

Soft-error Monte Carlo modeling program, SEMM

by P. C. Murley
G. R. Srinivasan

The application of a computer program, SEMM (Soft-Error Monte Carlo Modeling), is described. SEMM calculates the soft-error rate (SER) of semiconductor chips due to ionizing radiation. Used primarily to determine whether chip designs meet SER specifications, the program requires detailed layout and process information and circuit Q_{crit} values.

Introduction

SEMM is a computer program that calculates the soft-error rate (SER) of semiconductor chips from information about their design and the ionizing radiation to which they may be exposed; only a brief description of the program has previously been presented [1, 2]. The sources of ionizing radiation are radioactive trace contaminants in the chip which emit alpha-particles or cosmic rays which collide with nuclei within the chip and produce alpha- and other particles capable of ionizing silicon atoms. SEMM is used primarily to determine whether chip designs meet SER specifications and to enable chip designers to make changes in order to meet specifications. It is also used to determine specifications on the level of radioactive contaminants allowable in the materials from which the chips are built. Since SEMM uses information about the sensitivity of circuits on the chip to the charge collected from radiation-generated electron-hole pairs, circuit

parameter specifications related to soft-error sensitivity can be determined iteratively.

Early methods of SER calculation used explicitly or implicitly the concept of the *critical volume* of a device (e.g., [3]). If a given amount of charge (Q_{crit}) collected at a device terminal is known to be just enough to cause a circuit upset, and if that amount of charge can be produced by a given length of ionizing particle track, the volume which could just contain that length of track is a critical volume. For a particular semiconductor device structure, the frequency of regions adjacent to collecting junctions whose volumes exceeded the critical volume was typically empirically estimated. Once this frequency was determined, estimates of alpha-particle fluxes or cosmic-ray-induced secondary-particle fluxes were used to calculate the respective SER. This method was inherently subjective and made it difficult to compare different devices and circuits consistently.

The calculation of charge collection in SEMM begins with the definition and description of an electrically sensitive region of the chip. This requires constructing a three-dimensional rectangular mesh and assigning to each volume element physical properties representative of the corresponding portion of the chip. A Monte Carlo method is then used to simulate a number of radiation events consistent with some *a priori* estimate of the failure probability. In each event primary particles are selected

©Copyright 1996 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/96/\$5.00 © 1996 IBM

according to the specified type of radiation source, with suitably random energies and directions. Each ionizing particle has a randomly selected path through the 3D array, and generates excess electron-hole pairs in the array in accordance with built-in range/energy tables. In contrast to the practice of defining a critical volume and estimating the frequency of its occurrence in the chip, SEMM in effect averages the charge collected from a statistical sample of particles passing through the chip, counting only those instances where charge collected in a region exceeds the critical charge for that region. This makes possible the treatment of complicated chips for which the determination of a critical volume would be difficult.

As in many statistical experiments, the number of simulated random events required in a SEMM run for statistically significant predictions of SER depends inversely on the actual failure rate. If the failure rate expected from the simulation is 10^{-4} , for example, of the order of 10^5 simulated events would be appropriate to achieve statistical significance. The sample sizes typically needed in SER work preclude the direct use of a semiconductor device simulation program. Instead, a simplified charge representation is used in which excess carriers along an ionizing particle track are treated as a perturbation of an existing steady state. Their diffusion is simulated by using a random walk model, where the diffusivity is doping-dependent and the effect of the built-in field associated with doping gradients is included. Minority carriers are collected when they cross a junction or enter a depletion layer. The funneling effect (see below) is represented by rapid collection of electron-hole pairs from tracks which touch a junction region and which would be within the funneling region (separately defined in a device simulation program). Detailed information about the size of the funneling region as a function of doping, junction bias, and the angle between the particle track and the junction depletion edge is available in the experimental work of Hsieh et al. [4, 5].

