and accuracy of the codes depend on the expertise of the users or on preliminary trial runs. Key features which have been developed specifically to meet these requirements include:

(i) modular construction of codes;

(ii) simplified geometries or specially written geometry modules;

(iii) automatic variance reduction facilities;

(iv) a continuous energy treatment of data so that the user is not left with the choice of basic material parameters which affect the accuracy of the calculation;

(v) albedo methods for the solution of streaming problems.

It will be convenient to discuss the shielding code scheme under these five headings. At present, it comprises a series of modules from which a variety of Monte Carlo codes can be assembled. These can be used on a stand-alone basis. Alternatively, they can be coupled to diffusion modules, run in the adjoint mode, for the automatic generation of importance functions in bulk penetration calculations, or to kernel-albedo modules for semi-analytic

solutions of streaming problems, in a range of one-, two- or three-dimensional geometries.

3 THE WRS MODULAR CODING SYSTEM

3.1 Relation to other schemes

While the modular approach to programming is now common practice the particular forms which it takes are very varied, reflecting the special requirements of their users. Basically, the term "modular" implies one of the following approaches:

(i) the assembly of programmes from units or "modules" with data linkages achieved via the fast store;

(ii) the performance of a calculation by means of several separate computer runs with automated data transfers via backing storage.

The Winfrith Radiation Physics and Shielding Group (WRS) System [9] comes into the first category. Programmes assembled from WRS modules are Fortran programmes running under the normal operating system but using standardised coding techniques and standard routines for house-keeping. The aim is to improve the ease and flexibility with which codes can be constructed, especially for the complex series of steps frequently required in a shield design calculation. It is because these calculations include iterative processes, of which the Monte Carlo tracking cycle may be classed as one, that the WRS system concentrates on the construction of single codes rather than automated linkages between separate computer runs as in the second of the basic approaches mentioned above. However, for calculations which can be run efficiently as separate jobs this second approach can be made very powerful through the use of a highly organised data base with special data management routines as, for example, in the COSMOS system [10] which has been developed at Winfrith. The two approaches are complementary since the basic unit in the COSMOS scheme is a code, and this code may be one assembled from WRS modules. Thus, if the shield designers using WRS codes require extensive data banking facilities or automated links with core calculations then these are available through COSMOS.

3.2 Description of the WRS system

A WRS module consists of one or more subroutines. A programme is assembled from one or more modules together with a linking module (a Fortran main programme possibly supplemented by subroutines) which controls the logical flow of the calculation. The ease and flexibility of code construction is achieved through two main features. The first is the convention that a module must be capable of accepting its initial data either from cards or from a general location in core store or on backing store. A similar rule applies to the output, with the line printer replacing the card reader. This provides considerable flexibility since the same module can be used, for example, either on a stand-alone basis (with just a very simple linking module) or as part of a programme in which some or all of its initial data is calculated by other modules. The second feature is the method for the dynamic storage of data in core and on backing store. The data in core resides in a

single array in a common area which is partitioned into "data groups". The first partition is a directory containing information about all the data groups in core and also about those on backing storage. Subroutines are provided to perform such tasks as the allocation of space, the freeing of space no longer required, the communication of data between modules, and the transfer of data groups to and from core store. These routines are responsible for the upkeep of the directory, the details of which need not concern the programmer, who needs only to know how to call the routines.

3.3 Use in Monte Carlo programming

The WRS system is used for all coding in the two option scheme described in Section 2. In the Monte Carlo field it is proving to be of considerable value in the assembly of codes from modules typically having the functions of source sampling, tracking, collision mechanics, splitting, or scoring. Experience was gained earlier when the system was used for writing the prototype XY geometry code McBOX [11] which was developed primarily as a test-bed for variance reduction methods in one and two dimensions (see Section 5). Modular programming necessarily involves some penalties but these can be minimised in a well designed scheme. Three basic issues which must be considered are speed of execution, storage requirements, and programming complexity.

(i) The problem of execution speed is especially critical in Monte Carlo codes, in which modules are entered many thousands of times and any fixed overheads are multiplied correspondingly. Care was taken in the writing of McBOX to avoid these overheads during the generation of particle histories so that the efficiency of the code is not significantly impaired by its modular nature.

(ii) There are two factors which lead to an increase in the storage requirements of WRS programmes. The first is the presence of the system routines which occupy 25 K bytes, and the second is the presence of the full input and output options in every module. However, the latter are often written as separate subroutines so that space can be saved through the use of an overlay. In addition, the dynamic storage system ensures that very economical use can be made of the data storage area by keeping the storage compacted and by removing data when it is no longer required.

(iii) It is a feature of all but the simplest modular systems that the initial task of programming becomes protracted and more complex. This is acknowledged to be true of the WRS system, though by no means to an unacceptable degree.

Experience with the use of the system has confirmed the view that the extra initial coding effort is repaid ultimately during the assembly of large programmes. The ability to use modules on a stand-alone basis has been found useful during the testing stages, which are often especially lengthy and difficult in the case of Monte Carlo codes. The system also provides a very convenient way of linking Monte Carlo with the approximate methods for the types of hybrid calculation described in the Introduction.

4 GEOMETRY

In planning the geometric capabilities of Monte Carlo codes for shield design there are two possible approaches:

(i) to provide a very general three-dimensional capability;

(ii) to rely mainly on the use of simple (though up to three-dimensional) geometries, and to write special geometry routines for specific applications when necessary. (The term "simple geometry" is used here to signify that the material boundaries are surfaces of the form α = constant, where α represents any of the co-ordinates.)

The second of these methods has been adopted in the interests of computing speed and on the basis of past experience which has shown that many practical shielding configurations can be adequately represented by means of simple geometries and that the input data for such representations can be quickly set up.

Thus the modules now written allow the tracking of particles in XYZ and RZ geometries (and therefore also in X, XY or cylindrical geometries), and these can accommodate a wide range of practical design problems. However, in the fast reactor and HTR fields there is a need to represent internal shields and core/breeder configurations in triangular-prismatic geometry with provision for the introduction of cylindrical regions representing fuel channels, control rods or, in the case of a sodium-cooled fast reactor, materials testing rigs containing cylindrical specimens. The McRIG code [12] has accordingly been written for this purpose, initially in two-dimensional geometry in which the horizontal reactor plan consists of a mesh of equilateral triangles on which are superimposed circles which may be nested. Within each triangle the technique of Woodcock tracking [13] is used, in which the collision points are determined using the mean free path corresponding to the highest cross section of any material in the triangle. At each point the ratio of the true cross section to this maximum cross section is calculated, and using this ratio as a probability a choice is made between a real reaction and a pseudo-reaction in which no change of energy or direction occurs. This method has the advantage that it is not necessary to calculate distances to region boundaries within the triangle but only to test whether a co-ordinate is inside or outside a region. There is a danger that this saving of effort will be offset by the generation of large numbers of pseudo-collisions if there are materials present with widely different cross sections. However, this has not proved to be a problem in the application of McRIG to gamma-ray heating in fast reactor cores. This is probably due in part to the particular geometric configurations, and also to the fact that the energies at which widely different cross sections occur are unimportant.

While this approach provides geometry capabilities adequate for a large range of problems it is essential to have a more general code for those which are not amenable to a simplified representation and which are not sufficiently common to warrant specially written geometry modules. Reference has already been made to the general Monte Carlo code MONK developed at Risley. An example of this type of problem occurred in the design of the Prototype Fast Reactor when the general geometry code GEM [14] (the forerunner of MONK) was used in

calculations for the fuel storage rotor which involved both neutron multiplication and neutron penetration [15]. In the field of HTR shield design and in some calculations of gamma-ray heating in fast reactors (such as a partially inserted control rod) there is an increasing demand for three-dimensional calculations in triangular-prismatic geometry to be carried out on a routine basis. It is accordingly proposed to incorporate the Z co-ordinate into McRIG to simplify the task of specifying the geometry in such problems which are at present tackled using MONK.

5 VARIANCE REDUCTION METHODS FOR PENETRATION CALCULATIONS

5.1 The use of a calculated importance function

For use in shield design applications acceleration techniques employing splitting and Russian roulette have the advantages of being relatively simple to apply, of being sufficiently powerful for deep penetration problems, and of being useful in a wide range of applications. The method was implemented in a simple form in the RZ geometry programme McNID [16], in which splitting and roulette are applied at spatial surfaces with a fixed splitting power of two. This approach has the disadvantages that no allowance is made for the energy dependence of the importance function and that considerable initial experimentation is often required to find the best position for the splitting surfaces. A considerably improved method was used by Bendall in the prototype one-dimensional code McBEND [4] in which splitting and Russian roulette may occur at collision points and at points distributed uniformly along a particle track. At each point the degree of splitting is determined by reference to an importance function which may depend on position, energy, and direction. It was demonstrated that the use of an approximate importance function obtained from an adjoint removal-diffusion calculation was very effective in accelerating the calculation of a specific pay-off.

The extension of this method to higher dimensions was not immediately possible owing to the lack of a sufficiently fast two-dimensional removal-diffusion programme. This obstacle has now been removed by the development of a method of adjusted diffusion coefficients (ADC) [7] which eliminates the lengthy calculation of first collision sources, and also by the introduction of improved methods for the solution of the diffusion equations [17]. It is now proposed to investigate the further extension to three dimensions by using the diffusion programme SNAP [18] in adjoint mode to calculate importance functions.

5.2 A self-adjusting acceleration method

A method has been established for circumventing the difficulty of calculating importance functions in the rather different, but not uncommon, class of problem in which uniform accuracy is required over a wide range of energies and positions. This is frequently the case, for example, in a gamma-ray heating calculation in a reactor core, and it was the type of output needed from the one-dimensional reference calculations to which diffusion solutions were fitted in the production of adjusted coefficients for the ADC method mentioned above. The technique, which has been tested in the prototype XY geometry code McBOX, is to arrange for the programme continually to adjust importance values on the basis of the histories so far generated, with the aim

of assisting the penetration of particles to all parts of phase space in roughly equal numbers. The source module of McBOX samples from a distribution $S'(x,y,E)$, and if the true source function is $S(x,y,E)$ it assigns a weight $T = S(x,y,E)/S'(x,y,E)$ to the particle. For the purposes of the splitting and Russian roulette process the phase space of points (x,y,E) is divided into a number of regions, each of which has an importance value denoted by I. It is convenient to regard $W = 1/I$ as the "region weight". Particles have a tracking weight U which is initially 1 and which is separate from the source weight T so that the total weight used in flux estimation is the product of T and U. Splitting or roulette may occur when a particle enters a new region by crossing a spatial boundary, by suffering an energy loss which transfers it to another region, or when it is first generated as a source particle. The expected number of particles produced is equal to U/W, the ratio of the old particle weight to the weight of the new region. The actual number of particles is selected by random sampling from the integers on either side of this ratio, and the weights of the particles (if any) are divided by this integer. This method ensures that the expected values of the particle weights are equal to the region weights themselves, but the weights within any region are of course distributed somewhat because only integral numbers of particles can be produced by splitting and roulette.

