A Monte Carlo Code for Particle Transport

an algorithm for all seasons

John S. Hendricks

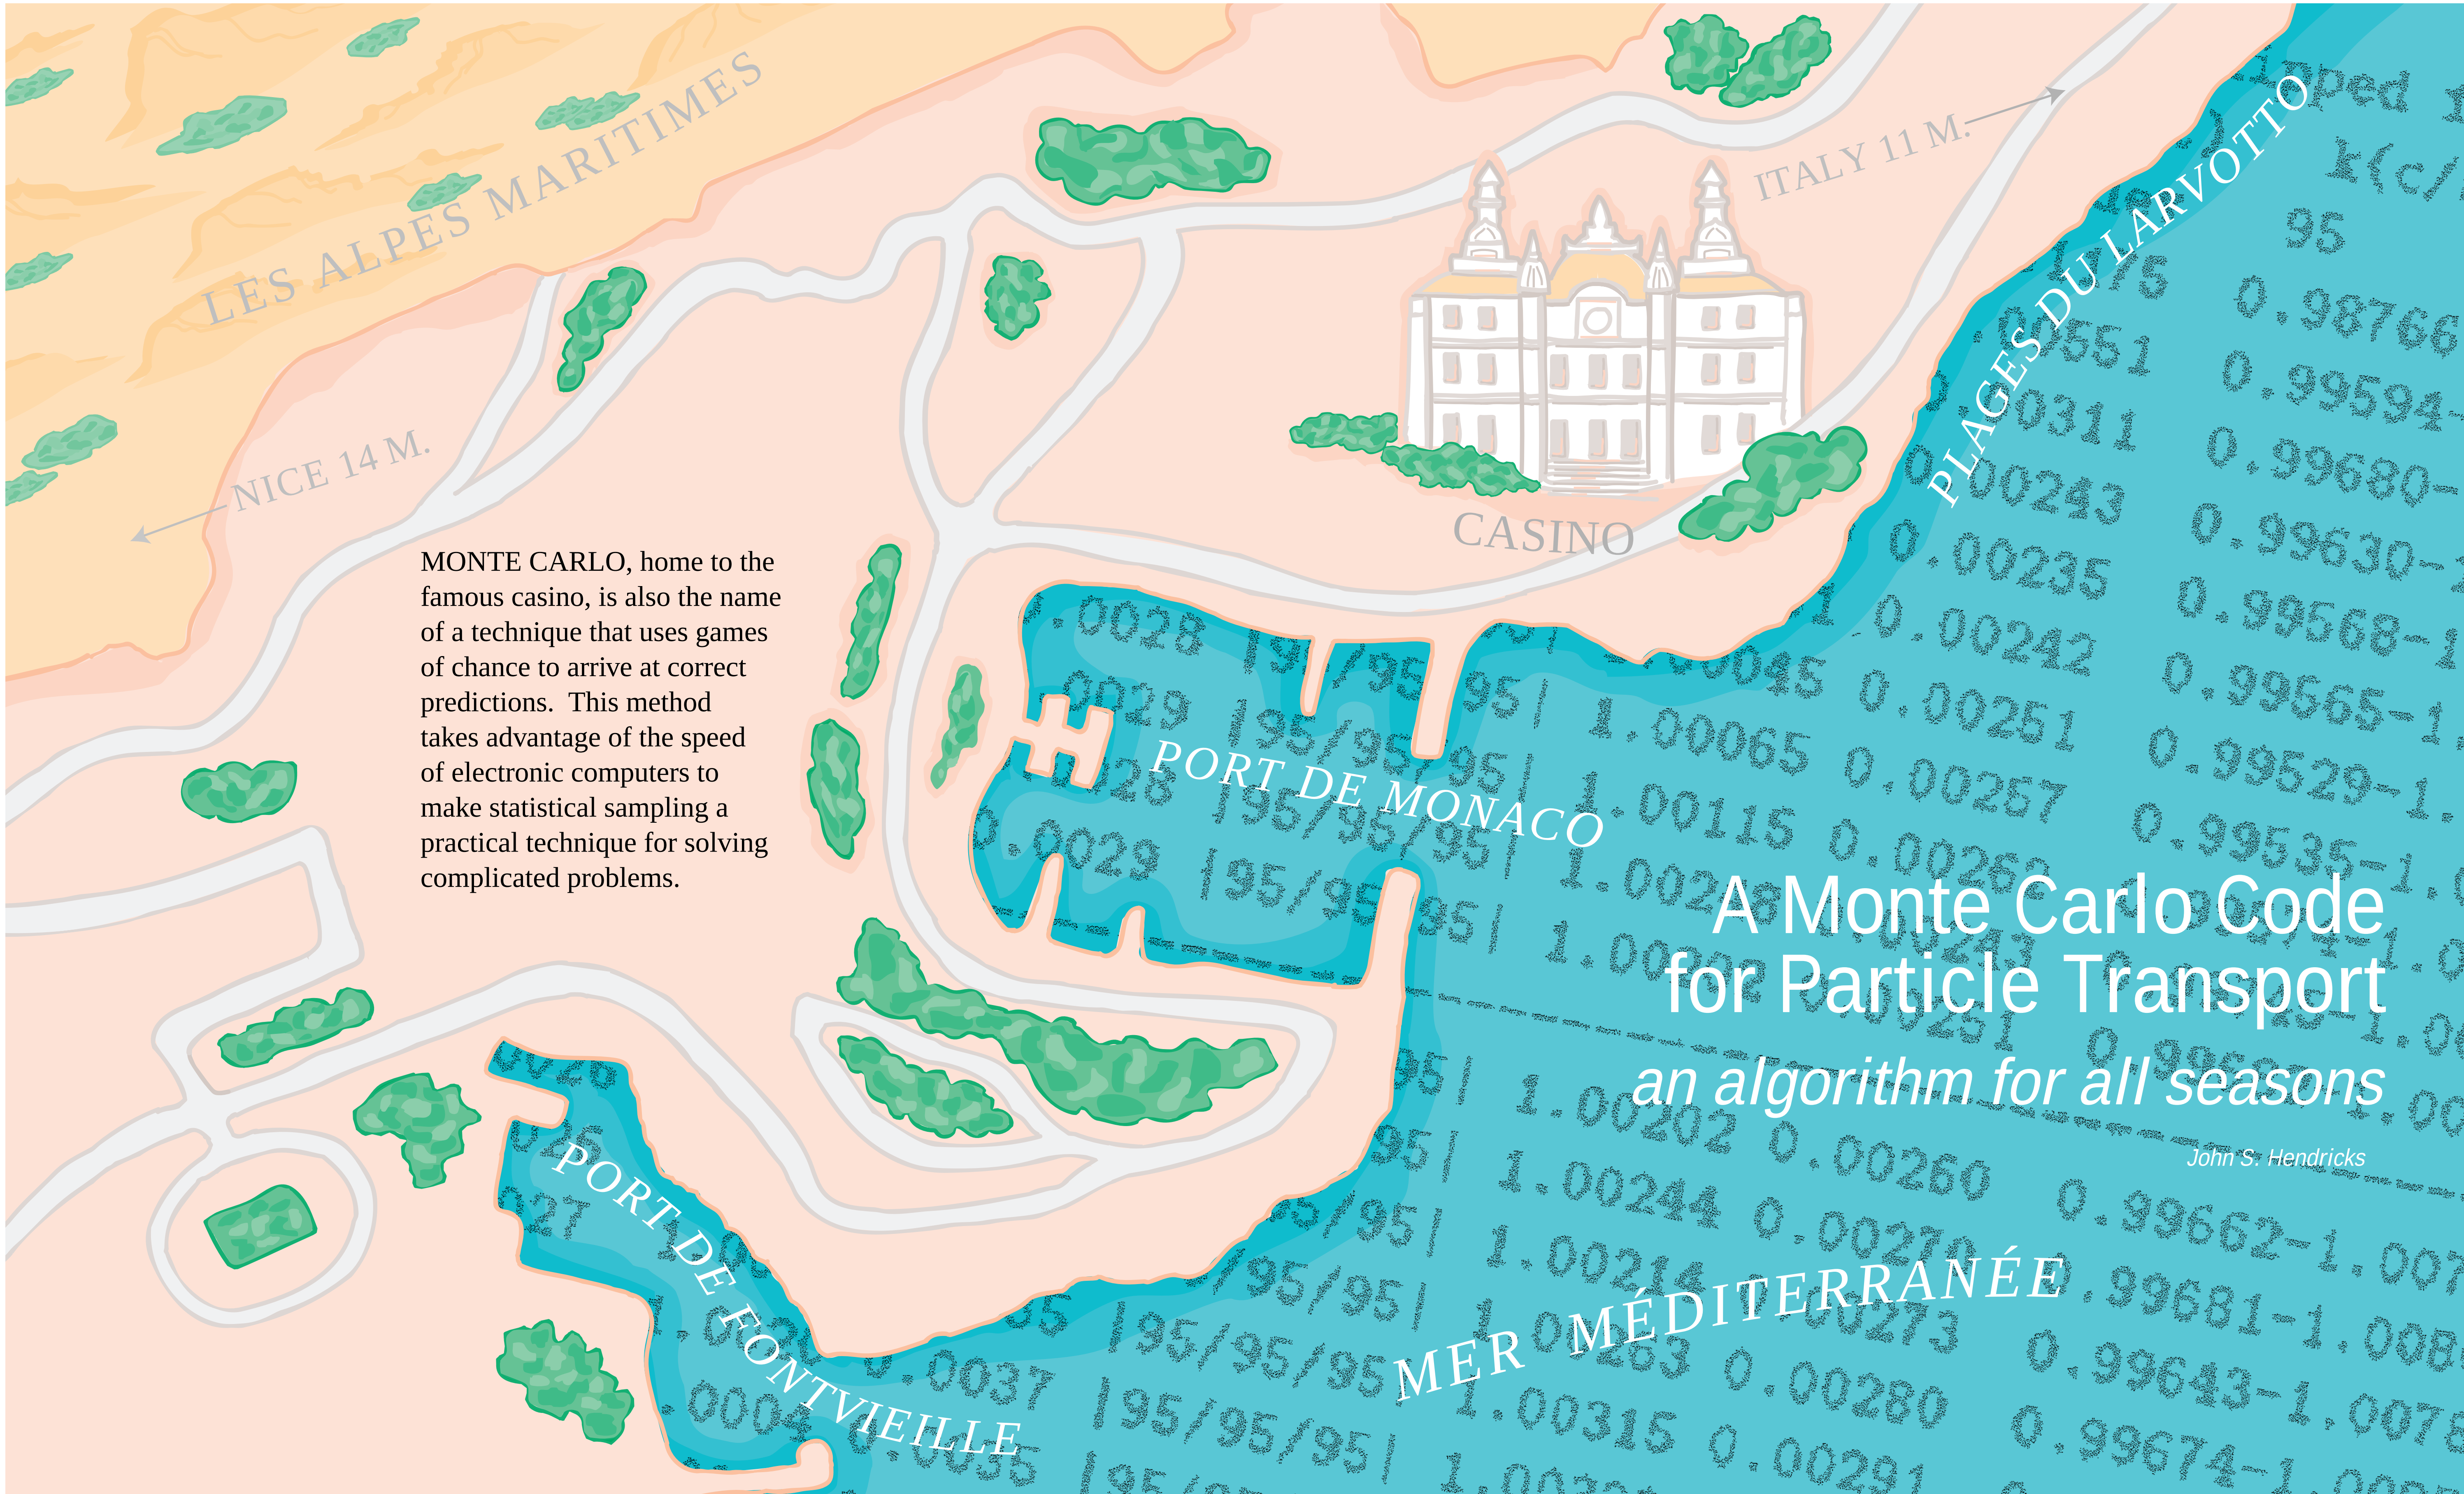

MONTE CARLO, home to the famous casino, is also the name of a technique that uses games of chance to arrive at correct predictions. This method takes advantage of the speed of electronic computers to make statistical sampling a practical technique for solving complicated problems.

The Monte Carlo method was first applied on the MANIAC computer at the Laboratory to predict the rate of neutron chain reactions in fission devices. Now, half a century later, its recent incarnation as the MCNP code has become a significant technology-transfer success.

Fast computers and sophisticated computational techniques have launched a revolution in scientific research. No longer does that endeavor depend only on physical experimentation as a basis for developing and refining theory. Instead, scientific theories can now be based on numerical experiments, which, compared with physical experiments, not only are less expensive, considerably safer, and more flexible but also provide more information, a better understanding of physical phenomena, and access to a wider range of experimental conditions. In addition, analysis of the results of numerical experiments can guide the selection and design of the physical experiments best suited to validating theories. Thus numerical experimentation (also referred to as numerical modeling or numerical simulation) has added a new dimension to scientific research.

Two general approaches to numerical modeling are available—deterministic methods and the Monte Carlo method. Deterministic methods involve solution of an integral or a differential equation that describes the dependence on spatial coordinates or time of some behavioral characteristic of the system in question. The equation is cast in an approximate form that permits calculation of the incremental change in the characteristic caused by an incremental change in the variable(s). The value of the characteristic itself is then calculated at each of successive points on a spatial or temporal grid. The accuracy of deterministic methods is limited by how well the equation approximates physical reality and by the practical necessity of making the spatial or temporal difference between grid points finite rather than infinitesimal. Well-known deterministic methods include the finite-difference and finite-element methods.