This charge representation was originally formulated by Sai-Halasz et al. [6] and successfully applied to calculate the SER due to alpha-particle emitters within a chip [7]. Calculations were programmed for the topology of one specific device at a time, such as a half cell in a memory array. This required each circuit designer to extensively modify a computer program and to become deeply involved in the details of SER calculations. The development of SEMM relieved circuit designers of this burden, since within the basic layout unit, a repetitive cell, there may be a number of devices of differing rectangular solid shapes. Various subregions within a device are defined using rectangular solids or triangular prisms. All such regions are specified by the coordinates of their vertices.

SEMM has the additional capability of calculating the SER caused by cosmic radiation. This component of the SER has become more important as increased attention to raw-material purity has reduced the SER due to alpha-particle emission from radioactive contaminants. To simulate the SER caused by cosmic rays colliding with atomic nuclei in the chip, SEMM requires statistical information about the ionizing particles created in such collisions. This information is obtained from a companion program, NUSPA (Nuclear Spallation Reaction Model/Codes) [8, 9], which simulates a large number of such collisions for cosmic rays of various discrete energies between 20 MeV and 900 MeV. This range includes all cosmic rays of interest; secondary particles from sub-20-MeV cosmic rays do not induce enough carriers to be a significant cause of upsets, and there are so few cosmic ray particles above 900 MeV that they do not contribute significantly to SER.

After determining the distribution of the excess carriers generated by each ionizing particle, SEMM simulates the drift and diffusion of holes and electrons as they move through the chip and records in a history file the number of carriers collected at each junction or contact as a function of time. Later this file is postprocessed to calculate any SER effects due to the collected charge, which can be greatly affected by the way the device terminals are connected to form a circuit. SEMM uses its own waveform-dependent method of postprocessing collected charge to determine its impact on circuit operation; the postprocessing itself may be statistical, facilitating the treatment of process tolerance effects upon SER. Since Monte Carlo simulation of the electron-hole collection process uses far more computer resources than the postprocessing, this is an efficient implementation.

The SEMM procedure

The complete SEMM procedure for predicting the SER for a particular chip region requires a number of distinctly different inputs to several independently executable programs, as depicted in **Figure 1**, where a program is represented by the icon for a computer. A geometrical description of the semiconductor region to be analyzed is used by a preprocessor (*pre* in the figure) to generate the topological input file for the main program (*SEMM* in the figure). At present, SEMM can treat up to 24 distinct 3D devices in a (repetitive) cell. Another file supplies SEMM with vertical device profiles (impurity concentration vs. depth) and miscellaneous control parameters. A separate execution of SEMM is needed for each on-chip source and each cosmic ray energy. Alpha-particle emission from an on-chip source is simulated by SEMM, using a geometrical definition of the source (such as a metallized layer or flip-chip solder ball pattern); there may be several of these




Figure 1

Flow diagram for soft-error Monte Carlo modeling program, SEMM. From [2], reprinted with permission; ©1994 IEEE.

sources. In the case of a cosmic source, for each primary nucleon energy a file tabulating the results of many thousands of random nuclear collisions is created by the NUSPA program if the appropriate file does not already exist; files currently exist for enough discrete energies to cover the energy range of interest. (The NUSPA icon in the figure is shown with dashed lines, indicating that NUSPA need only be run occasionally.) For each primary collision event, the energies and directions of resulting secondary particles are input from NUSPA's cosmic history file; these particles include a recoiling nuclear fragment and may include one or many alpha-particles.

Accordingly, SEMM is typically executed with many different simulated radiation sources; each simulation may require 10000 events (or more, if the SER is low), with CPU time measured in hours. For each event resulting in significant charge collection, an entry is made in the

collection history file; this entry gives transient details of the charge collection at all affected locations on the chip. Postprocessing of a saved charge-collection history file, which takes minutes, may be carried out at any later time. (The postprocessor is called *POST* in Figure 1.) The chip designer uses wiring rules which define how charge collected at the terminals of the devices along certain paths in a circuit combines (positively or negatively) to form a composite charge. The definitions of critical charge (Q_{crit}) at each of the sensitive terminals must be provided to the postprocessor; these are generally dependent on the shapes of the current pulses applied to those terminals. A Q_{crit} specification may include a (Gaussian) 1σ tolerance percentage, in which case the postprocessing is statistical; such Q_{crit} tolerances are often implied by statistical variations of chip process parameters or power supplies. Definitions of Q_{crit} are derived from (iterative) circuit