The quantity which the programme calculates for use as a region weight is the "current" entering the region. This is a quantity which would be observed with a true source $S'(x,y,E)$ and no splitting. When splitting is in operation the current from a given number of source particles is estimated by $(\Sigma U)/V$ where V is the region volume and the summation is taken over all particles entering the region. Thus, when the weight adjustments have settled down we have

$$W = (\Sigma U)/V$$

and since the particle weights have a mean value W this implies that the number of particles entering any region is proportional to its volume. This is the quantitative interpretation which has been given to the requirement for particles to be distributed evenly throughout phase space. While this criterion probably gives an approximation to the optimum conditions for most problems of this type, it cannot of course always guarantee equal variances in all regions. For example, the errors still usually increase with distance from the source since the tracks tend to be more highly correlated at larger distances. It is also not always possible to apply the splitting method without modification. It has been found necessary, for example, to limit the number of particles produced by splitting to a maximum of 10 in order to avoid inefficiency in problems in which there is the possibility of transfers between regions of very widely differing weights.

Subject to these qualifications the method has proved to be useful, especially in easing the task of the user in this type of problem. It is, of course, essential for efficiency that the weights should be substantially determined within some reasonable fraction of the total running time. This has been achieved in the applications so far tested in which, for example, the weights had settled down to reasonable values within four minutes of a 30 minute run for neutron penetration through 80 cm of iron.

6 THE HANDLING OF DATA FOR PENETRATION CALCULATIONS

6.1 Neutron data

The choice of a scheme for handling interaction data must be made in the face of conflicting demands.

(i) The methods and data must be simple enough to ensure speed of execution and economy of storage.

(ii) The treatment must be sufficiently detailed to give the required accuracy in a wide range of problems.

The scheme which is used in the Monte Carlo shielding codes at Winfrith is the DICE system [19] which consists of two parts:

(i) a programme MOULD which processes data from the UK Nuclear Data Library into a form suitable for use during Monte Carlo tracking and writes it onto a tape;

(ii) a suite of routines which are loaded with the Monte Carlo programme and are called to calculate mean free paths and to determine the outcome of interactions.

In this system energy is treated as a continuous variable but the cross sections (both total and partial) are assumed to be constant in each of a number of groups. These groups, which must be the same for all elements, can be chosen when the data tape is generated by the MOULD programme, which calculates the cross sections by averaging the UKNDL values over each group. Angular distributions are also converted by MOULD into a more convenient form which consists of equiprobable ranges within each of which the probability is assumed to be constant. 32 equiprobable ranges have normally been used. Energy distributions for scattered neutrons are taken from the UKNDL without alteration except that continuous probability distributions are again expressed in terms of equiprobable ranges (64 in number).

As a compromise between the conflicting requirements noted above, the cross sections have normally been represented in 300 DICE groups covering the range 0.025 eV to 15 MeV with equal lethargy intervals. This group size is sufficiently narrow not to require the group averaging process to take account of broad variations in the shape of the flux spectrum. On the other hand, the groups are too wide to follow cross section changes in the vicinity of narrow resonances, and it is therefore necessary to consider the effect of using average values in such regions and to examine the averaging process used. The programme MOULD calculates the group averaged microscopic cross sections for each element using as its weighting factor the reciprocal of the total cross section for the element. Thus, for a pure element the weighting factor approximates to the asymptotic infinite medium flux shape $(E \xi \Sigma_t)^{-1}$, since the energy E and the mean logarithmic energy decrement ξ are almost constant over such a narrow group. In this case the error will therefore depend on how far the spectrum differs from this asymptotic shape. In the case of a mixture of elements an error is introduced since the averaging is done before the macroscopic cross section is calculated and it does not therefore take account

of the environment of the element.

6.2 Tests of neutron data

These effects are currently being examined, principally in the context of benchmark experiments [20] performed in the ASPIS shielding facility at Winfrith in which the source is a conventional fission plate. The standard scheme with 300 DICE groups has been used in calculations for an experiment in which spectra were measured in a thick iron shield. As can be seen in Figure 1, the calculated fluxes below a few hundred KeV are higher than the spectrum obtained from proton-recoil counter measurements by up to a factor two. In order to find out whether these differences can be attributed to the cross section representation in this region of narrow resonances and minima some tests have been made in which very much finer DICE groups were used. For these tests the MOULD programme was rerun to process the same basic data into about 1000 groups above 10 KeV. Figure 2 shows in a small but typical energy range how the standard group scheme fails to follow the cross section variation, and that while even the 1000 group scheme is not completely adequate it is sufficient to provide a test of the effect of a finer representation. A one-dimensional case with a plane fission source was used and the flux at up to 70 cm from the source was calculated in fairly broad groups (similar to those used in the ASPIS calculation). The McBOX programme was run for long enough (15 minutes CPU on an IBM 370/165) to give standard deviations of about 10% for the group fluxes. No statistically significant differences were found between the results obtained with the standard scheme and the fine group scheme. With the broad group scoring it was not possible to observe the differences of fine structure which were undoubtedly present, but the result shows that a finer cross section representation does not significantly alter the general flux level. It is concluded that the observed differences from the measured spectra are not due to the cross section representation, and that in problems similar to this the normal DICE group scheme is adequate for predicting at least those response functions which do not show a rapid variation with energy.

6.3 Data sensitivity calculations

Tests of this kind involving the comparison of independent Monte Carlo runs are often difficult because the differences are masked by statistical errors. The problem arises also in calculating the sensitivity of results to small changes in the basic data itself. A way of overcoming this difficulty is to use the same set of histories for both cases and to take account of data changes by means of weighting factors [21]. Consideration is being given to the use of this method for the data sensitivity calculations which are required in shield design problems with complicated geometries in order to identify the major sources of error and to make error estimates. For the method to be useful it is necessary to be able to look at the effects of a large number of different data changes, and this is best done by writing details of the basic set of histories onto a magnetic tape which can then be processed as many times as necessary.

6.4 Gamma-ray data

Data handling is very much simpler for gamma-ray tracking than it is in the case of neutrons. The GAMBLE system [22] was written as a gamma-ray counterpart

of the DICE system, and was designed so as to be interchangeable with it. It has been found necessary to consider only three reactions: Compton scattering, pair production (in which the photons are assumed to be produced at the collision point), and photoelectric absorption. The justification for omitting Bremsstrahlung effects is based on the work of Dutton [23] who showed their contribution to be negligible in reactor shield design calculations, when the principal sources are usually below 6 MeV. The GAMBLE system uses a version of the routine HEITLER [24] to derive cross sections for any element from data in the UK Data Library. During programme execution the cross sections are stored point-wise, and linear interpolation is employed.

7 THE ALBEDO METHOD FOR STREAMING CALCULATIONS

The efficiency of reactor shielding is always reduced to some extent by the presence of essential passages and void-spaces containing gases or other low density materials. The coolant passages through the shield of a gas-cooled reactor provide an obvious example, but the problem is not confined to this system since access spaces and structural clearance gaps must also be present in the shields of water-cooled reactors and of sodium-cooled fast reactors. In the scheme described in Section 2 the codes provided for streaming calculations in the survey option are based on the kernel-albedo method. This technique is, however, restricted to the simpler geometries and it is therefore necessary to provide a Monte Carlo option for the correct treatment of more complicated systems. For this purpose, methods have been developed by Miller [25] and have been evaluated by means of a code, RANCYN, for cylindrical ducts and a more general geometry code RANSORD.

7.1 The form of the albedo

The general form of the albedo (representing the probability of reflection of radiation from a surface) may be expressed as

$$\beta(\underline{r}_o \rightarrow \underline{r},\ E_o \rightarrow E,\ \underline{\Omega}_o \rightarrow \underline{\Omega})$$

where $\underline{r}_o$, E_o and $\underline{\Omega}_o$ are the position, energy and direction of the incident radiation and $\underline{r}$, E and $\underline{\Omega}$ are the corresponding quantities for the emergent radiation. For both neutron and gamma-ray data the following general simplifications have been made:

(i) Reflected particles are assumed to emerge at the point of incidence, ie $\underline{r} = \underline{r}_o$.

(ii) The angular distribution of the emergent radiation is assumed to be independent of both the energy and the direction of the incident radiation.

(iii) The energy distribution of the emergent radiation is assumed to be independent of the direction of the incident radiation.

The remaining functional dependences are then expressed in the following

particular forms:

(i) Energy is represented group-wise.

(ii) The angular terms depend on a power of the cosine.

Thus, the albedo for a given wall material may be written as

$$\beta(g_o \rightarrow g)\, A(n)\mu_o^{-n(g_o)}\, B(n)\mu^{1-m(g)}$$

where g represents a group number and μ is the cosine of the angle to the surface normal. For neutrons the β term has been obtained from diffusion calculations while the values of A and n or B and m have been derived by fitting the expressions to the results of Monte Carlo calculations. For gamma-rays the functions have been derived from published compilations of differential albedos.

This degree of simplification has been used not only to ensure rapid sampling from the reflected distribution but, more importantly, to reduce the task of calculating albedo data for the wide range of wall materials encountered in practical problems. In future developments it is anticipated that the azimuthal dependence can be introduced into the albedo used with the Monte Carlo method without an appreciable loss of speed or an unacceptable amount of work to calculate the additional data. However, it would not be possible for all the simplifications to be removed, nor does this appear to be necessary on the evidence of the successful applications of the simplified albedo. Two examples of such applications are illustrated in Figures 3 to 6. Figure 3 shows the one-, two-, and three-legged 3 ft. square concrete-walled ducts studied extensively at ORNL [26]. The source from the Tower Shielding Reactor is a collimated neutron beam incident at 45° on a side wall. In Figure 4 the results of RANSORD calculations are compared with measurements of the thermal flux due to incident thermal neutrons. Figure 5 illustrates a three-legged duct of rectangular cross section assembled into the thermal column of the GLEEP reactor at Harwell. A Monte Carlo calculation for this duct was performed with the cylindrical geometry RANCYN code by representing the duct cross section by a circular cross section of the same area. The predictions of the response of a BF_3 chamber are compared with the measurements and with the predictions of the kernel-albedo code MULTISORD in Figure 6.

7.2 Variance reduction techniques

In most streaming problems some form of acceleration is required if the Monte Carlo method is to be of practical use, and various techniques have been examined.

(i) The first methods investigated involved biassed sampling from the angular distribution for reflected radiation. Three types of biassing were examined:

(a) The azimuthal angle was sampled from a distribution with a linear bias which increased the probability of reflection in the forward direction by a factor N_c rela-

tive to the backward direction.

(b) A polar angle biassing in the form of an inverse power of the cosine of the polar angle (μ^{-N_D}) was used to increase the probability of sampling at glancing angles.

(c) Correlated sampling of the azimuthal and polar angles was used, with stronger biassing being applied near the start of the leg than near its exit.

It was found that the optimum choice of biassing functions increased the efficiency of a three-legged duct calculation by a factor 60. Polar angle biassing was most effective, and surprisingly correlation of the azimuthal and polar angle sampling did not improve efficiency.

(ii) The method of splitting and Russian roulette applied at each point of reflection was examined. The importance function used was independent of energy and was an exponential function of the distance along each leg of the duct ($e^{\alpha z}$). The function was adjusted in trial runs in order to give an approximately uniform flow of particles along the duct. In a three-legged problem with a mouth source the use of splitting improved the efficiency by a factor of 500.