The Monte Carlo method involves calculating the average or probable behavior of a system by observing the outcomes of a large number of trials at a game of chance that simulates the physical events responsible for the behavior. Each trial of the game of chance is played out on a computer according to the values of a sequence of random numbers. For that reason a Monte Carlo calculation has been defined in general as one that makes explicit use of random numbers. The Monte Carlo method is eminently suited to the study of stochastic processes, particularly the process called radiation transport—the motion of radiation, such as photons and neutrons, through matter. Another well-known use of the Monte Carlo method is the Metropolis technique for finding the equilibrium energy, at a given temperature, of a system of many interacting particles. The method also has a more strictly mathematical application, namely, estimating the value of complicated, many-dimensional integrals.

The principle behind the Monte Carlo method—statistical sampling—dates back to the late eighteenth century, but it was seldom applied because of the labor and time required. However, the mathematician Stanislaw Ulam, after returning to the Los Alamos laboratory in mid 1946, realized that the electronic computer, which had only recently become a reality, could turn statistical sampling into a practical tool. Ulam discussed his idea with the mathematician John von Neumann, a consultant to Los Alamos, who proceeded to outline a computerized statistical approach to a problem of immense interest to designers of nuclear weapons—neutron diffusion and multiplication (by fission) in an assembly containing a fissile material. Von Neumann sent his outline to the head of the Los Alamos Theoretical Division, and the idea was pursued enthusiastically. Among the Laboratory scientists who pioneered development of the Monte Carlo method at Los Alamos, Nicholas C. Metropolis was the one who gave it its entirely appropriate and slightly racy name. Continuous effort at Los Alamos National Laboratory since those early times has culminated today in the computer code called MCNP (for Monte Carlo N-particle), perhaps the world's most highly regarded Monte Carlo radiation-transport code.

This article is of necessity too short to do justice to the Monte Carlo method. It cannot properly acknowledge the scientists who have devoted their careers to developing it or those who have successfully applied it in a variety of fields. All that will be attempted here is to describe the method, showcase a few of the many applications of MCNP, and to explain what is involved in developing and maintaining a modern Monte Carlo radiation-transport code such as MCNP.