Figure 2

Example of digitized arrays for a cell (basic repetitive unit).

analysis program procedures (*Circuit* in Figure 1). For design purposes, one may wish to change the wiring rules or Q_{crit} definitions and repeat the postprocessing. By default, the result of postprocessing a charge-collection history file is the probability of *at least one* failure, given that an alpha-particle or cosmic nucleon has interacted with the defined chip region; however, one may ask for the probability of at least two (or more) failures in the same event.

In the case of an on-chip source, this probability is converted into an SER by multiplying it by the alpha-particle flux (number per unit area per unit time) characteristic of the source material. For several sources the on-chip SER is the sum of the separate rates, since the sources are statistically independent. Calculation of cosmic SER requires one additional step, using a simple program called SERCALC (not shown in Figure 1). SERCALC has a built-in representation of terrestrial cosmic flux as a function of energy [10]. It computes terrestrial cosmic SER by integrating (over the discrete set of nucleon energies) the product of the cosmic flux and the cosmic failure probability from the postprocessor. The total SER is just the sum of on-chip and cosmic SERs.

• Chip geometry

The particular region on an actual chip that is to be analyzed is henceforth referred to simply as the *chip*; the term *cell* denotes the basic repetitive unit of a chip layout, not necessarily a *memory* cell. The chip is assumed to be covered with identical cells, which are repeated by translation. A cell contains one or more sensitive devices, usually surrounded by passive areas. A radiation source sees the chip boundaries—for example, an alpha-particle emitted from a solder ball at a shallow angle might miss the chip—but once a particle penetrates the chip it is treated by SEMM as if there were no chip boundaries. This is a good approximation for layouts of many cells in each lateral direction. The chip is treated as a planar layer of active semiconductor material on a semiconducting substrate. On top of the chip there are layers of metals, separated by planar passive layers of oxide; above them may be solder balls and a ceramic header.

The top (metallized) surface of the chip is assumed to lie in the y - z coordinate plane, as shown in Figure 2. The x -axis points downward into the chip. A cell, whose layout is described by a planar rectangular grid, contains one or more active device areas (up to 24), and may contain passive areas. The latter collect no charge, but allow diffusion and recombination of carriers. Each rectangular element in the grid is given a one-character device code (*A* or *B* in the example) or passive-area code (not shown). Below each device area is a rectangular solid which has the same lateral grid and is divided by a vertical grid into rectangular blocks. Each block is given a one-character region code (0, 1, 2, ..., *F* in the example) that

identifies the function and physical properties of its three-dimensional region; currently there may be up to 36 region codes. To run SEMM one must obtain the top-view coordinates of all of the corners of distinct regions within the cell, as well as the depths of diffused, implanted, and insulating regions in the chip, in order to complete the 3D spatial description. One must also have impurity concentration profiles perpendicular to the surface, which influence the diffusion and recombination of minority carriers, and the thickness and composition of each package layer and the dimensions of solder balls (if pertinent).

• Region of analysis

For computational efficiency, during each event SEMM analyzes only a limited region of the chip surrounding an alpha-particle or cosmic ray strike; an example is shown in Figure 3. A cosmic ray is assumed to produce multiple secondary-particle tracks whose horizontal projections lie within an area marked off by three cells in each coordinate direction. That area, bordered by an extra strip of cells, constitutes the region which is analyzed and for which charge-collection results are saved. The border allows some carriers to diffuse beyond the area originally affected.