(iii) With the high intergroup transfer probabilities in the albedos it is usually a problem to obtain accurate results in the high energy groups. There is a straightforward way of avoiding this difficulty if the emergent angular distribution in the albedo is independent of energy since in this case the same set of particle tracks may be used for all groups and the group dependence from the other terms in the albedo can be taken into account by carrying a separate particle weight for each group. This "group correlation" approach has been extended to albedos in which the angular distribution does vary with energy, when these variations are allowed for by means of additional weighting factors.

Table I summarises the relative efficiencies of RANSORD calculations for the three-legged ORNL duct using the various angle biassing and splitting options. It can be seen that for this case it was splitting which was the most powerful technique, and this enabled the thermal neutron calculation described above to be completed in 3 minutes on an IBM 360/75 machine. The group correlation technique was used in the calculation for the GLEEP rectangular duct, in which the Monte Carlo results in Figure 6 were obtained in 1 minute. The increase in efficiency was estimated here as a factor 3. Thus, although the major improvement in efficiency achieved with this method occurs at high energies, the calculation of low energy fluxes is also accelerated significantly.

8 SUMMARY

The special features required of Monte Carlo codes to be used by shield designers have been identified as a flexible mode of construction, efficient operation, and ease of application. In the scheme described in this paper the flexibility is provided by the WRS modular coding system, which facilitates the assembly of Monte Carlo codes and their linking with the approximate methods required for many practical applications, and allows alterations or additions to be made without the need for major reprogramming. The requirements for speed and for economy of storage have influenced not only the variance reduction methods but also the choice of simplified treatments of geometry, neutron cross section data, and albedo data. The need to facilitate the task of the programme users has led to the development of automatic variance reduction facilities. It has also set a limit on the degree to which cross-section data can be simplified, to ensure that the user is not left with a choice of parameters such as the number of groups or the cross section averaging procedures, which may be problem dependent.

The aim of providing Monte Carlo methods which can be used efficiently by engineers for routine shield design calculations on commercial reactors presents considerable problems. It is believed, however, that the techniques described in this paper represent significant progress towards this goal.

ACKNOWLEDGMENTS

The author is grateful for the co-operation of the other members of the Winfrith Radiation Physics and Shielding Group whose work is described in this paper. Thanks are due especially to Mr D E Bendall and Dr P C Miller, and to Dr J Butler for his guidance during the preparation of the paper.

REFERENCES

1. L M C DUTTON, "Some Problems Associated with Calculating the On-Load Dose Rates in the Pile Cap of an Advanced Gas-Cooled Reactor", in Proceedings of Fourth International Conference on Reactor Shielding, Paris (1972), Vol. 3, p 771.

2. I G PUGH, R S HOVERD and A F AVERY, "Investigations of Shielding Problems Arising from Modifications to Refuelling Machinery at TRAWSFYNYDD Power Station", in Proceedings of Fourth International Conference on Reactor Shielding, Paris (1972), Vol. 3, p 973.

3. A F AVERY et al, "Radiation Level Predictions for Shielding of Pressurised Water Reactors", Trans. Am. Nucl. Soc., 17, 555 (1973)

4. D E BENDALL and A K McCRACKEN, "McBEND - A Prototype Code Utilising both Removal-Diffusion and Monte Carlo Methods", in Proceedings of Conference on the Physics Problems of Reactor Shielding, Harwell (1967), Vol. 1, p 41.

5. M J GRIMSTONE et al, "The Use of Modular Codes in a Unified Scheme of Shield Design Procedures for Commercial Fast and Thermal Reactor Plant", in Proceedings of Fourth International Conference on Reactor Shielding, Paris (1972), Vol. 2, p 409.

6. A F AVERY, J CLARKE and MRS A HARTLEY, "COMPRASH", AEEW-M 648, Atomic Energy Establishment, Winfrith (1966).

7. D E BENDALL and S J CRIPPS, "A Modified Form of Diffusion Theory for Use in Calculating Neutron and Gamma-ray Penetration in Practical Shields", in Proceedings of Fourth International Conference on Reactor Shielding, Paris (1972), Vol. 2, p 501.

8. J G MOORE, "The General Monte Carlo Code MONK", paper presented at this meeting.

9. D E BENDALL, ENEA Computer Programme Library Newsletter No. 11, March 1971.

10. R J BRISSENDEN, ENEA Computer Programme Library Newsletter No. 11, March 1971.

11. D E BENDALL, Atomic Energy Establishment, Winfrith, Private Communication (1973).

12. M E R VEALE, Atomic Energy Establishment, Winfrith, Private Communication (1973).

13. E R WOODCOCK et al, "Techniques Used in the GEM Code for Monte Carlo Neutronics Calculations in Reactors and other Systems of Complex Geometry", in Proceedings of the Conference on the Application of Computing Methods to Reactor Problems, Argonne National Laboratory (1965), ANL-7050, p 557.

14. P J HEMMINGS, "The GEM Code", AHSB(S) R 105 (1967)

15. J BUTLER, "A Review of Contemporary United Kingdom Shielding Techniques with an Indication of Future Trends", in Report on a Joint ENEA/IAEA Specialist Meeting on the Physics Problems of Reactor Shielding, Paris (1970), p 109.

16. D E BENDALL, "McNID - A Monte Carlo Programme for Calculating the Penetration of Neutrons in Systems with Cylindrical Symmetry", AEEW-R 308, Atomic Energy Establishment, Winfrith (1966).

17. J K REID, Atomic Energy Research Establishment, Harwell, Private Communication (1972).

18. C W J McCALLIEN, UKAEA Risley, Private Communication.

19. J B PARKER (Ed.), "DICE Mk V: The Preparation of Nuclear Data into a Form Suitable for Monte Carlo Calculatings Using an Electronic Computer", AWRE Report No. O-27/66, Atomic Weapons Research Establishment, Aldermaston (1966).

20. J BUTLER and R NICKS, Report of Joint NEA/EURATOM Specialist Meeting on Shielding Benchmark Experiments, Ispra (1974), to be published by EURATOM as an ESIS Newsletter.

21. G GOERTZEL and M H KALOS, "Progress in Nuclear Energy, Series 1, Volume 2", Pergamon Press (1958), p 315.

22. D E BENDALL and J M PARKER, Atomic Energy Establishment, Winfrith, Private Communication (1968).

23. L M C DUTTON, "The Importance of Bremsstrahlung in the Shielding of Gamma-rays having Energies less than 10 MeV", AEEW-R 675, Atomic Energy Establishment, Winfrith (1969).

24. A FODERARO, "Subroutine HEITLER", AERE-M 1956, Atomic Energy Research Establishment, Harwell (1967).

25. P C MILLER, "Radiation Streaming in Ducts and Void Spaces", Ph.D. Thesis, University of London (1973)

26. R E MAERKER and F J MUCKENTHALER, "Calculations using the Albedo Concept, of Thermal-Neutron Fluxes and Fast Neutron Dose Rates along the Centre Lines of Square Concrete Ducts: Comparison with Experiment", ORNL-4147, Oak Ridge National Laboratory (1967).

TABLE I

Relative Efficiencies of the RANSORD Acceleration Techniques for a Three-Legged 3 ft x 3 ft Concrete Duct Geometry

Importance Splitting	Reflected angle biassing			Relative Efficiency
	Polar angle N_D value	Azimuthal angle N_c value	Correlation	
yes	0	1	no	500
yes	1.6	100	yes	160
no	1.0	1	no	65
no	1.5	5	no	65
no	0	1000	no	25
no	1.5	5	yes	18
no	0	10	no	17
no	0	100	no	17
no	1.6	1000	no	11
no	1.6	100	no	6
no	1.7	10	yes	4
no	1.7	100	yes	4
no	0	2	no	4
no	1.6	1	no	3
no	1.5	1	no	2
no	0	1	no	1

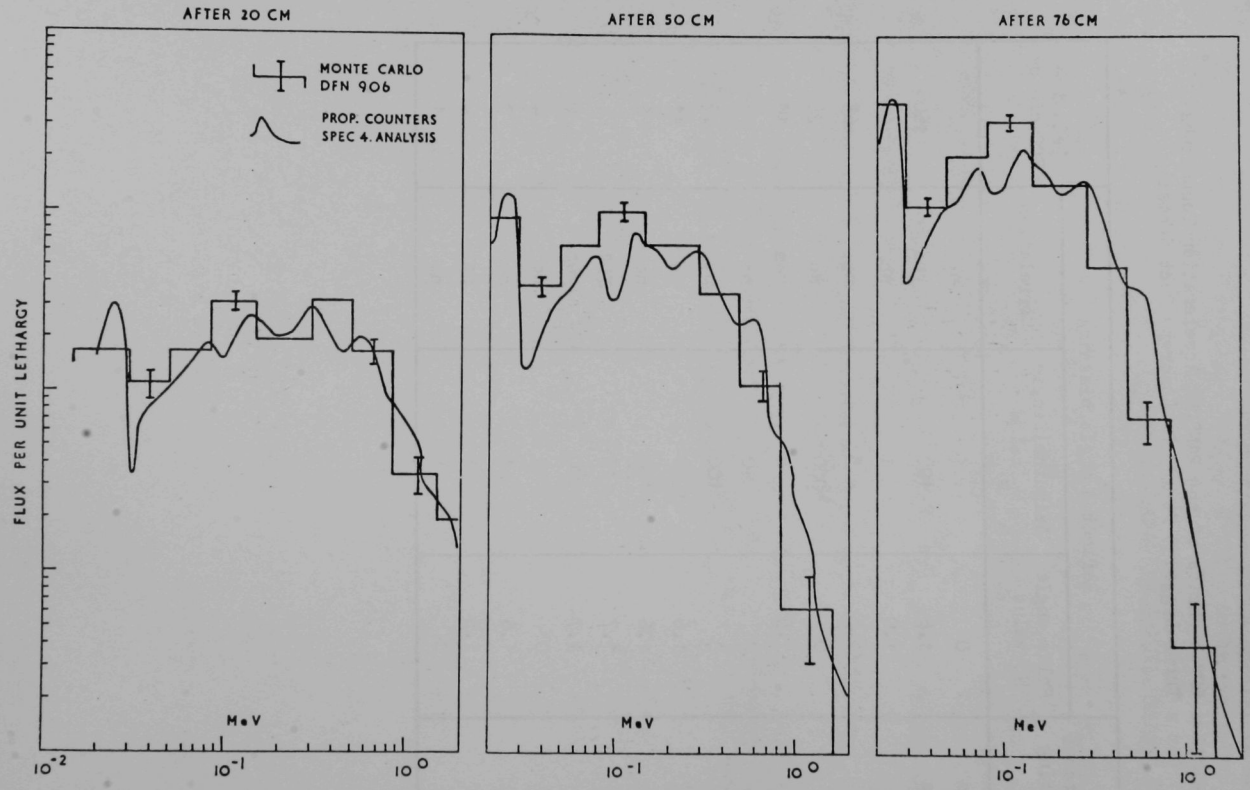

Fig. 1. Comparison of measured and predicted neutron spectra in a mild steel shield in ASPIS.