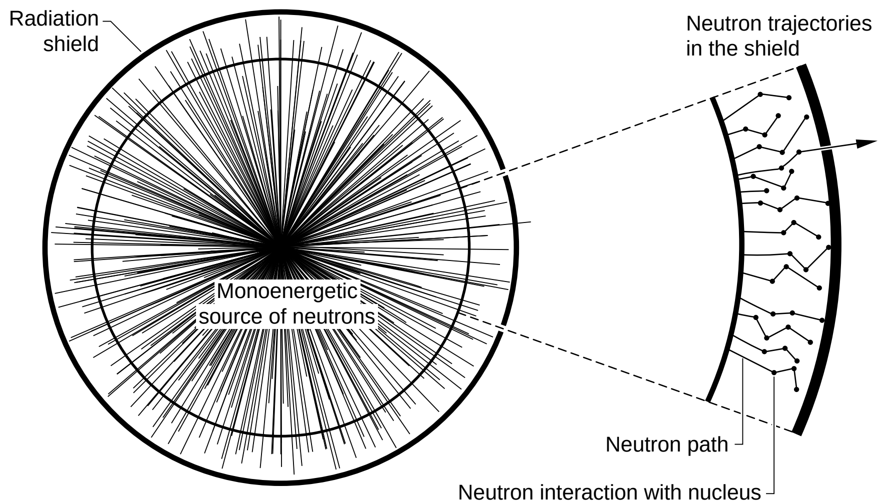

Figure 1. A Simple Monte Carlo transport Problem
In the problem discussed in the main text, a source of mono-energetic neutrons is at the center of a thin, spherical radiation shield. The object is to determine the probability that any given neutron emanating from the source will pass through the shield. As shown above, after entering the shield wall, each neutron will follow a succession of straight-line paths whose lengths and directions are in many respects random relative to each other. In other words, the neutron's trajectory resembles a "random walk". Both changes in direction and terminations of a given trajectory result from the neutron's interaction with the nuclei in the shield. The possible interactions with the nuclei in the shield are: elastic scattering, which changes the neutron's direction of motion but not its energy; inelastic scattering, which changes both the direction of motion and the energy of the neutron; absorption, which terminates the neutron's trajectory as the neutron is absorbed into the nucleus; and fission, which produces additional neutrons but occurs only if the shield contains certain isotopes. The length of the path between one interaction and the next as well as the outcome of each interaction are described by probability distributions that have been determined experimentally. The Monte Carlo method is used to construct a large set of possible trajectories of a neutron as it travels through the shield; the experimental probability distributions are sampled during the construction of each trajectory. The neutron's probability of escape is then determined based on the outcomes in that set of trajectories. The figure shows some possible neutron histories in the enlarged view of the shield wall. Most of them scatter a number of times before being absorbed and one escapes through the shield.

The Monte Carlo Method

Portraying the essence of the Monte Carlo method is perhaps best accomplished by focusing on its application to radiation transport. For concreteness, let us focus in particular on the problem of estimating the probability that a neutron emanating from some source passes through some radiation shield. For simplicity assume that the source is an isotropic point source, that it emits monoenergetic neutrons, and that it is located at the center of the shield, which is a relatively thick spherical shell made up of matter containing only one isotopic species. The physics of the problem is well known: Each neutron emitted by the source follows a trajectory within the shield that consists of a succession of straight-line paths whose lengths and directions appear to be random relative to each other (Figure 1). That "random walk" is the

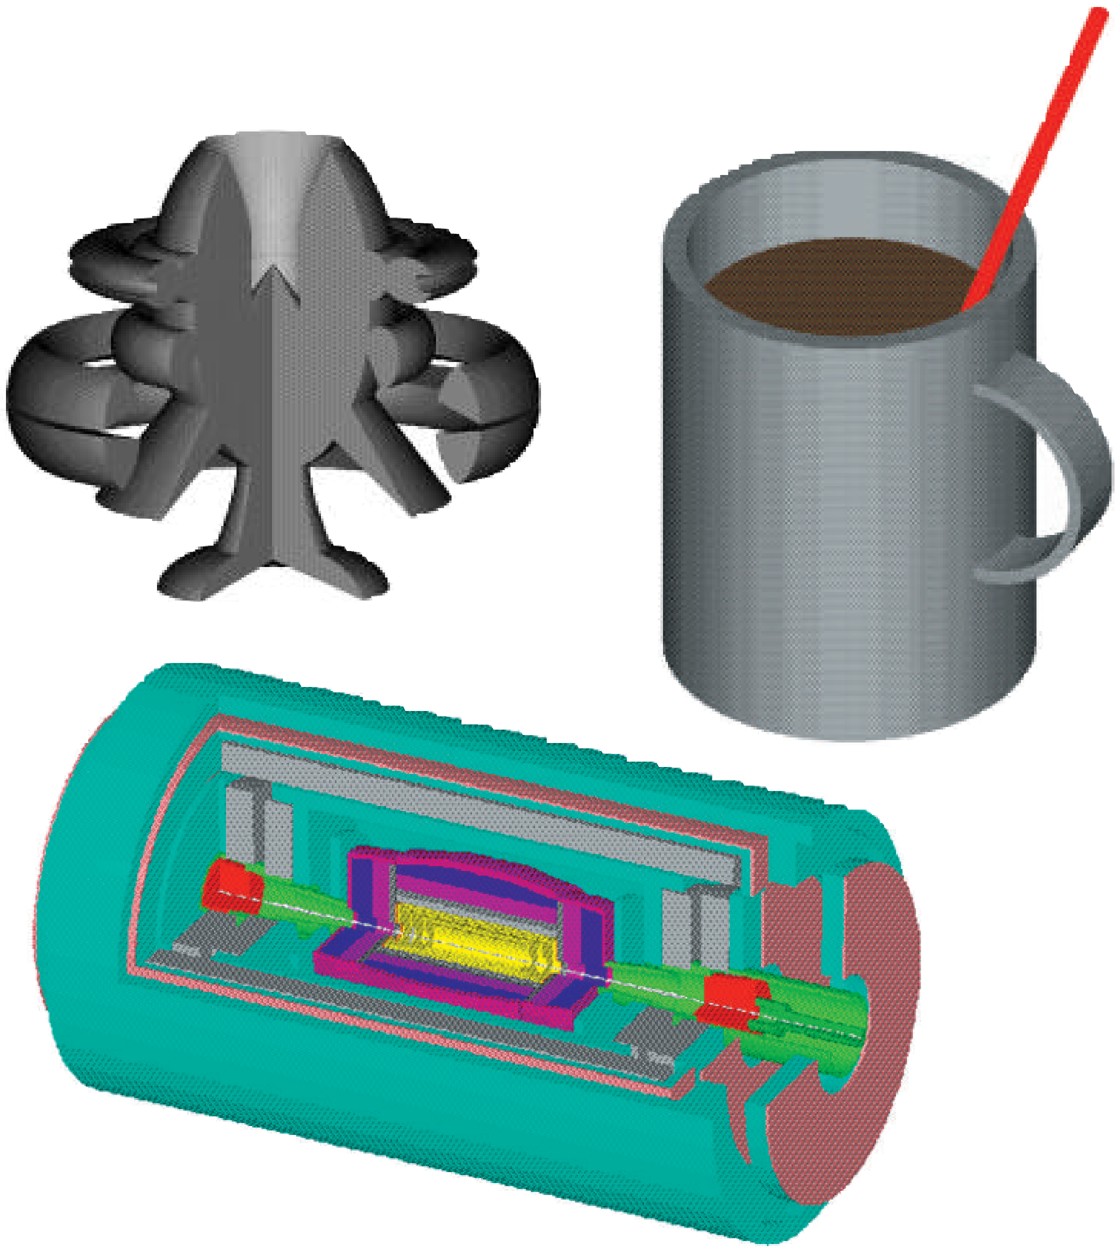

Figure 2. Geometry Modeling in MCNP

Shown here are examples of the geometry-modeling capability that can be accessed for MCNP calculations. The geometries were created and visualized with the graphics software known as Sabrina, which was developed by James T. West and Kenneth Van Riper. Each example is depicted in a different visualization format. Unlike the other examples, the example in color, a model of the SDC detector for the Superconducting Super Collider, was constructed not simply to demonstrate the capabilities of Sabrina but for use in an MCNP simulation. The detector, shown here in cutaway view, contains dozens of components and is very large (the outer cylindrical surface shown has a diameter of nearly 22 meters). The accelerator was designed to produce two counter-propagating beams of extremely energetic (20-TeV) protons that travel along the axis of the detector and collide at its midpoint. The detector model was used in MCNP simulations to track the products of the collisions and to calculate the levels of the radiation induced in the detector components by the collision products. The radiation levels are necessary to assess component damage and background signals in the detector.

result of interactions of the neutron with nuclei within the shield. The possible types of interactions of a neutron with a nucleus include elastic scattering, inelastic scattering, absorption, and fission. For simplicity, assume that the only interactions that occur are elastic scattering, which changes the direction but not the speed (and hence energy) of a neutron, and absorption, a nuclear reaction that swallows up, or "kills," a neutron. Whether an interaction results in a neutron's being absorbed or scattered can be predicted only probabilistically, as can the scattering angle if the neutron is scattered. The probabilities of a neutron's being absorbed by various isotopes have been measured, and so have the probabilities of its being scattered through various angles, all as functions of neutron energy. Those probabilities, or cross sections, for the shield nuclei are necessary input to solving the problem at hand. Also needed is the probability density function for the distance a neutron travels in the shield without undergoing an interaction with a nucleus (in other words, the probability density function for the lengths of the straight-line paths composing the neutron's trajectory). It is known that the probability density function for the "free-path" length in any material decreases exponentially. In particular, the probability density that a neutron will travel a distance x before undergoing an interaction is given by $\rho\sigma e^{-\rho\sigma x}dx$, where ρ is the density of nuclei and σ is the total cross section (here the sum of the scattering cross section integrated over scattering angle and the absorption cross section).