• On-chip radiation sources

SEMM models the emission of alpha-particles from sources made up of rectangular bars, cylinders, or spheres, which may be solid or hollow and are specified by their individual positions, thicknesses, and materials. A metallized interconnection layer can be represented as a pattern of solid bars or (for economy of input) as a single solid bar thin enough to have the same volume as the patterned layer. Interconnection layers are separated by oxide layers. Solder balls are usually treated as solid cylinders, and ceramic headers as solid bars. Each interconnection layer or other source normally requires a separate SEMM execution; any layer between a source and the silicon chip acts as a passive absorber of alpha-particle energy.

Emission of an alpha-particle occurs at one or more discrete energies, depending on the emitting material. Polonium in the lead-tin solder ball emits only at 5.3 MeV; the bulk emission rate for low-contamination lead is of the order of one alpha-particle per cm^2 per hour. Aluminum in the interconnect layers is assumed to have the same discrete energy spectrum as thorium, ranging from 4 MeV to almost 9 MeV; emission occurs at a rate of about 0.01 per cm^2 per hour, so low as to be difficult to measure. The emission spectrum of ceramic material is assumed to be equivalent to equal parts of thorium and uranium, covering the same energy range as thorium alone; its emission rate has been measured at 30 alpha-particles per cm^2 per hour.




Figure 3

Example of region around particle strike which is to be analyzed using SEMM (darker dashed lines).




Figure 4

Range of an alpha-particle in various materials of a chip [12, 13].

• Cosmic radiation effects

Cosmic particles in the form of neutrons or protons can collide randomly with silicon nuclei in the chip and

fragment some of them, producing alpha-particles and other secondary particles, including the recoiling nucleus [8, 9, 11]. These can travel in all directions with energies which can be quite high (though of course less than the incoming nucleon energy); alpha-particle tracks so produced can sometimes extend a hundred microns through the silicon. SEMM follows all of the secondary-particle tracks in such a point burst.

• Charge creation and particle energy losses

A high-energy particle penetrating a silicon chip moves approximately in a straight line, losing energy as it moves along. Shown in Figure 4 are curves of alpha-particle range vs. energy currently used in SEMM and obtained from TRIM, a PC program which simulates ion implantation in amorphous materials [12, 13]. The track of an ionizing particle may extend a fraction of a micron to many microns through the chip volume of interest, generating in its wake electron-hole pairs at the rate of one pair per 3.6-eV loss of energy. In SEMM the ionization is considered to be almost instantaneous and the track width less than the smallest grid spacing. A typical track might represent a million pairs of holes and electrons. SEMM simulates their movement through the chip by drift or diffusion until they recombine or are collected at device junctions or contacts. The CPU time is proportional to the number of pairs, so each pair *used* in the simulation often is chosen to represent 1000 or more actual electron-hole pairs without noticeably affecting the results. Such a "super pair" is termed a *charge packet*.

• Charge collection and diffusion

The charge generated by a primary-particle strike is accounted for by collection or recombination before the next (independent) strike occurs. There is no accumulation of charge from one strike to the next, since even for high-flux sources the time interval between strikes is long in comparison with the switching speed of today's circuits. In a neutral region the electrons and holes generated by a particle diffuse slowly away from their origin. This movement does not produce a net electric current because the electrons and the holes tend to move together, almost maintaining charge neutrality. However, when the electron-hole pairs drift into the depletion region of a p-n junction, a separation of charges occurs; the electrons move to the n-side of the junction, and the holes to the p-side. The separation occurs almost instantly for pairs initially in or "funneled" (see below) to a depletion region.

This separation of charges is represented in an equivalent circuit by a current source connected between the external terminals of the n- and p-regions. The positive direction of the current is from the n- to the p-terminal. Such induced leakage currents can cause unintended changes in the state of a cell that serves as a storage

element. For each radiation event simulated in the Monte Carlo portion of SEMM, a number of entries are made in a charge-collection history file to record the number of charge packets in each interval of simulation time at each collecting junction or terminal in the active chip area analyzed.