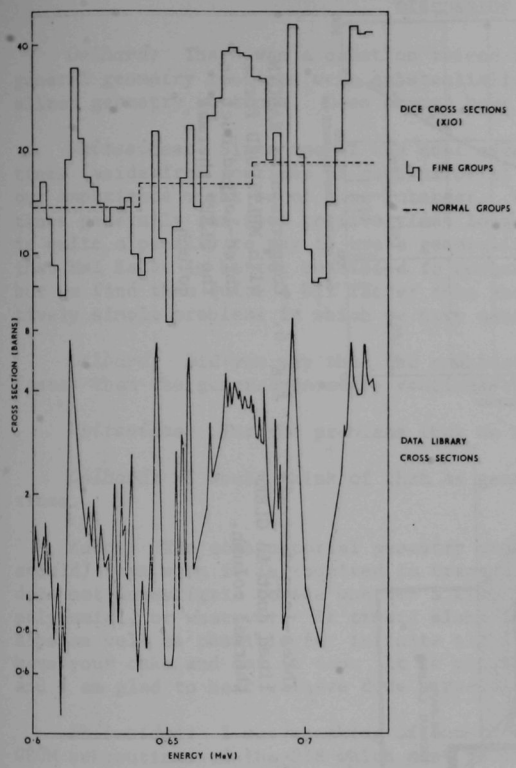

Fig. 2

Total cross section for iron.

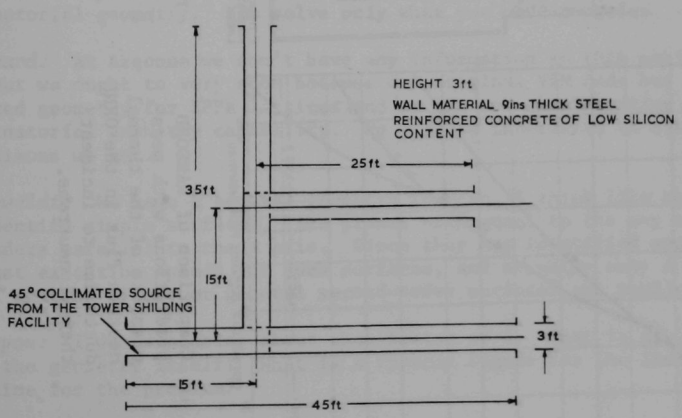

Fig. 3. Geometry of the ORNL concrete ducts (ORNL-4147).

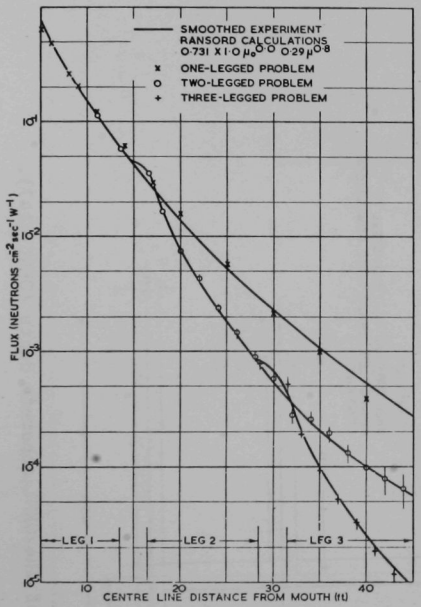

Fig. 4. Comparison of RANSORD predictions with measurements of the thermal neutron flux in the ORNL ducts due to incident thermal neutrons.

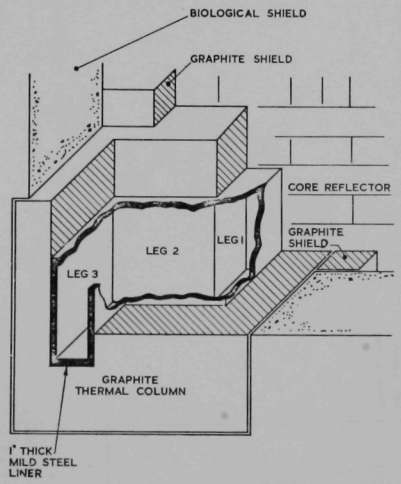

Fig. 5. Three-legged steel-lined duct in GLEEP thermal column.

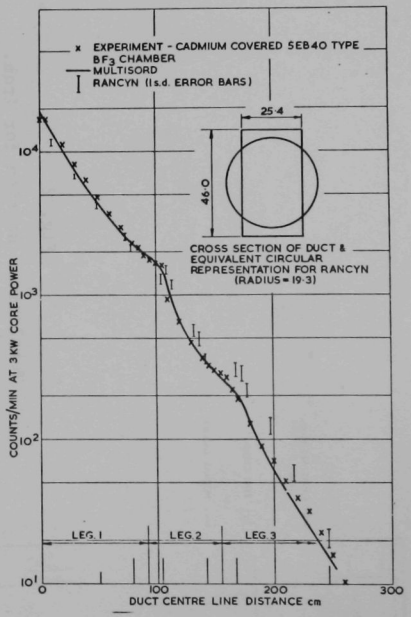

Fig. 6. Comparison of RANCYN with MULTISORD for a three-legged duct in the GLEEP thermal column.

DISCUSSION

Gelbard: There was a question raised initially in this talk as to whether general geometry routines were substantially more expensive to use than specialized geometry routines. Does anyone have any information on this subject?

Whitesides: Since one of the most extensively used general geometry routines (aside from routines of combinatorial type) was written at Oak Ridge, our experience might be of some interest. We find that general geometry routines generally run four to five times longer than specialized routines. That is quite a penalty to pay to use a generalized geometry routine. I believe that Mal Kalos is better qualified to comment on the combinatorial routines, but we find them quite a bit faster than the generalized routines for the relatively simple problems in which we have used them.

Gelbard: Did you say that the combinatorial routines were quite a bit faster than the general geometry routines that you have been using?

Whitesides: For the problems that we have run, that seems to be the case.

Gelbard: I would think of them as general routines, not specialized routines.

Kalos: The combinatorial geometry code has the property (as such codes should) that when it is required to treat a slab, it does so very well. It does not investigate to see whether a slab is, perhaps, a general n-*th* degree polynomial, or whatever. It treats slabs in a rather good way, though not perhaps as well as possible for infinite slabs in all cases. It is possible to have your cake and eat it too. It is possible to make reasonable compromises and I am glad to hear we have done a reasonable job in combinatorial geometry.

Whitesides: I was speaking of combinatorial geometry, as opposed to the GEOM subroutines of the 05R which most people are using. GEOM solves a general equation for every surface, and you may not have to do this, as Mal Kalos says, in combinatorial geometry. You solve only what you need to solve.

Gelbard: At Argonne we don't have any information on this subject at this point. But we ought to very soon because our original VIM code has a very specialized geometry for ZPPR lattices and we are now supplementing this with the combinatorial geometry capability. We will be interested to see what sort of comparisons we get.

Borgwaldt: We have a general geometry routine, I think like that in 05R. But we identify simple surfaces, like planes orthogonal to the x-y and z axes, and cylinders parallel to the z axis. Since they are identified separately, we get fast execution speeds for such surfaces, and normally only a very small number of general planes or general second-order surfaces are really used.

Coveyou: I was wondering about that factor of 4. That is the execution time for the geometry itself? What is a typical figure for the increase in running time for the problem?

Whitesides: I am not sure I know ...

Kalos: In a lot of problems, in our experience, tracking through the geometry takes a very large fraction of the running time — at least half, perhaps more.

Whitesides: Yes, that is right.

Gelbard: What is it in RECAP? Is it about 40%? I think that you mentioned that earlier.

Gast: Yes, that is correct.

Gelbard: So, one-half is not unusual for a wide range of codes.

Kalos: I would like to add one comment about the use of complex geometry routines. In particular, in combinatorial geometry, the running time, the speed with which you treat a complex geometry, can well depend on how you set the problem up. For example, you may have a particularly complex domain which is not encountered very often and which may require a great deal of computation. If you use the simple device of putting it in a box, then you never have to deal with the surfaces in this domain unless you first encounter the box. The geometry computations will, then, speed up considerably. I don't know how to give general prescriptions for the average user, but a word to the wise will go a long way.

Gelbard: The other suggestion that has been mentioned is that simple surfaces be identified specifically in general geometry.

Coveyou: When I designed this sort of routine originally it was based on a block and zone concept, so that you in fact never did use the more complex parts of the routine unless they were needed. Has that feature been dropped?

Whitesides: No it has not been dropped, obviously. But the problem is, as Kalos correctly pointed out, that the efficiency of the routine in a specific case depends on the care and time you invest in writing out the geometry description. In general, most people find that it is easier to let the machine solve the geometry problem than to try to treat the geometry carefully themselves.

Gelbard: I would like to rush the discussion today, because I want to stick strictly to schedule and would like to cover some other questions. The biasing that you discussed initially was one where the machine adjusted the biasing so that you got equal density particles in equal volume of phase space?

Grimstone: Yes.

Gelbard: Have you had any opportunity to notice what effect this has on the error estimates? Have you had any trouble with error estimates with this sort of biasing.

Grimstone: No, we have not really looked at this in any great detail.

Kalos: I want to ask first of all a technical question about this adaptive procedure. When you use this method, do you start with some sort of guess

as to what would be a good set of weights, or do you start with flat weights and go on from there?

Grimstone: We can do both; but I think in most applications we have, in fact, started with flat weights.

Kalos: So that means that your information grows, diffuses out slowly from the source, and it also means that you have to execute some sort of cycle, stopping, taking a tally of what has happened, and improving and going on. About how many such bootstraps seem to be necessary in a penetration problem where you have an attenuation of, say, 10^6?

Grimstone: Generally we make weight readjustments after every particle history, and do this throughout the entire calculation.

Kalos: Then the procedure is biased, your answers are biased.

Gelbard: You recall that he said that the arrangement settled after 5 minutes of computation, that although one continues to adjust the weights in principle, in practice it is claimed that the adjustment process settles after 5 minutes.

Kalos: So that, in practice, he is claiming that, very likely, the bias is small. I don't deny that.

Gelbard: This does depend on the settling time.

Grimstone: Presumably the bias would be smaller if we were to discard that first 5 minutes of history, which we don't, in fact, do at the moment.

Kalos: If you did that, then it would be unbiased.

Grimstone: Or almost unbiased ...

Kalos: No it would be unbiased, completely unbiased if you discarded the settling period and did not readjust weights afterwards.

Grimstone: If we continued to refine the weights somewhat then it would be somewhat biased, but probably not to a serious extent.

Kalos: I regard the fact that you are able to bootstrap from zero information to a reasonably satisfactory scheme as a significant triumph. In the very early days we tried manually to do things like this and our experience seemed discouraging. I am very encouraged and hope that we can do some experiments along similar lines. We have had some ideas about bootstrapping which are not quite the same as yours, but they rely on one's ability to do this kind of thing, and I think that your success means that we will have success.

Coveyou: Have you performed the experiment of doing this kind of calculation twice, and comparing the weights you get to see if they vary very much from one run to another.

Grimstone: Yes, I have tried to see how rapidly the weights appear to be converging. The weights that we use are, in fact, very closely related to,

and very similar in properties to the flux estimates themselves, and therefore behave in roughly the same way.