Application of the Monte Carlo method to the problem above involves using a sequence of numbers uniformly distributed on the interval (0, 1) to construct a hypothetical (but realistic) history for each of many neutrons as it travels through the shield. (To say that a sequence of numbers is uniformly distributed on (0, 1) means that any number between 0 and 1 has an equal probability of occurring in the sequence. Such numbers, when generated by a computer, are called pseudorandom numbers.) The ratio of the number of neutrons that escape from the shield to the number of neutrons whose histories have been constructed is an estimate of the answer to the problem, an estimate whose statistical accuracy increases as the number of neutron histories increases. Details of the process can be illustrated by following the construction of a single neutron history.

Constructing the first step of a neutron history involves deciding on a value for its first free-path length x_1. As pointed out above, the sequence of pseudorandom numbers generated by the computer is uniformly distributed on (0, 1), whereas the free-path lengths are distributed according to $e^{-\rho\sigma x}$ on $(0, \infty)$. How can the sequence of uniformly distributed numbers ξ_i be used to produce a sequence of numbers x_i whose distribution mirrors the experimentally observed distribution of free-path lengths? It can be shown that the transformation $x_i = -(1/\rho\sigma)\ln(1 - \xi_i)$ yields a sequence of numbers that have the desired inverse exponential distribution. So the first free-path length is obtained by setting $x_1 = -(1/\rho\sigma)\ln(1 - \xi_1)$. The second step in the neutron history involves deciding whether the neutron's first interaction with a nucleus scatters or kills the neutron. Suppose it is known from the cross sections for the shield nuclei that scattering is nine times more likely than absorption. The interval (0, 1) is then divided into two intervals, (0, 0.1) and [0.1, 1). Assume that x_2, the second pseudorandom number generated by the computer, is 0.2. Because 0.2 lies within the larger subinterval, the neutron is scattered rather than absorbed. The third step in the neutron history involves deciding through what angle it is scattered. Again some transformation must be performed on the third pseudorandom number, a transformation that changes the uniform distribution of the ξ_i into a distribution that mirrors the observed distribution of scattering angles (the scattering cross section as a function of scattering angle). Further steps in the history are generated until the neutron is absorbed or until the radial distance it has traveled within the shield exceeds the thickness of the shield. The histories of many more neutrons are generated in the same manner.

Assume that N neutron histories are generated and that n of the histories terminate in escape of the neutron from the shield. To calculate an estimate for the probability that any single neutron escapes, assign a "score" s_i to each neutron as follows: $s_i = 0$ if the neutron is absorbed within the shield, and $s_i = 1$ if the neutron escapes. Then the estimated probability of escape is given by the mean score $\bar{s}$, that is, by $(1/N)\Sigma s_i = n/N$. The relative error (relative statistical uncertainty) in that probability estimate is related to the so-called variance of the s_i, $\mathrm{Var}(s_i)$, which can be approximated, when N is large, by the difference between the mean of the squares of the scores and the square of the mean score:

$$\mathrm{Var}(s_i) \approx \frac{1}{N}\sum_i s_i^2 - \left(\frac{1}{N}\sum_i s_i\right)^2$$

$$= \frac{n(N-n)}{N^2}$$

The relative error in the probability estimate is then given by

$$\frac{\sqrt{\mathrm{Var}(s_i)/N}}{\bar{s}} = \sqrt{(N-n)/Nn}$$

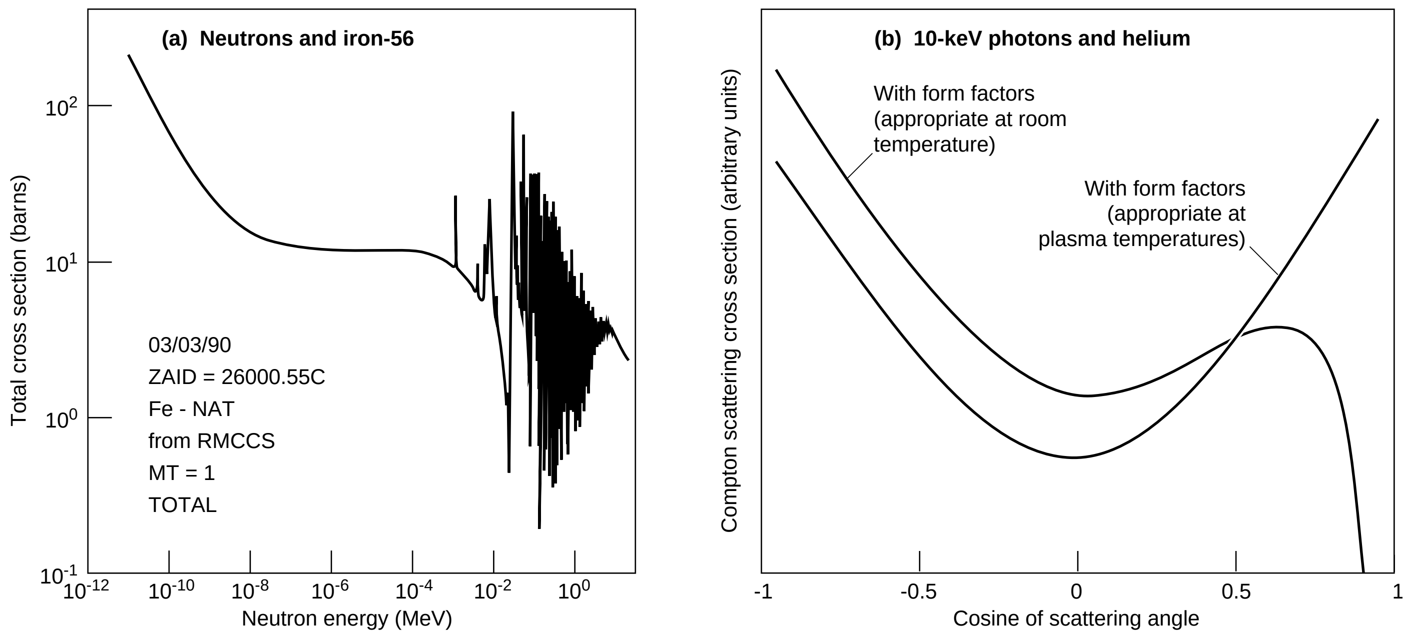

Figure 3. A Few Selections from the MCNP Data Library

The MCNP data library contains a vast amount of information about the interaction of radiation with matter. Shown here are three items included in that library: (a) the continuous-energy representation of the total cross section, as a function of neutron energy, for the interaction of neutrons with the most abundant stable isotope of iron, iron-56; (b) the "differential" cross section (cross section as a function of scattering angle), with and without the form factors that account for electron binding, for Compton scattering of 10-keV photons off helium; (c) the cross sections, as functions of neutron energy, for various interactions of neutrons with uranium-235. Total, elastic, and inelastic cross sections are shown, as well as cross sections for capture of the incoming neutron accompanied by the emission of a photon and for reactions in which one incoming neutron results in two or three outgoing neutrons, (n, 2n) and (n, 3n).