In the SEMM approximation of the diffusion of the minority carriers in a charge packet, a representative carrier from the packet takes a sequence of random steps. A step moves the carrier a randomly positive or negative distance S in each coordinate direction; S normally has a value comparable to the vertical grid spacing near the silicon surface. If D is the local carrier ambipolar diffusivity, the simulated time for the step is $S^2 \div (6D)$ [14]. A step adds its time to the age of the minority carrier to determine whether the carrier should recombine, terminating the walk. The carrier's new x - y - z coordinates are used to determine whether it has crossed a p-n junction; if so, the entire packet is *collected*, and the walk is terminated. The walk is also terminated if the carrier has diffused so far from the region of interest that its return would be highly unlikely. The diffusion steps are modified by a factor calculated from built-in electric fields perpendicular to the surface; this enhances charge transport, sometimes aiding charge collection and sometimes reducing it.

• Funneling

If an ionizing particle intersects a reverse-biased asymmetrically doped p-n junction, the junction's depletion layer is temporarily distorted into the low-doped side by the ionized track. The electrostatic equipotentials, initially approximately parallel to the junction, become stretched out in the vicinity of the track, as shown schematically in **Figure 5**. The term "funneling" was defined in 1981 [4, 5] to describe this effect because of the shape of these equipotentials (calculated in transient device analysis programs). The modification of the equipotentials is strongly affected by background doping and moderately affected by the angle of intersection between the track and the plane of the junction [5, 15]. If funneling occurs, the electric field in that part of the depletion layer is reduced in strength, but acts upon the ionized carriers over a greater distance than the initial depletion layer would. As the drift field pulls minority carriers to the metallurgical junction and consumes them, the equipotentials return to their equilibrium positions. The disturbance builds up and subsides in fractions of a nanosecond, but despite its brevity often represents the dominant mode of charge collection (as opposed to diffusion). The drift of minority carriers to a junction causes a pulse of current into an associated terminal; the pulse typically has a short rise time and decays exponentially from its peak; this is shown in **Figure 6** for an alpha-particle penetrating a junction on a uniformly low-doped (10^{15} cm^{-3}) substrate. Vertical




Figure 5

"Funneling": Distortion of p-n junction equipotentials due to an alpha-particle emitted from a chip package and penetrating through a p-n junction into a low-doped substrate [5].

funneling depth is defined for such a planar p-n junction as the difference between the abrupt-junction depletion layer width for the amount of charge actually collected by drift after a particle strike and the abrupt-junction depletion layer width at an equilibrium bias.

Evaluation of failure rates

• Critical charge considerations

Failure rates are computed from the collected charge by postprocessing the Monte Carlo output. Much of the recent development of SEMM has been in the functions of the postprocessor, which reads a charge-collection history file event by event and counts the failures. For each event, it computes an *effective* charge by summing the charges collected in a cell, using various weights for the various terminals. The weights, which may be positive, zero, or negative, are supplied by the chip designer; they determine the circuit paths, directions, and relative sensitivities, and are used in conjunction with the critical charge values for those terminals. Each value of Q_{crit} is commonly calculated with a circuit simulator by applying a current pulse to a terminal and varying its magnitude until its charge is just enough to cause an upset [16].

For early chip designs, it was commonly accepted that the sensitivity to radiation of a typical memory cell or a latch could be expressed with a single value of Q_{crit} , independent of the *waveform* of the radiation-induced pulse of current into a terminal, but possibly using *a priori* weights as just described. For waveform-independent Q_{crit} values, the postprocessor reads an event record from the collection history file and computes the maximum effective




Figure 6

Waveform of a calculated current pulse influenced by funneling of charge after alpha-particle strike into a substrate doped with boron at 10^{15} cm^{-3} .

charge over all affected cells on the chip. After repeating this for all event records, the postprocessor tabulates internally the maximum effective charge as a histogram. Its upper partial sum (from highest to lowest charge) is then output as a table of the probability that a particle strike on the chip will produce *at least* a certain effective charge. The designer can then look up in the table the probability corresponding to a particular Q_{crit} .