Coveyou: I wonder whether such a procedure has, shall we say, metastable states. Would you get a set of weights on one run, and yet another set of weights and another set of answers if you were to rerun the whole problem again, with a different set of starting weights and random numbers? I am wondering if it is sufficient to rely on the fact that the weights seem to have settled down.

Grimstone: Well when I say they settle down what I mean by that is that they might only have reached, say, within a factor of 2 of their final values. I am not putting any very stringent convergence criteria on the weights.

Gelbard: I want to direct a brief question to Whitesides. In Grimstone's study of the effect of cross sections, it was found that the cross sections for iron in a big hunk of iron did not have to be represented very accurately. I recall Monte Carlo studies of this sort at Oak Ridge, and I seem to remember that they came to other conclusions. Is that not true?

Whitesides: That is correct. It just depends upon what sort of problem you are looking at. If you are looking at a deep penetration problem through iron, and you have many neutrons in the vicinity of the windows, they are going to stream through. One thing that I have personally been concerned about for some time (and this applies to criticality problems where you have heavy metal reflectors) is that it is very, very difficult, I believe, to define a multigroup set that is going to work. You are simply going to run into streaming through windows when you least expect it, and when this streaming might cause you the most difficulty.

Gelbard: The problem that was described here, was it not also a problem of penetration through a large chunk of iron?

Grimstone: That is right, yes.

Gelbard: So it is not really clear why the conclusions are so different. The Oak Ridge conclusion has been that you need an accurate cross section representation.

Whitesides: Obviously you can get a set that would give you the right answer.

Gelbard: But the windows all disappeared in the smooth set that Grimstone used.

Whitesides: That is so.

Gelbard: If you can close these windows with impunity, then the cross-section problem becomes very much easier. The S_n method then becomes a much more important competitor. We don't have time to discuss this, but I think it is really important to people in our business to think of the role of other computational methods. If windows don't matter too much, S_n becomes a tool which is much more valuable.

MCT, A MONTE CARLO CODE FOR TIME-DEPENDENT NEUTRON THERMALIZATION PROBLEMS

E. Cupini and R. Simonini

Comitato Nazionale Energia Nucleare
Centro di Calcolo
Bologna, Italy

* * *

Presented by R. Simonini

MCT, A MONTE CARLO CODE FOR TIME-DEPENDENT NEUTRON THERMALIZATION PROBLEMS

E. Cupini and R. Simonini

Comitato Nazionale Energia Nucleare
Centro di Calcolo
Bologna, Italy

ABSTRACT

In the Monte Carlo simulation of pulse source experiments, the neutron energy spectrum, spatial distribution and total density may be required for a long time after the pulse. If the assemblies are very small, as often occurs in the cases of interest, sophisticated Monte Carlo techniques must be applied which force neutrons to remain in the system during the time interval investigated. In the MCT code a splitting technique has been applied to neutrons exceeding assigned target times, and we have found that this technique compares very favorably with more usual ones, such as the expected leakage probability, giving large gains in computational time and variance. As an example, satisfactory asymptotic thermal spectra with a neutron attenuation of 10^{-5} were quickly obtained.

INTRODUCTION

Time-dependent neutron transport calculations by the Monte Carlo method are of considerable interest in studying the approach to equilibrium of a neutron pulse in a moderating assembly. In these problems the neutron energy spectrum, total neutron density, and spatial neutron distribution are to be calculated at different time intervals after the neutron pulse. Results of special interest are those for very small samples, due to the fact that beyond a critical buckling no discrete decay constant exists. To investigate the process of thermalization in these small samples the Monte Carlo code technique, known as "expected leakage probability", is usually applied, forcing the neutron emerging from a collision to remain in the assembly. In our Monte Carlo code a splitting technique has been adopted which, in the cases examined so far, compares very favorably with respect to the expected leakage probability. By this technique a neutron is split when assigned target times are exceeded, so that it is possible to follow particles for a long time after the pulse without large fluctuations of the weights scored per history provided the number of particles born from splitting at each target time is properly chosen, for example on the basis of a previous analog Monte Carlo calculation.

THE MCT CODE

The purpose of the MCT code is the simulation of the time-dependent neutron transport of pulsed-source experiments in finite or infinite assemblies.

The geometries that the code can handle were chosen in order to ensure both the possibility of comparison with typical experimental assemblies and the speed of calculation: at present, single spheres, cylinders, and parallelepipeds containing a homogeneous mixture can be treated, together with the infinite medium. It is also possible to treat in a single run several assemblies of the same type and different sizes, each internal to the other, a neutron being considered internal to all assemblies from which it has not yet escaped.

The code assumes that an assigned neutron point source pulses at the time $t = 0$, emitting a monoenergetic burst of neutrons of energy $E \geq E_{th}$, where E_{th} is the thermal energy upper bound (the source point location is not important if only the equilibrium distribution is to be investigated).

For the calculations source neutrons are followed, all information required being collected at fixed target times. A history is stopped when either the last target time is reached, or the neutron escapes from the system, or the statistical weight falls under a given cut-off value.

Two processes, the slowing-down and the thermalization and diffusion, are considered. During the slowing-down process, i.e. for neutron energy greater than E_{th}, nuclei are assumed to be at rest in the laboratory system, and only elastic scattering with constant cross section, isotropic in the center-of-mass system, is considered. During the thermalization and diffusion process, the thermal motion of the nuclei is taken into account. For the thermal collision simulation the code utilizes a very refined point description of the differential scattering cross sections $\sigma(E \to E')$ and $\sigma(E \to E', \cos\theta)$, where r is the scattering angle in the laboratory system. From the scattering matrix $[\sigma_{ij}]$, with $\sigma_{ij} = \sigma(E_i \to E_j)$, a table is set up of N equiprobable energy intervals with bounds (E_k, E_{k+1}) for each E_i defined by

$$\frac{1}{\sigma(E_i)} \int_{E_k}^{E_{k+1}} \sigma(E_i \to E)\, dE = 1/N , \qquad k = 0, 1, \ldots N - 1 ,$$

where (E_0, E_N) is the total post-collision energy range and $\sigma(E_i)$ the total cross section for the energy E_i. For the incoming energy E, $E_i < E \leq E_{i+1}$, two random accesses to the tables and successive linear interpolations provide the post-collision energy E'.

Similar tables are set up for the angular distributions of the energy transitions $(E_i \to E_j)$. Here in addition, thanks to the "detailed balance" relation

$$E\, e^{-E/T}\, \sigma(E \to E', \cos\theta) = E'\, e^{-E'/T}\, \sigma(E' \to E, \cos\theta) ,$$

the equality

$$\frac{\sigma(E_i \to E_j, \cos\theta)}{\sigma(E_i \to E_j)} = \frac{\sigma(E_j \to E_i, \cos\theta)}{\sigma(E_j \to E_i)}$$

shows that it is sufficient to set up tables only for the transitions $(E_i \to E_j)$ with $E_i \geq E_j$.

Absorption in the thermal zone is taken into account by a 1/v-type cross section.

The code is completely dimension-flexible. All variable-length quantities are stored in a single one-dimensional array whose length can be changed by simply substituting a "DIMENSION" declaration in a dummy subroutine.

Besides the analog simulation, three techniques can be used in the code to treat small geometries and time intervals sufficiently large:

(a) The usual "expected leakage probability" method (see for instance Ref. 1), which forces the neutron to stay within the assembly and compensates the biassing of the game by an appropriate weight. When using this technique, only one assembly may be treated in a run.

(b) The "expected value" calculation of the contributions to scores at the target times of interest: if a neutron emerges from a collision with parameters $(\vec{x}, v, \vec{\Omega}, t, w)$ ($\vec{x}$ position, v speed, $\vec{\Omega}$ direction, t time, w statistical weight), the probability p for the particle to exceed the *i-th* target time t_i, $t_i > t$, without undergoing a collision is given by $e^{-\Sigma v(t-t_i)}$, where Σ is the total macroscopic cross section, if the point $\vec{x}' = \vec{x} + v\vec{\Omega}(t_i - t)$ is still inside the assembly. In this case scores are carried out for a particle with parameters $(\vec{x}', v, \vec{\Omega}, t, wp)$. The original neutron then suffers a new collision or escapes without further contributing to the scores.

(c) The "splitting technique": to each of some assigned time steps t_i, $i = 1, \ldots, m$ (which may coincide with some target times), a splitting factor n_i is associated. When a neutron enters and emerges from a collision, with parameters $(\vec{x}, v, \vec{\Omega}, t, w)$ and $(\vec{x}, v', \vec{\Omega}', t, w')$, respectively, a check is made to see whether in the last free-flight it has exceeded any of the t_i. If so, the particle is split into n_k daughter particles, t_k being the largest time step exceeded. The parameters of the *j-th* particle, $j = 1, \ldots n_k$, are $(\vec{x}, v', \vec{\Omega}_j', t, w'/n_k)$, where $\vec{\Omega}_j'$ is such that $\vec{\Omega}_j' \cdot \vec{\Omega} = \vec{\Omega}' \cdot \vec{\Omega}$, i.e., a new azimuth is chosen for each particle in order to improve exploration of the geometry. Parameters about $n_k - 1$ particles are stored in a bank and the remaining one goes on with the flight, producing possibly new daughter particles until it is removed from the assembly; then the last particle stored is picked out from the bank and followed as a new one.

RESULTS AND CONCLUSIONS

To test the code, many calculations were performed for several moderators and for assemblies of different sizes. Results are quoted here for a 9.12-cm radius sphere of heavy water at 300°K. The differential cross sections were

calculated with the Haywood model [2], and used to set up tables of equiprobable intervals for 100 energy and 10 cosine points for the energy and angle, respectively. Source neutron energy and thermal threshold were fixed to 10 eV and 1 eV; the particles started at the center of the sphere.

The neutron densities at various target times, calculated by applying the splitting technique, are quoted in Fig. 1. The decay constant turned out to be 1.51×10^4 sec^{-1}. The asymptotic average speed was found to be 2.13×10^5 cm/sec. The neutron attenuation per source particle in the time interval examined is of the order of 10^{-5}, and the computing time needed to run the 30,000 histories was about 10' for an IBM-360/75 computer.

In Figs. 2 and 3, asymptotic spectra are shown at the time t = 500 μs, calculated by the expected leakage probability and the splitting technique, respectively. Besides the obvious gain in shape due to richer statistics, one must keep in mind that the latter was obtained with a computational time of about one-third that of the former.

Extensions of the code are planned to include the treatment of several, even multiplying materials.

REFERENCES

1. E. D. CASHWELL and C. J. EVERETT, *A Practical Manual on the Monte Carlo Method for Random Walk Problems*, Pergamon Press (1959).