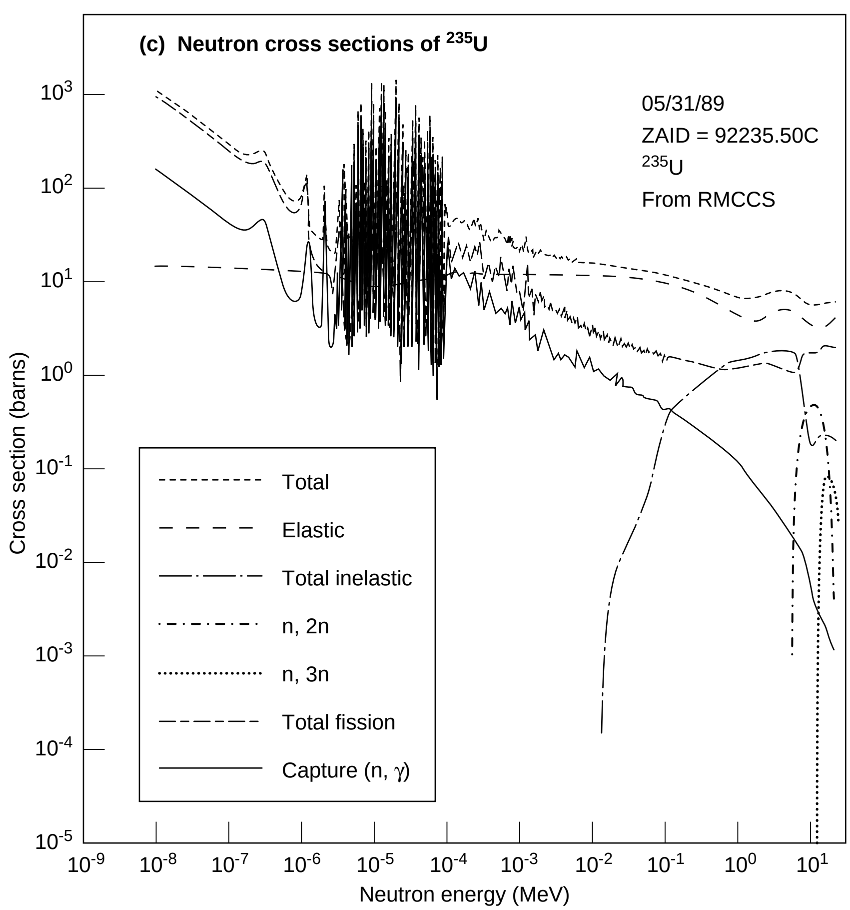

For example, if $N = 100$ and $n = 47$, the probability of escape is estimated as 0.47 and the relative error is a little less than 11 percent.

Now is a good time to point out a major difference between a Monte Carlo calculation and one based on a deterministic method. If the calculation outlined above is repeated, the second calculation is highly unlikely to show that exactly 47 of the 100 histories result in escape. In other words, repetitions of a Monte Carlo calculation yield only approximately the same answer, whereas repetitions of a calculation based on a deterministic method yield exactly the same answer. That difference is due to the element of chance inherent in a Monte Carlo calculation. Thus deterministic methods provide more exact solutions of approximate models, whereas Monte Carlo methods provide approximate solutions of more exact models.

The simplifying assumptions invoked in the above example of a Monte Carlo radiation-transport calculation do not of course hold in general. The radiation may consist of particles other than neutrons (or of more than one type of particle), and other types of interactions may be involved (neutron-induced fission of fissile nuclei, for example). The radiation may not be monoenergetic, and it may emanate from a source that is neither point-like nor isotropic. The material through which the radiation travels may be nonuniform in composition and intricate in geometry. All those additional complexities can be handled provided the necessary input data are available.

Perhaps the greatest advantage of using the Monte Carlo method to simulate radiation transport is its ability to handle complicated geometries. That ability rests on the fact that, even though the geometry in question may be complicated in its entirety, only the geometry in the vicinity of the particle for which a random walk is being constructed need be considered at any point in the construction. A given geometry can be modeled in its entirety in two ways: as a "combinatorial" object or as a "surface-sense" object. As its name implies, a combinatorial object is constructed by combining relatively simple geometric entities, such as rectangular parallelepipeds, ellipsoids, cylinders, cones, and so on. Combinatorial objects are easy to construct but are less general than the surface-sense objects provided by MCNP. A surface-sense object is constructed by combining bounding surfaces, each of which is assigned one of two sense values to indicate on which side of the bounding surface the object lies. Any combination of linear, quadratic, or toroidal surfaces in any orientation or even skew can be accessed by users of MCNP. Figure 2 shows examples of the intricate geometries that can be constructed for MCNP calculations.

It is worth noting that a Monte Carlo radiation-transport code is more than a list of instructions for executing a certain set of arithmetic and logic operations. It is also a repository of experimental data about and theoretical understanding of radiation transport accumulated over the years. That knowledge is accessed unwittingly by users whose expertise may lie in areas far removed from radiation transport. The data are not built into modern Monte Carlo radiation transport codes but rather are available as separate data libraries. That stratagem permits the data libraries to be upgraded independently of the operational portion of the code. Many older Monte Carlo codes include "multigroup" data, which are averaged over ranges of energy or of some other parameter. That practice saves considerable data-storage space but introduces approximations into the calculations. MCNP and other modern Monte Carlo radiation-transport codes offer libraries of unaveraged, or "continuous-energy" data. Shown in Figure 3 are a few specific examples of the data contained in the MCNP data library.

The example of neutron transport through a spherical-shell shield shows that a Monte Carlo radiation-transport calculation is similar to an opinion poll, such as a Gallup poll. Both use the technique of statistical sampling to obtain a probabilistic answer to some question. Just as a Gallup poll attempts to predict the outcome of, say, an election by questioning only a small portion of the voting population, so also does a Monte Carlo radiation-transport calculation attempt to simulate the outcome of the interaction of 10^{15} to 10^{25} neutrons with some material by constructing histories for 10^5 to 10^8 neutrons. And just as the voters questioned during a Gallup poll are not chosen randomly but are carefully selected to mirror demographic characteristics of the entire voting population, so also are the neutron histories not truly random walks but histories that mirror the observed phenomenology of neutron interactions with matter. The analogy between a Gallup poll and a Monte Carlo radiation-transport calculation will be continued in the following discussion of variance reduction.

Variance Reduction

As pointed out above, the result of a Monte Carlo calculation has associated with it a statistical uncertainty. How can that uncertainty be reduced and the result thereby be made more accurate? One obvious way to do so is to increase the number of neutron histories generated (or voters questioned). But that "brute-force" approach is costly in terms of computer time (or pollster time). Fortunately, more sophisticated techniques are available to achieve a lower uncertainty without increasing the number of histories (voters questioned) or to achieve the same uncertainty from a smaller number of histories. Four types of such "variance-reduction" techniques are available: truncation, population control, probability modification, and pseudodeterministic methods.

Truncation involves ignoring aspects of the problem that are irrelevant or inconsequential. For example, the source-and-shield assembly described above may include structural elements that position the source at the center of the shield. Because the nuclei in the structural elements, like the nuclei in the shield, interact with the neutrons, the structural elements should be included in the simulation. Suppose, however, that the structural elements are very fine rods and hence are considerably less massive than the shield itself. Then truncating the problem by ignoring the existence of the structural elements would have little effect on the results. An example of truncation in a Gallup poll is to not include among the sampled population those who live abroad and yet are eligible to vote because such persons are unlikely to affect the outcome of an election.

Population control involves sampling more important portions of the sampled population more often or less important portions less often. For example, suppose the neutrons that escape from the left half of the spherical-shell shield are of greater interest (say, because someone's office is located there, whereas a little-used stairwell is located to the right) and that the left half of the spherical shell is composed of a material more effective at absorbing neutrons. Each neutron that has a possibility of reaching the region of greater interest (any neutron that is emitted toward the left) is "split" into m neutrons ($m > 1$) and assigned a "weight" of $1/m$. The scores of the histories of the split neutrons are multiplied by their weight so that the splitting stratagem does not alter the physical situation but does allow the sampling of more of the more important neutrons. The corresponding technique for sampling fewer of the less important neutrons is referred to as Russian roulette. Applied to the same example, Russian roulette involves specifying that the neutrons emitted to the right have a probability of $(1 - 1/m)$ of being terminated immediately upon entering the shield. A neutron whose history begins with immediate death is of course tracked no further. Those neutrons that do not suffer immediate death, $(1/m)$ of the neutrons emitted to the right (in the limit of large N), are assigned a weight of m. Thus the simulation of the real physical situation is unaltered. An application of population control in a Gallup poll might be as follows. Suppose 40 percent of the voters are Democrats, 40 percent are Republicans, and 20 percent are independents. Then the outcome of the vote on an issue such that Democrats and Republicans are likely to vote the party line is determined primarily by the independent vote. Therefore a good polling strategy would be to question a sample consisting of 20 percent Democrats, 20 percent Republicans, and 60 percent independents (such a sample could be achieved by including every independent on a list of 300 registered voters but including each party member only if the roll of a die yielded a chosen one of the six possible outcomes) and weight the opinions of the 20 Democrats and the 20 Republicans by a factor of 2 and the opinions of the 60 independents by a factor of 1/3. The discarding of 100 Democrats and 100 Republicans corresponds to their death by Russian roulette.