2. B. C. HAYWOOD and D. I. PAGE, "Neutron Thermalization and Reactor Spectra," Vol. I, 361 (IAEA, Vienna (1968).

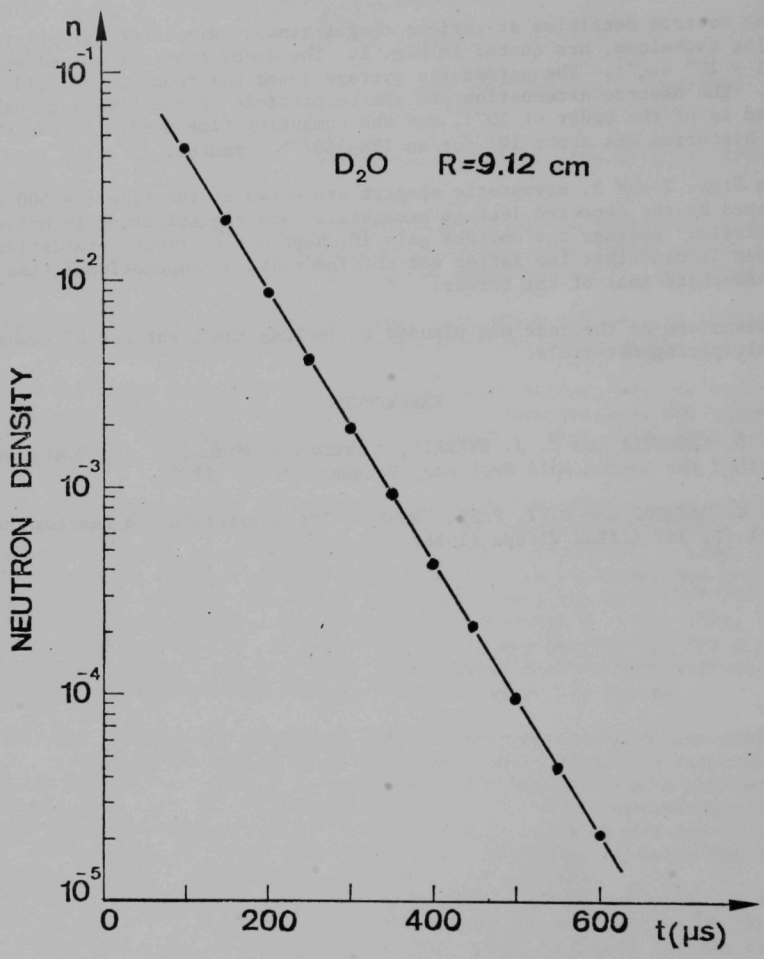

Fig. 1. Neutron density, as a function of length time, compiled by splitting.

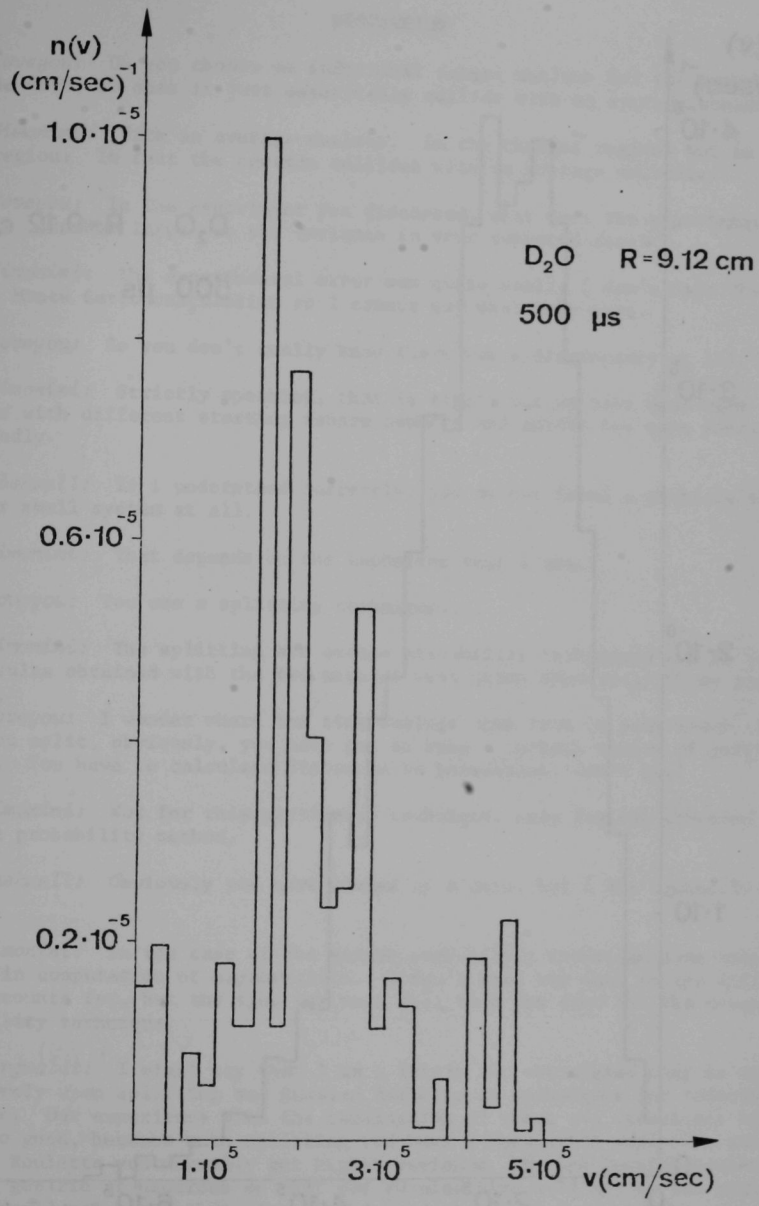

Fig. 2. Asymptotic spectrum t = 500 μs computed by the expected leakage probability method.

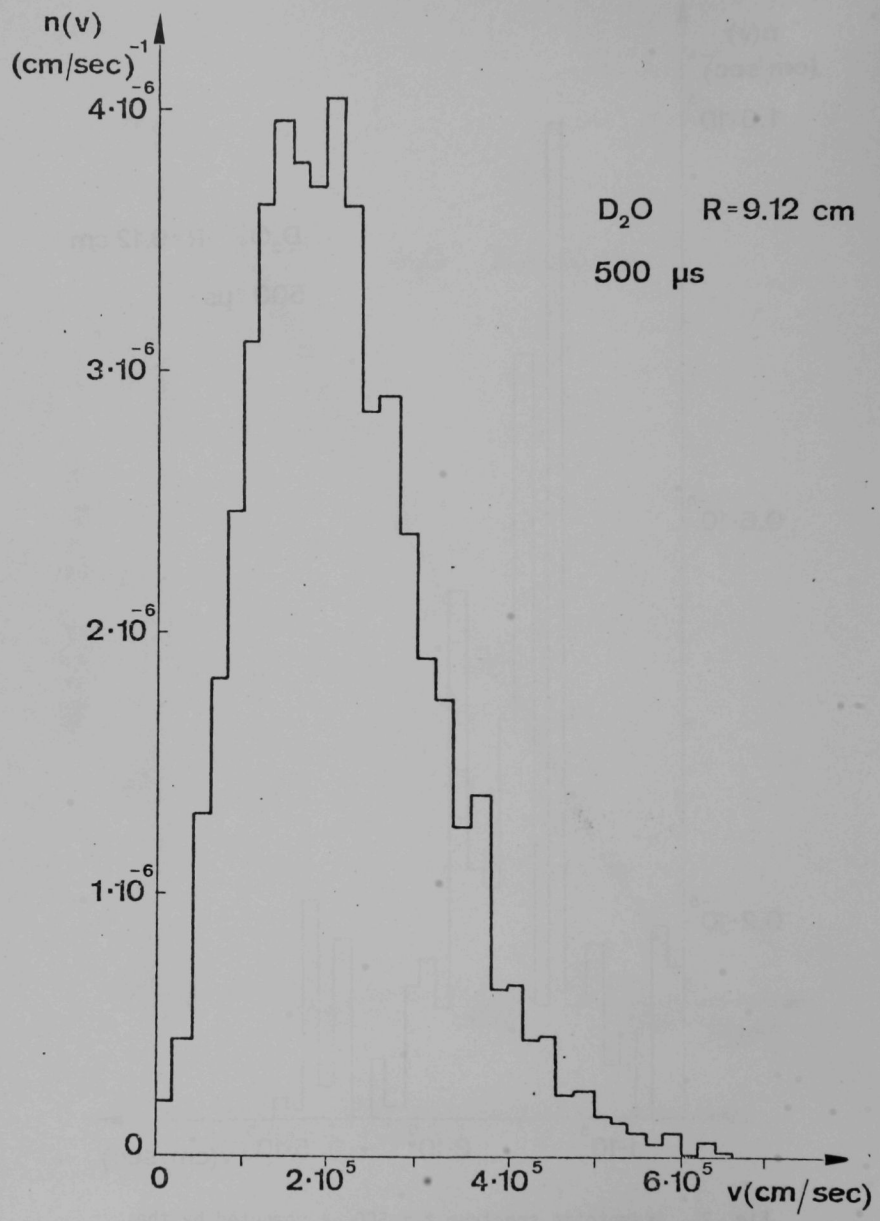

Fig. 3. Asymptotic spectrum at 500 μs computed by splitting.

DISCUSSION

Coveyou: Do you choose an individual target nucleus for the neutron to collide with, or does it just essentially collide with an average nucleus?

Simonini: With an average nucleus. In the thermal region, not in the fast region. In fact the neutron collides with an average molecule.

Coveyou: In the experiment you discussed, what were the experimental errors? And how large was the variance in your computed result?

Simonini: The experimental error was quite small; I don't have the errors in the Monte Carlo calculation so I cannot say what they were.

Coveyou: So you don't really know there was a discrepancy at all.

Simonini: Strictly speaking, that is right; but we have made many calculations with different starting random numbers and gotten the same discrepancy repeatedly.

Cashwell: If I understand correctly, you do not force a particle to stay in your small system at all.

Simonini: That depends on the technique that I use.

Coveyou: You use a splitting technique....

Simonini: The splitting and escape probability techniques can be combined, but results obtained with the two methods were shown separately in my paper.

Coveyou: I wonder where the time savings come from in your code? Because when you split, obviously, you have got to keep a certain number of particles around. You have to calculate distances to boundaries, don't you?

Simonini: Not for this particular technique, only for the expected leakage probability method.

Cashwell: Obviously you have picked up a gain, but I don't exactly see why.

Simonini: In the case of the escape probability technique some time is wasted in computation of exponentials. I don't know how much of the difference that accounts for, but the time was much less than the time for the escape probability technique.

Borgwaldt: I might say that I am a little bit astonished that so many people rely upon splitting and Russian Roulette as techniques for reducing variance. Our experience with the combination of these two techniques has not been too good, because with splitting you have a population explosion and with Russian Roulette you normally get higher variance. We use Russian Roulette only to get rid of neutrons so that our running time will not be too high. But I should point out that, for time-dependent problems, there are two other techniques that can be used in place of the extrapolated leakage probability method. One technique involves introducing an artificial anisotropy, an

angular biasing that keeps the neutrons inside the sphere or the pulsed body. Another technique is one which we have called the modified extrapolated leakage probability method. This second technique has been described in a paper given by me at the Idaho Falls conference in 1971, and we have found that this technique is really very efficient. If one looks at it very closely one sees that it is, from a mathematical point of view, a combined track length truncation and angular biasing technique that has the effect that, on the one hand, all the particles are kept strictly within the body in which you want to have them for your time-dependent problem. On the other hand, if a particle makes a collision near a boundary it is preferentially scattered backwards into the body, and this is done with a minimum of computation. It is only that instead of computing one exponential you have to compute two exponential terms. This method has been introduced into our code and I think that it has been successful.

Coveyou: I thought it was well known that Russian Roulette always increases the variance per history, and I think it is true, also that splitting always decreases the variance. I know that this fact seems to be ignored most of the time. People tend to say that they are reducing the variance either by splitting or Russian Roulette. That is not what happens. The Russian Roulette may result in a gain in efficiency, but only by decreasing the average running time per history. The variance per history just has to be increased by Russian Roulette and I think, although I am not quite so certain of this, that splitting always decreases the variance per history. But the point is that it produces more particles so that it takes you longer to finish your calculations.

Cashwell: You want to decrease the number of particles when you use Russian Roulette.

Coveyou: You want to do this to decrease the running time.

Cashwell: You could use the method that Bending talks about. When you get close to the boundary your preferential direction is back towards the center of a sphere, if you have a sphere. Generally it would be back towards the center of the body.

Borgwaldt (in summary): Perhaps I could briefly sketch how our method works. I might add that our system really is simpler. Consider, first, the ordinary extrapolated leakage probability (ELP) method. After each collision one chooses the next flight path and forces a neutron collision on that part of the path which lies within the body. What we do is slightly different. For simplicity let us assume the scattering is isotropic in the lab system. If this is not true our method becomes slightly more complicated, but it is still pretty simple. If the scattering is isotropic then, after each collision we select two possible flight paths in opposite directions. Both these directions are equally probable. Now we compute the conditional probability that a collision will occur on either of the two flight paths, assuming it occurs on one of them. We select one of the two paths from this conditional probability and force a collision, along that path, and within the body. If the last collision has occurred near the boundary, then the next collision will tend to be pushed back towards the interior of the body. And so we get a preferential scattering in the backward direction. This doesn't take that much computation, but it gives you twice as much information on the leakage probability information which you can use to estimate leakage. Further, it

works. It does not give you as large a scatter of neutron weights as you get in the normal ELP method, which sometimes forces a high density of collisions near a boundary. This artificially high density is associated with very low neutron weights. With our method the weights are more likely to remain normal.

Hoogenboom: I don't understand why it does not take much computer time. It seems that you have to examine the whole system, from the collision point first in one direction and then the other.

Borgwaldt: Well I must add that in practice there are, of course, some refinements. This technique is only used in designated outer regions. One can recognize quite easily that all such techniques make no sense when you are far from a boundary. In general, one might say that the amount of computation which you need is twice the amount needed for Rief's extrapolated leakage probability technique because you have to compute two exponentials. On the other hand, if you look at the paper I mentioned, published in the *Proceedings of Conference on New Developments in Reactor Mathematics and Applications*, CONF-710302 (1971), Vol. 2, p. 857, you can see that the behavior of the error in the ELP technique becomes catastrophic when you go to extremes, when you want to follow the decay of neutrons over several decades. Whereas, in our case, you lose computation time when you are not looking for large decay, but you gain considerably when you follow neutrons over several decades of intensity.

Gelbard: You cut off the thermal group at 1 eV?

Simonini: Yes.

Gelbard: So you also have the problem that I referred to, in another connection, earlier. There must be some binding effects epithermally, above 1 eV. Do you treat those?

Simonini: Above the thermal band we assume a constant cross section and isotropic scattering. The code is designed to investigate only the thermal properties of materials like heavy water and light water, nothing else.

Gelbard: When you actually compute a period, do you fit an exponential to the time behavior?

Simonini: Yes, of course.

Gelbard: There is another sort of procedure that has been used a great deal for computing, not periods, but decay lengths. I think it would be applicable here also. Rather than fitting the decay curve you can compute an absorption rate over the energy spectrum which, by neutron balance, gives you a period. I think the problem of fitting an exponential to a stochastic variable is probably a pain in the neck, whereas the other method directly gives you a period by balance and is probably somewhat easier. Before we leave this topic, I want to ask you whether the discrepancies you mentioned have been observed by other people.

Simonini: I don't know.

APPENDIX

GENERAL DISCUSSION
ON MONTE CARLO METHODS

GENERAL DISCUSSION

Gelbard: Coming back to the shielding calculation for a moment, I would like to ask Elliot Whitesides what he thinks is the proper role of Monte Carlo methods. At Oak Ridge you also use S_n in shielded calculations. What do you use S_n for, and what do you Monte Carlo for?

Whitesides: I guess I have been fortunate to see what I think is a switch in the roles of the two methods. I guess that I believe now that Monte Carlo is better applied to criticality problems, and that the discrete ordinates method is better applied to shielding problems. At Oak Ridge we use our discrete ordinates code for just about all of our shielding problems, with the single exception of problems that are so geometrically complex that we cannot describe them properly in an S_n calculation. For such cases we have a code called Domino which we use in coupled calculations. We use techniques that were developed by Tom Hoffman in his dissertation, I guess following along the lines suggested by Lichtenstein in his dissertation. In our coupled Monte Carlo and discrete ordinates calculations we couple the forward and adjoint calculations, and we think it is proper to use Monte Carlo in shielding design problems solely in situations that cannot be described geometrically with the discrete ordinates methods.

Gelbard: You use Monte Carlo for streaming in ducts, and for treating little complicated details of the shield? Is that what you are referring to?

Whitesides: That is right. In fact for the FFTF or the fast breeder reactor program we will do part of the problem with DOT, couple into MORSE, and then go back to DOT so that we single out isolated regions for the Monte Carlo treatment.

Gelbard: Generally, then, you find that complicated cross-section effects are not very serious as far as you are concerned.

Whitesides: Right. We have a fairly extensive sensitivity program under way in which we look at the effects of the cross-section representation in multigroup forms, and we rely on this sort of information to help us come up with multigroup cross-section sets that we will find to be adequate for our problems.

Gelbard: I recall that, at Bettis also, all the shielding work was done with S_n. There was no use of Monte Carlo at all. So we are gradually putting ourselves out of the shielding business. What, then, should we be doing. Should we, for example, be concentrating on the analysis of criticals? Some difficulties involved in the Monte Carlo analysis of criticals have already been pointed out.

Borgwaldt: Perhaps I could comment on this question. We developed our code, originally, with the idea that most people would use option 1, the source iteration option. We still think this option would be particularly useful for computing the flux and power density near the core-blanket interface, where the treatment of resonance cross section becomes really rather important and difficult. I must admit that we have not had much of a response. The other option, the census time option, is used to a large extent by experimental physicists who do the analysis of time-dependent leakage spectra in order to extract

information on inelastic scattering cross sections. This is the same type of experiment that has also been performed at the Livermore Laboratory. Other people have been using Monte Carlo to analyze transmission experiments. They have to do calculations for shadow cones, to treat shadows in the stream of neutrons from the source into the detector. Monte Carlo has also been used to analyze the time jitter which is introduced in the lithium glass detector by covering it with a foil of polyethelene. This type of application is really required by one group of users. We cannot create any excitement in the other group so far.

Gelbard: I have the feeling that we should have been more successful in getting Monte Carlo used for methods testing. At Bettis, shortly before I left, I was involved in some work which I might mention as an example of the sort of methods testing I have in mind. We were studying the transient absorption rate near the boundary between a water region and a fuel region. The resonance flux is not self-shielded as it goes from the water into the fuel region, so you get a sharp transient in the absorption rate. To study this phenomenon you can write a code in which you assume that the resonance material does not moderate at all. The only materials that moderate, let us say, are hydrogen and oxygen. Assume that their cross sections are constant in the resolved resonance range. Now all the scattering cross sections are constant and isotropic, so it is easy to write an adjoint code. And if you work in lethargy, as has been pointed out here, you have no difficulty with the adjoint method, you don't get divergent integrals. You only have to adjoint isotropic scattering in the center-of-mass. With such an adjoint code you can get a point plot of the activation rate as you approach the boundary. It is perfectly possible to do this, and it can be done very cheaply with such a completely artificial adjoint code. You then compare this activation plot with what you would get for the same problem by whatever analytical approximations you would want to make. There are many other cases like this, I think, where (if you have the people available) you can write a code especially to treat a simplified problem. You can then test any proposed approximate computational technique on the same simplified problem.

Cashwell: May I say a few words about this problem. Apparently at Los Alamos we see a little of the reverse happening, in that Monte Carlo is being used much more than ever before. Now, obviously a lot of our problems have very difficult geometries and, of course, that is an overriding reason for using Monte Carlo. In addition, some people, who don't necessarily need complicated geometries, are not willing to compromise on the treatment of the cross section. A lot of our work is motivated by concern over the validity of multigroup methods. I know that some Livermore codes use as many as 2000 groups, so it seems that they share this concern. So I think we see an increase in the use of Monte Carlo, and this is a little bit counter to what is going on in the reactor business.

Gelbard: It can have something to do with the amount of money available. I think that there is a sharp line between applications in which you are studying weapons, and civilian applications, and there seems to be a very large step jump in what you can afford to do when you get into any sort of military work.

Cashwell: I would say that machine time in the weapons business is not the main factor at Los Alamos. We have machine time. At Livermore they have even more machine time.

Whitesides: I feel I should take issue with what I understood you to say about the success of Monte Carlo codes in criticality work. I guess when we do a head-on comparison with other methods, for the same geometry and the same cross sections, Monte Carlo does as well as any other method. The difficulty only comes about in trying to use what we consider to be more refined data and more precise data in Monte Carlo, when we use ENDF data directly. In fact we do quite extensive comparisons, back and forth, between our discrete ordinates methods and our Monte Carlo, to check the accuracy of the discrete ordinates calculations. We get good results, and for most criticality type problems I would not run DOT.

Gelbard: The problem that keeps coming up, apparently in the fast reactor field, is not something that I had noticed before in the thermal reactor field. Fast reactor people always seem to insist that they want to calculate minute perturbations and that is not the kind of application that I have seen in thermal reactor work. It has something to do, I think, with the shocking fact that, if all the sodium boils out of a fast reactor, the reactor hardly even knows it. That seems to have something to do with it. If you were to empty all the water out of a thermal reactor you'd have a big blob of nothing.

Whitesides: That is true and as I mentioned the other day, we have been doing perturbation-type calculations on fast systems. I did get some data which I did not have with me the other day. Our computer time on the IBM-360/91 is less than 10 min. to calculate an error in k_{eff} of 1/2% for an FFTF-type engineering mockup.

Gelbard: But you also said that you were not looking at the most extreme cases where we have some concern as to whether the complete correlation method is even feasible.

Whitesides: Let me separate the two problems. I am saying that we calculate k_{eff} to within 1/2% in 10 minutes. And, at the same time, we calculate the change in k_{eff} due to some configuration change. What we have been looking at is a 4 by 4 element or plate or drawer, whatever you want to call it, in the system. In the experiment just one of these is moved out, but we have, in some of our calculations, been moving drawers out in some kind of uniform fashion around that single element, so as to enhance the perturbation. The sort of Δk's that we have been looking for are in the neighborhood of \$1 to \$6 in reactivity, and beta for this particular system is 0.003.

Gelbard: You have been computing three-tenths of a percent change in reactivity for the system, and you do this by enhancing the perturbation, by taking more drawers out than you actually would in the experiment?

Whitesides: That is the sort of thing we have been looking at.