Probability modification involves sampling from a fictitious but convenient distribution rather than the true distribution and weighting the results accordingly. For example, instead of applying splitting and Russian roulette to the neutrons emitted to the left and right, respectively, by the isotropic neutron source, the spatially uniform neutron distribution is replaced, for the purpose of constructing histories, by a distribution such that more neutrons are emitted to the left. The "bias" that such a strategy would introduce into the result is removed by appropriately weighting the scores of the neutron histories. An example of probability modification in the Gallup-poll analogy might be the following. Suppose only 40 percent of the voters questioned by a pollster happened to be women. Instead of questioning sufficiently more women to bring the sexual distribution of the sampled population to a 50-50 distribution, the response of each woman could be weighted by a factor of 5/4 and that of each man by 5/6.

Pseudodeterministic methods are among the most complicated variance-reduction techniques. They involve replacing a portion (or portions) of the random walk by deterministic or expected-value results. Suppose, for example, that the spherical shell is surrounded by further shielding material with complex geometry. Instead of transporting each neutron via a random walk through the spherical shell to the more complex region of the shield, the neutron may simply be put at the interface between the two shield components and assigned a weight equal to the (presumably known) probability of its arriving there. Similarly, it is known that not all those eligible to vote are equally likely to carry out their civic duty. So the response of each person polled may be assigned a weight equal to the probability, based on factors such as gender, race, age, and place of residence, that he or she will indeed vote. The difficulties encountered when using pseudodeterministic methods arise in assigning the probabilities.

The use of modern variance-reduction methods has allowed Monte Carlo calculations to be carried out many orders of magnitude faster and yet maintain the same statistical accuracy. In fact, many calculations that once would have required prohibitive amounts of computer time are now routine.

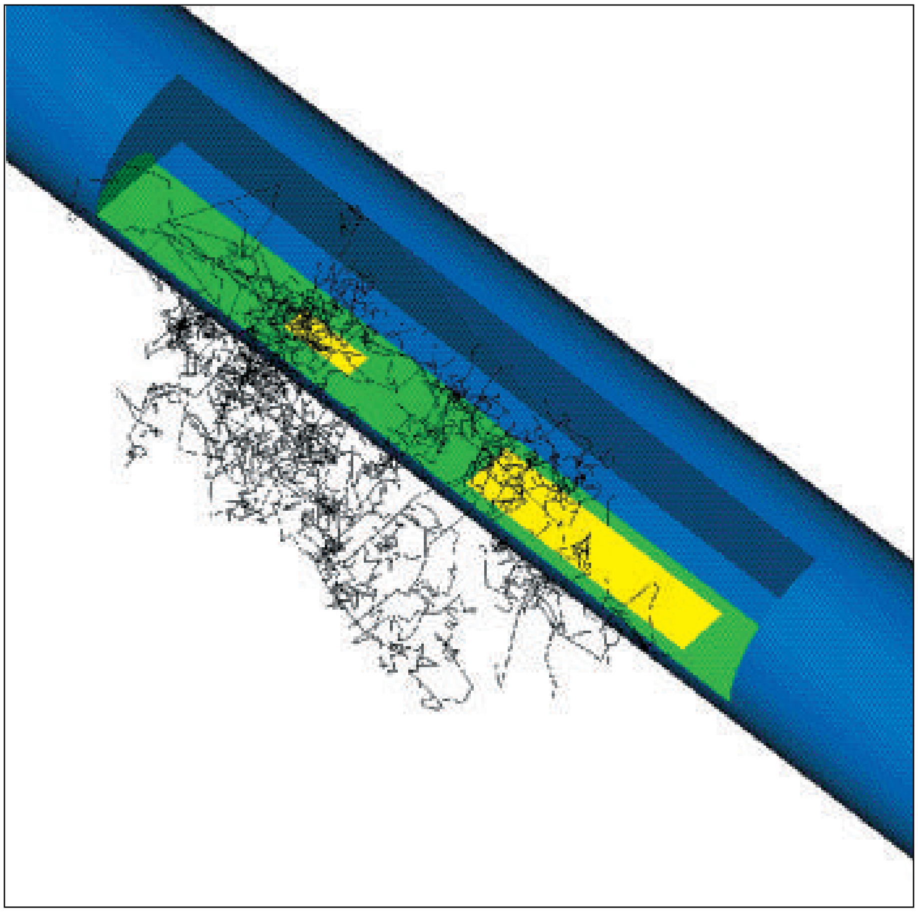

Figure 4. Nuclear Tool for Oil-Well Logging
This three-dimensional view of a generic nuclear tool for oil-well logging was generated with the graphics software Sabrina, which is linked to MCNP. The cutout displays a helium-3 neutron-detector assembly (yellow) encased in iron (green) and immersed in a water-filled borehole (blue), which is surrounded by a limestone formation (brown). Neutrons emanating from a source are scattered from the rock formation into the detector assembly. The red lines are neutron paths, as simulated by MCNP.

Applications of Monte Carlo Radiation-Transport Codes

The many applications of the Monte Carlo method in radiation transport reflect the pervasiveness of radiation in nature and in established and emerging technologies. A few of the applications of the Laboratory's MCNP code are discussed below.

Criticality safety. Preventing the inadvertent assembly of critical masses of fissile material is a grave concern at facilities that process fissile materials, such as the laboratories of the DOE's nuclear-weapons complex and facilities associated with the generation of nuclear power, including not only the power plants themselves but also waste-storage sites and fuel-fabrication plants. An MCNP simulation is a safe, reliable way to assess, for example, whether a given amount of fissile material assembled in a given geometry constitutes a critical mass or whether a chemical reaction involving a fissile material or a change in physical state of a fissile material can lead to the assembly of a critical mass. Answering such questions by physical experimentation is so dangerous as to be out of the question in many instances.

The MCNP code is used worldwide to predict the transport of radiation in a dazzling array of applications from oil-well logging and fusion research to the most advanced medical diagnostics and treatments. It is adapted to run on supercomputers, but even more important, it can be used by nonspecialists on ordinary personal computers. With over a thousand users at a hundred institutions around the world, the MCNP code is one of the most successful in the history of scientific computing.