Gelbard: In other words, what you are doing is not analyzing the true experimental situation, but trying to analyze something that is a little easier. This sort of approach seems to lie halfway between analyzing experiments and using Monte Carlo to test methods.

Whitesides: If you are really talking about changes in a small, a very small region, you might not ever be able to analyze such perturbations by Monte Carlo, just simply because you cannot sample adequately.

Gelbard: I suspect that Bernnat is going to point out that the adjoint method, which he has described, is really designed for just this sort of case, and we have not really tried it.

Bernnat: The use of the adjoint method is only one part of the method.

Gelbard: The Green's function .....

Bernnat: Right. If I have a large region in which a perturbation takes place, naturally I will try to treat this with a forward calculation. But I also want to point out that, when I compare Monte Carlo reactivity calculations with experiments, I never know how accurate the experiments are. When one is dealing with reactivities of $6 or more, what measurements technique can one rely on? How accurate are the different experimental methods in such situations? Pulse measurements are not very exact in this range. They involve some uncertainties in β_{eff}, etc; and rod drop measurements or inverse kinetic measurements are often biased because of the location of the detector, or the point at which the system is pulsed.

Gelbard: That is true, but the point of view that I was taking was that you would like the Monte Carlo to be exact, in the sense that the Monte Carlo error is given reliably by a stated error estimate. If this is true, then the Monte Carlo answer is at least one fixed point in your analysis. If it is not true, then in addition to a whole bunch of discrepant experimental results, you now have one more number, a discrepant Monte Carlo result; and instead of adding some sort of firm information you have simply put in another unknown. I would like to be able to say, without hedging, that for a given set of cross sections the eigenvalue lies in a certain band, with a known probability.

Coveyou: I think that I agree with what you said. I have always looked at the Monte Carlo calculation as an experiment in a universe which we have tried to make as close to the real physical universe as possible, though we don't know whether it is close enough or not. And in these experiments we should be able to compute our experimental errors very precisely. There is a point I would like to make, and this is sort of a philosophical point, with regard to many of these questions that we have raised. There seems to be some worrisome, but probably not terribly important, questions about how you estimate k, and so on. It has struck me that, because most of us have done fixed-source calculations, we have gotten used to the fact that, in essence, in the background, the systems that we are working with obey some kind of ergodic law. In other words, we think of our system as a member of an ensemble of systems in a certain physical situation. What we really want to do is find the average value of some quantity over the systems of this ensemble. Instead we pick one of these systems and examine its time behavior. Now, in the fixed source case, most of the ergodic theorems happen to be true, and therefore we do not have any great difficulty. However, in the case of multiplying systems, the ergodic hypothesis in almost any form that you can imagine, is just not true at all. The point here is that if we have a just-critical reactor, and we follow it long enough, the estimate we are going to get for k is zero. If the reactor is just critical it is easy to prove that a time will come when all the neutrons disappear.

Gelbard: That all depends on how close your model is to an analog model.

Coveyou: No, that is not the point. The point here is that these are very rare events. But I think that this is the basis for our difficulties for biases and other worrisome problems. The estimators we tend to choose are probably those that are appropriate for what I might call ergodic systems.

Gelbard: I think a more serious problem than the bias, in practice, really is the problem of statistical fluctuation.

Coveyou: That too.

Gelbard: This is a terrible problem in thermal reactors.

Coveyou: In fact it is probably also true that the correlations you find by what is essentially a time analysis of a single reactor are related, in an extremely complicated fashion, to the correlations you would find in the original ensemble of possible systems. And again the ergodic hypothesis comes in. The point is we have to make some kind of a hypothesis. We have to prove some kind of ergodic theorems. Either they have not been proved, or we are not aware of them and don't know how to use them.

Gast: May I make one comment here. I think that if the Monte Carlo process is sufficiently nonanalog, then there is no reason to get the zero multiplication factor or the zero flux after a certain number of iterations. For example, since source neutrons are always borne in fuel, you could use some fairly elaborate track-length-type estimator on the first flight. You could always revert to this first-flight track-length estimator if you get no fissions using some simpler estimation process. By such a procedure you could always bypass any termination of the fission chain.

Gelbard: In other words, if you use a track length estimator to generate offspring you will never get zero offspring.

Gast: Well in any case that to me is sort of an academic point. It is a very troublesome point when you try to prove convergence mathematically, so it is something that you do have to consider.

Borgwaldt: I would also agree that this seems to me to be an academic point because if we look at a Boltzmann equation in its normal formulation, we know that there is an eigenvalue. I would say that it is not our task to imitate a reactor, but to find solutions to the Boltzmann equation. We should do everything along the lines that Gast mentioned to avoid this situation which can occur in a theoretical reactor, but which would never be a reality because of cosmic radiation and all sorts of processes which will produce neutrons.

Coveyou: I cannot resist this. There have been two comments to the effect that the subject I introduced was an academic subject. At this meeting, and at three or four other meetings, a certain plausible method has been introduced for making estimates, and each time I have tried to show that the proposed method produces biased estimates. It is precisely the fact that these questions are considered to be academic that results in the reappearance of such methods. I don't say that this particular method is completely faulty; but what I do say is that the authors in both cases were thoroughly convinced that the method was unbiased, when in fact it was biased. I don't think that these discussions are entirely academic.

Gast: An unbiased estimator, by definition, has the following property. If one repeats the whole Monte Carlo calculation over and over again with different starting random numbers, and then averages all the different Monte Carlo results, the overall mean approaches some "true value" stochastically. I consider the question as to whether an estimator in an eigenvalue problem is biased or not, I consider this also an academic question. To me the very crucial question is that of fair game or consistant estimate.

Coveyou: I have yet to see a computational procedure which I was convinced actually satisfied your fair game property. I do not mean to imply that such procedures do not exist. I don't mean to imply that yours is not a fair game. I do mean to imply that I am not convinced that it is. I also mean to imply that I regard this property of being a fair game as being a far more academic property than the property of unbiaseness. Actually, if an estimate is biased, you are making a systematic error. If you claim an estimate is unbiased, when in fact it is biased, then every time that estimator is used there will be an error in it. The estimate will be in error by a certain fixed amount, namely the amount of the bias. I don't see anything academic about that.

Gelbard: Incidentally, I always thought that a fair game meant a method which is unbiased if you carry out an infinite number of generations. Is that what you mean?

Gast: Yes, as the number of generations goes to infinity one will approach the true answer. But I want to point out that an eigenfunction strategy can never be unbiased. Suppose you start with a given source guess, and keep repeating the problem back from the beginning with new random number sequences. Say you go through 100 iterations and start averaging the results. At the end of the 100 iterations you will still be sensing the source guess; but you can make the game a fair game. You can arrange the process in such a way that it will converge in the limit as the number of iterations goes to infinity. Now the only thing that is really crucial is to have an estimate of uncertainty at whatever point you stop the computation. To me the question as to whether a method is unbiased is academic because we have basically only one problem that we are considering. We have one final accumulated result and we desire a confidence interval large enough to cover the true value.

Gelbard: You are saying that it is important to have reliable error estimate and that if you do not have a fair game there is no way to get such an estimate?

Borgwaldt: My feeling, which I have often expressed, at Karlsruhe, is that if I make an estimate of any quantity by Monte Carlo, or by any other method, I commit some errors. These errors may consist of three contributions. One is an error in the input data, which we must eliminate. The other error is an error of numerical kind. In Monte Carlo this is a statistical error. In other methods it is some sort of truncation error due, perhaps, to the fact that I have used an S_n approximation of insufficiently high order, or a space grid which is not sufficiently fine. Then there may be an error due to a fundamental weakness in the method which I use. This would be a bias in a Monte Carlo calculation; and what we have to do is to minimize the total error. Our problem so far, I think, is to minimize the statistical error. Because if one has a reasonably sound eigenfunction strategy there is certainly a bias but it is obviously lower than the 0.5% or the 0.2% statistical error which we can get in a k_{eff} calculation.

Gelbard: Just one comment. My coauthor pointed out to me a very interesting fact. We talk frequently about the bias in eigenvalue. Very few people have given any consideration to the corresponding bias in the eigenvector. If the bias in the eigenvalue is 0.1%, what is the bias in the eigenvector? I don't know. I think that it is worth pointing out that we should not think only about the eigenvalue.

OFFICIAL PARTICIPANTS

Richard C. Bending
U. K. Central Electricity
Generating Board
Berkeley Nuclear Laboratories
Berkeley, Gloucestershire, England

Wolfgang Bernnat
Institut für Kernenergetik
Universitat Stuttgart
7000 Stuttgart 80
Pfaffenwaldring 31, Germany

Horst A. Borgwaldt
Institut für Neutronenphysik
und Reaktortechnik
Kernforschungszentrum Karlsruhe
Postfach 3640
75 Karlsruhe 1, West Germany

Edmond D. Cashwell
Los Alamos Scientific Laboratory
P. O. Box 1663
Los Alamos, New Mexico 87544, USA

Robert R. Coveyou
Nuclear Division
Union Carbide Corporation
Oak Ridge, Tennessee 37830, USA

Ely M. Gelbard
Applied Physics Division
Argonne National Laboratory
Argonne, Illinois 60439, USA

Robert C. Gast
Bettis Atomic Power Laboratory
Westinghouse Electric Corporation
P. O. Box 79
West Mifflin, Pennsylvania 15122, USA

James M. Grimstone
Atomic Energy Establishment
United Kingdom Atomic Energy Authority
Winfrith, Dorchester, Dorset, England

J. Eduard Hoogenboom
Interuniversitair Reactor Instituut
Delft Technological University
Berlageweg 15
Delft, The Netherlands

Malvin H. Kalos
Courant Institute of
Mathematical Sciences
University Computing Center
New York University
251 Mercer Street
New York, New York 10012, USA

John G. Moore
Safety and Reliability Directorate
United Kingdom Atomic Energy Agency
Wigshaw Lane, Culcheth, Warrington
Lancashire WA3 4 NE, England

Rolando Simonini
Comitato Nazionale per
l'Energia Nucleare
Centro di Calcolo
Via Mazzini
2 - 40138 Bologna, Italy

Herbert Steinberg
Mathematical Applications Group, Inc.
3 Winchester Plaza
Elmsford, New York 10523, USA

Antonio Taormina
Eidg, Institut für Reaktorforschung
Institut Federal de Recherches
en Matiere de Reacteurs
5303 Wurenlingen, Switzerland

G. Elliot Whitesides
Oak Ridge National Laboratory
P. O. Box X
Oak Ridge, Tennessee 37930, USA

OTHER ATTENDEES

Norman E. Banks
USA Ballistic Research Laboratories
Aberdeen Proving Ground
Maryland 21005, USA

William Chandler
Lawrence Livermore Laboratory
University of California
P. O. Box 808
Livermore, California 94550, USA

Thomas J. Hoffmann
Department of Nuclear Engineering
University of Tennessee
Knoxville, Tennessee 37916, USA

Lewis Milton
Applied Physics Division
Argonne National Laboratory
Argonne, Illinois 60439, USA

Richard E. Prael
Applied Physics Division
Argonne National Laboratory
Argonne, Illinois 60439, USA

ARGONNE NATIONAL LAB WEST
3 4444 00007225 6