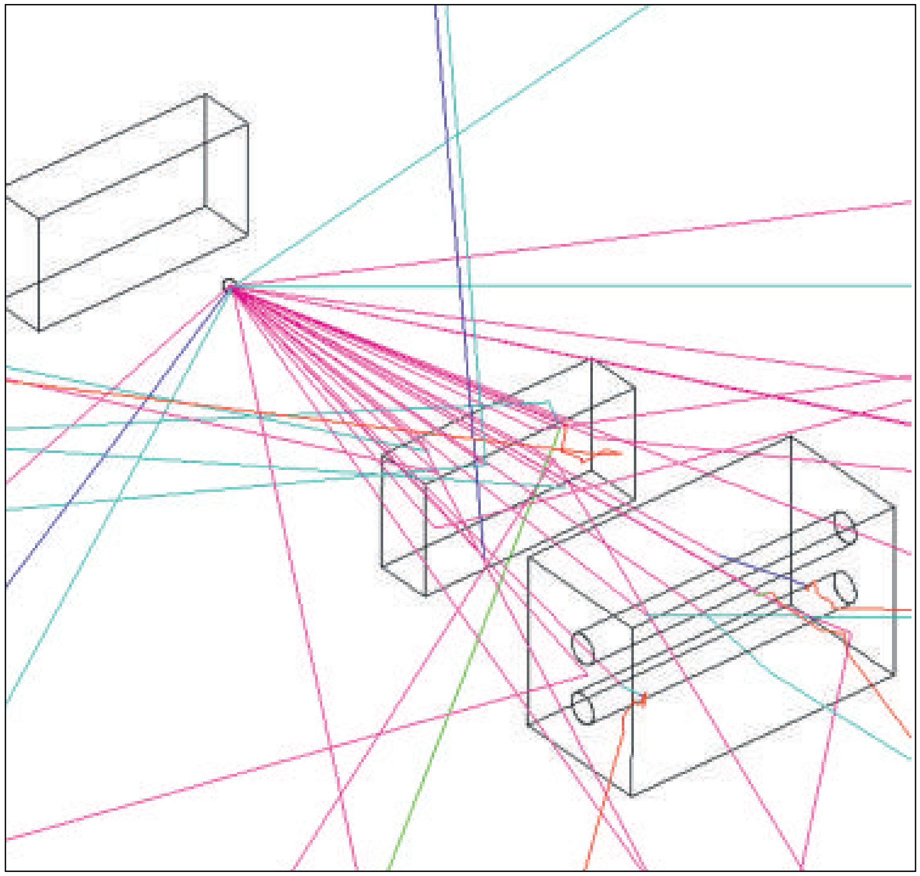

Figure 5. Example of a Particle-History Plot
Shown here are the MCNP-simulated histories of neutrons as they travel from a source, pass through a lead brick, and interact with helium-3 within two neutron detectors (thin cylinders). The neutron tracks are color-coded according to energy. The plot was generated with the graphics software Sabrina.

Oil-well logging. Surprising as it may be, radiation and hence MCNP play important roles in the search for oil. (Gone are the days when "finding oil" was equated with the gushing of oil from an exploratory hole.) The pattern of scattering of neutrons and gamma rays from rock formations can indicate the presence or absence of oil. MCNP can predict the scattering pattern characteristic of an oil-bearing rock formation, which is then compared with the scattering pattern obtained with a "logging" tool inserted into an exploratory bore hole drilled in the rock formation. As shown in Figure 4, the logging tool contains a source of neutrons and a detector assembly for observing the radiation scattered from the rock. Such a nuclear logging tool is particularly effective when used in conjunction with other logging tools that contain a source of radio waves or sound waves, whose scattering provides additional information. MCNP is used by many companies around the world in the search for oil. The Gas Research Institute and Schlumberger-Doll Research, a major oil-service company, have contracts with the Laboratory for continued development of MCNP.

Nuclear energy. MCNP serves a variety of purposes in the hundreds of nuclear power plants around the world. It is used, for example, to design shielding and spent-fuel storage ponds and to estimate the radiation dose received by operating personnel. General Electric Nuclear Energy of San Jose, California, and other fuel-assembly manufacturers use MCNP to design assemblies so that the fuel is "burned" more efficiently. More efficient burning implies cheaper electricity and less nuclear waste.

Nuclear safeguards. The signing of the Treaty on the Nonproliferation of Nuclear Weapons in 1968 was one of the most important political accomplishments of this century. Those signatory nations that had not yet developed their own nuclear weapons agreed to forego such development in return for access to peaceful nuclear technology. A cornerstone of the treaty is inspection of nuclear facilities by the International Atomic Energy Agency to ensure that no "special" nuclear materials are stolen from the facilities. The inspectors use nondestructive techniques based on radiation detection to assay the uranium-235 or plutonium-239 content of, say, fuel assemblies, barrels of nuclear waste, and workers' lunch boxes. The nondestructive assay techniques are based on detecting either the radiation emitted by those isotopes themselves in the course of radioactive decay or the radiation emitted by the products of nuclear reactions induced in the isotopes by irradiation with neutrons or gamma rays. Many of the detectors used by the inspectors were designed with the help of MCNP.

Fusion research. For years scientists have been investigating the possibility of producing energy by nuclear fusion, the mechanism that powers our sun. Nuclear fusion offers two advantages over nuclear fission as an energy source: It does not produce hazardous fission products, and it consumes a naturally abundant fuel—deuterium, an isotope of hydrogen that can be extracted from sea water. The most advanced strategy for fusing deuterium nuclei involves containment and compression of a plasma of deuterium ions by a toroidal magnetic field. Since MCNP is one of the few Monte Carlo radiation-transport codes capable of handling the challenging fourth-order toroidal geometry, it has long been the premier code for studying the transport of the neutrons and photons produced by the fusion reaction. Another use of MCNP in fusion research is studying ways to prevent damage to personnel and equipment by the fusion-produced radiation.

Medical technology. Few people realize that medicine is the third largest application of nuclear reactions, ranking behind only energy and defense. Among the nuclear medical technologies are boron neutron-capture therapy and positron-emission tomography. Boron neutron-capture therapy involves injecting a cancer patient with the stable isotope boron-10 and then irradiating the patient with neutrons. Because boron-10 collects preferentially in tumor cells and has an exceptionally high cross section for capturing (absorbing) neutrons, the therapy selectively kills tumor cells. MCNP is used to determine the neutron dose and energy spectrum that will kill the tumor and not the patient. Positron emission tomography is a nondestructive technique for observing metabolism in situ. It involves ingestion by the patient of water containing the radioactive isotope oxygen-15 rather than the usual, nonradioactive isotope oxygen-16. As the water is used by the body, a scanner detects the gamma rays resulting from interaction of the positrons emitted by the oxygen-15 with electrons in cells. For example, a PET scan may reveal that part of a heart is dead and that thus the heart requires replacement rather than repair. MCNP is used to properly interpret PET-scan images, which are blurred by the effects of electron and photon scattering. MCNP benefits many other medical technologies, including computer-assisted tomography, radiation therapies for cancer, and protection of medical personnel from nuclear radiation and x rays.

Space exploration. Among the harsh features of space are intense bursts of radiation. Simulating the effects of that radiation on equipment and personnel is a crucial preliminary to the exploration of space, particularly in light of its high cost. Calculations with MCNP and other codes have helped design shielding of minimum weight to protect astronauts from cosmic rays. In one case a "storm shelter" had been designed to protect astronauts from solar flares, but calculations with MCNP showed that although the proposed shielding would block most of the incident protons, so many gamma rays would be generated in the shield that the astronauts were safer outside the storm shelter than inside it! MCNP calculations have also shown that nuclear power will be essential to extended space exploration. The extra radiation astronauts would receive from a nuclear reactor is far less than the additional cosmic-ray radiation they would be exposed to during a slower, longer journey. Of course, MCNP plays a role also in designing nuclear-power systems for use in space and in assessing the survivability of sensitive electronic components and means to protect them from damage.

Applications at the Laboratory

Since MCNP was developed here to help carry out the Laboratory's primary mission, it is not surprising that the number of MCNP users here constitutes a large fraction of the total worldwide. Some of the applications at the Laboratory have already been mentioned: criticality-safety studies and the design of detectors for implementation of the Nonproliferation Treaty and for space exploration. The long list of other Laboratory uses of MCNP includes design of the neutron-producing target at the

Manuel Lujan, Jr. Neutron Scattering Center (LANSCE); design and evaluation of apparatus to carry out the proposed ATW (accelerator transmutation of waste) scheme for burning surplus plutonium and transmuting long-lived fission products into stable or short-lived isotopes by irradiation with accelerator-produced neutrons; design of the dual-axis radiographic hydrodynamic test (DARHT) facility for x-ray imaging of explosions; and evaluation of the radiation hazard posed by the plutonium-preparation line at the Laboratory's plutonium-processing facility.

Challenges of MCNP Development

The maturity of the Monte Carlo method and the high satisfaction level of MCNP users might imply that a large investment in future development is unnecessary. After having gone only a few years ago through the formidable process of certifying that version 4 satisfied rigid quality-assurance standards, why should we now be willing to go through the same process for version 4A? Or why should a user already happy with MCNP version 4 bother to obtain and implement version 4A? And finally why should the sponsors of MCNP continue to spend money on a code that is already "good enough"? The answers to those questions lie in the challenges presented by the shifting computer hardware and software environments, by the many users of MCNP, and by the demand for a code of increasing versatility and sophistication.

The MCNP of ten years ago would not run on today's computers because the architectures have changed. MCNP must constantly be adapted to new architectures and other new aspects of the computer world. But care must taken not to squander limited development funds on inappropriate adaptations. For example, we have deliberately avoided adapting MCNP to massively parallel computers that do not use UNIX, standard FORTRAN 77, and MIMD (multiple-instruction, multiple-data) architectures. If massively parallel machines are to become the computers of the future, they will have to accommodate to UNIX and FORTRAN 77, most likely in a MIMD architecture. For similar reasons we have also avoided adapting to systems with immature software and compilers. However, MCNP is usually among the first production physics codes to become available on any commercially viable, state-of-the-art computer architecture. For example, the latest version of MCNP (version 4A) takes full advantage of parallel processing on a cluster of workstations, an architecture that offers many times the computing performance of a single-processor Cray supercomputer.

Strange as it may seem, adapting MCNP to avant-garde architectures is currently of lesser interest than adapting it to the most primitive of computers, the IBM-PC and its clones, which use an operating system, MS-DOS, not designed for scientific computing. The performance of MCNP on such a machine, equipped with a larger memory and a faster processor (such as the 486 chip) and costing less than $3000, approaches its performance on today's supercomputers. In fact, the performance-to-cost ratio is about 10,000 times better than that available on the supercomputers of just a few years ago!

Updating the graphics package included in MCNP also requires considerable effort because graphics packages are not as standardized as operating systems and languages. A major enhancement to the latest version of MCNP is the addition of color graphics and an X-windows-based graphics package. Additional links have also been made to an auxiliary code, Sabrina, which allows geometries to be sketched (see Figure 2) and publication-quality plots of particle histories to be generated. Figure 5 shows an example of such a particle-history plot.

Of all the physics computer codes that have been developed over the years, MCNP is among the leaders in number of users—over a thousand at about a hundred institutions. The Laboratory is obligated, for scientific, ethical, and regulatory reasons, to make MCNP generally available. But meeting the challenge of providing users with a minimum level of support despite the lack of explicit funding for such support remains a dilemma.

Continued maintenance and development are also required to meet new criteria, which may be regulatory, technical, or even political. Higher standards of accountability, particularly in the area of health, safety, and environment, have generated an avalanche of government requirements concerning the quality-control standards imposed on computer codes, including physics codes. Most of the requirements are long overdue, and most have already been met by MCNP. All require voluminous paperwork and are an increasing part of the MCNP-maintenance effort. New technical criteria come about from advances in science and new applications. For example, as new libraries of data about the interactions of neutrons, photons, and electrons with matter become available, MCNP must be upgraded to accommodate different sampling techniques and data formats. New applications may require different code features, such as periodic boundaries and enhanced tally options. Modifications to MCNP are also required for linkage to related computer codes, such as the Sabrina graphics code.

Our current priorities for MCNP are quality, value, and new features.

"Quality" encompasses not only rigorous quality control but also demonstrations of the code's validity and accuracy by comparison with experiment. "Value" encompasses good documentation, timely and controlled releases of code versions, and sufficient standardization and portability that the code can be run on the computers favored by MCNP users. "New features" are the essence of code development but cannot supersede quality and value in importance.

The Future

The future of the Monte Carlo method is secure. As computers become cheaper and faster, it will more and more become the method of choice for solving a wide range of problems. And as the applications of radiation increase, Monte Carlo radiation-transport codes will become more and more widely used. The future of MCNP at the Laboratory is also bright, despite reductions in nuclear-weapons research. In fact, the code is likely to become a more important tool, as physical testing of nuclear weapons is phased out and simulations come to constitute a larger fraction of weapons research. And every enhancement in MCNP benefits not only the Department of Energy's defense programs but also its programs in many other areas: nuclear-criticality safety, environmental restoration, nuclear-waste management, prevention of nuclear proliferation, accelerator production of tritium, accelerator transmutation of nuclear waste, space nuclear power, nuclear safeguards, fusion and fission energy, arms control, intelligence, and on and on. Furthermore, MCNP is one of the Laboratory's most promising candidates for technology transfer to industry. Many companies have already expressed interest in Cooperative Research and Development Agreements on utilizing or jointly developing MCNP for specific applications. Those collaborations and contacts can provide a wide range of industries with a marvelous introduction to the Laboratory. ■

Acknowledgements

Sabrina illustrations provided by Kenneth Van Riper. MCNP is the product of many people over many years, and the following list of present and past contributors to its development is bound to be incomplete. Unless otherwise noted, all are or were affiliated with Los Alamos National Laboratory. Present principal developers: Thomas E. Booth, Judith F. Briesmeister, R. Arthur Forster, John S. Hendricks, H. Grady Hughes, Gregg W. McKinney, Richard E. Prael. Recent significant developers: Thomas N. K. Godfrey, Leland L. Carter (Westinghouse Hanford), David G. Collins, Guy P. Estes, Henry Lichtenstein, Robert C. Little, Shane P. Pederson, Roger R. Roberts, Robert G. Schrandt, Robert E. Seamon, Edward C. Snow, Patrick D. Soran, William M. Taylor, William L. Thompson, Kenneth Van Riper, Laurie S. Waters, James T. West III (IBM), William T. Urban. Graduate students and other contributors: Ronald C. Brockhoff, Robert S. Brown, David A. Cardon, Peter P. Cebull, David E. Hollowell, J. Steven Hansen, Glen T. Nakafuji, Scott P. Palmtag, Everett L. Redmond II, James E. Sisolak, Jennifer L. Uhle, Todd J. Urbatsch, Jason W. VanDenburg, John C. Wagner, Daniel J. Whalen. Distinguished "alumni": Edmond D. Cashwell, C. J. Everett, Enrico Fermi, Nicholas C. Metropolis, John von Neumann, Robert Richtmyer, Stanislaw Ulam.

Further Reading

N. Metropolis and S. Ulam. 1949. The Monte Carlo method. *Journal of the American Statistical Association* 44: 335–341.

W. W. Wood. 1986. Early history of computer simulations in statistical mechanics. In *Molecular-Dynamics Simulation of Statistical-Mechanical Systems. Proceedings of the International School of Physics Enrico Fermi*, course 97, edited by G. Ciccotti and W. G. Hoover. North-Holland Physics Publishing.

Roger Eckhardt. 1987. Stan Ulam, John von Neumann, and the Monte Carlo method. *Los Alamos Science* 15: 131–137. Los Alamos National Laboratory.

N. Metropolis. 1987. The beginning of the Monte Carlo method. *Los Alamos Science* 15: 125–130. Los Alamos National Laboratory.

L. L. Carter and E. D. Cashwell. 1975. *Particle-Transport Simulation with the Monte Carlo Method.* U.S. Energy Research and Development Administration. Available as TID-26607 from National Technical Information Service, Springfield, Virginia 22161.

K. Binder, editor. 1984. *Applications of the Monte Carlo Method in Statistical Physics*. Topics in Current Physics, volume 36. Springer-Verlag.

R. A. Forster, R. C. Little, J. F. Briesmeister, and J. S. Hendricks. 1990. MCNP capabilities for nuclear well logging calculations. *IEEE Transactions on Nuclear Science* 37: 1378–1385.

Ivan Lux and Laszlo Koblinger. 1991. *Monte Carlo Transport Methods: Neutron and Photon Calculations*. CRC Press.

J. F. Briesmeister, editor. 1993. MCNP: A general Monte Carlo N-particle transport code, version 4A. Los Alamos National Laboratory report LA-12625.

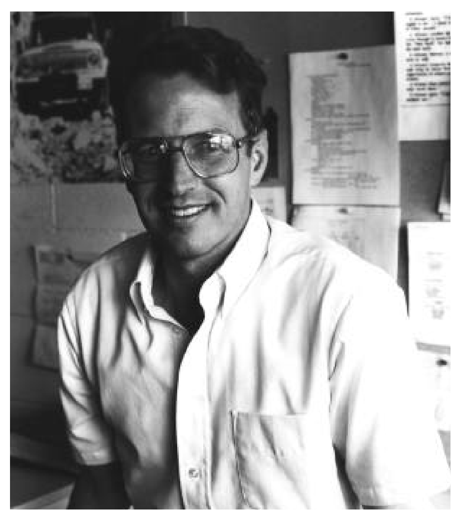

John S. Hendricks received his B.S. and M.S. in nuclear engineering from the University of California, Los Angeles, in 1972 and his Ph.D. in nuclear engineering from the Massachusetts Institute of Technology in 1975. He joined the Laboratory in 1975 as a staff member in the Radiation Transport Group and has remained with that group ever since, except for a two-year stint on the staff of the Laboratory's director as the Congressional liaison. He is currently team leader for Monte Carlo in the Radiation Transport Group.