For many contemporary applications, however, the charge required for an upset is strongly dependent upon the waveforms of pulses applied to a sensitive terminal. The pulses in general have a rapid rise followed by an exponential decay, and are characterized by their decay time-constants. A circuit which recovers quickly from a disturbance may have a lower Q_{crit} for a spike of current than for a slower pulse. The critical charge may also vary from one terminal to another. Therefore, to apply SEMM to a typical chip design one must have available the Q_{crit} as a function of its exponential decay time-constant at each of the sensitive terminals in the circuit. Usually, the Q_{crit} is a roughly linear increasing function of the time-constant.

One must also be aware of the effects of chip process parameters and power supply variations on these Q_{crit} values. All of these Q_{crit} effects are discussed at length in the companion paper by Freeman [16].

If Q_{crit} values are waveform-dependent, one also inputs for each sensitive terminal a table of Q_{crit} vs. time-constant. The postprocessor monitors the current arriving at each terminal, fits its decay with a time-constant, looks up the corresponding Q_{crit} in the table, and normalizes the *collected* charge by dividing it by that Q_{crit} . The normalized charges are added as described above, but now a sum of normalized charges greater than 1 for any event inherently denotes a *failure*. The postprocessor's direct output is the probability that a particle strike on the chip will produce at least one failure. The result of processing a charge-collection history file is represented in Figure 7, which indicates the probability of being at least a certain amount over (or under) the critical value of charge. Primary interest is in the probability defined by $(Q - Q_{\text{crit}}) = 0$, indicated by the dotted line in Figure 7 (in this example the probability value is 3×10^{-4}).

For postprocessing, the Q_{crit} at any terminal is allowed to vary statistically to reflect variations in wafer processing or power supplies. (Lithography variations would require complete SEMM runs.) In this case, the Q_{crit} for any particular time-constant is assumed to have a Gaussian distribution about the nominal value. For each SEMM event, the postprocessor nominal calculation proceeds as described above. Then the calculation is repeated with random Gaussian samples of the Q_{crit} values for the same SEMM event. After processing all the events, the program prints the mean and standard deviation of the distribution of failure probabilities, which are saved for further postprocessing.

• Cosmic SER calculations

Model verification via particle accelerator SER requires the value of the beam flux, F , measured consistently with the SER experimental conditions, and the SEMM-simulated failure probability, P . The calculated SER in this case is just $P \times F$.

For terrestrial cosmic-induced SER calculations, fluxes vary with altitude and geographic location [10]; expressions for these variations are part of the SEMM package, as is a formula for the energy-dependence of the sea-level flux at New York City [10]. A program, SERCALC, is used to compute terrestrial cosmic SER from the failure probabilities previously calculated by the postprocessor for a set of individual cosmic nucleon energies. Assuming that between tabulated energies the failure probability P and the flux F each behave as some power of the energy, SERCALC integrates $(P \times F)$ over the energy domain for the nominal Q_{crit} results and also for any statistically varied Q_{crit} results. The principal input data to SERCALC are

the set of failure statistics saved by the postprocessor. Additional input data items required are chip dimensions, the number of bits in the chip (for memory arrays), and the altitude, latitude, and longitude of the location at which to evaluate the terrestrial SER.

• On-chip SER calculations

The output from the postprocessor run for a particular radiation source is the probability P of failure, given that the chip area has been struck. To calculate chip SER (fails per unit time), one multiplies P by the number of particle strikes per unit time. The number of particle strikes per unit time is the chip area multiplied by the flux impinging on that area. Usually, the flux from a bulk sample of a package material is measured with a large-area zinc sulphide detector and reported as so many particle strikes per cm^2 per hour or per thousand hours. The detector has a typical threshold energy (minimum particle energy for detection) of 1 or 2 MeV. An uncertainty exists for the value of the flux from package materials, either because their radioactive contamination level varies or because the emission rate is so low that it is masked by background noise. Therefore, up-to-date and suitably authenticated flux values must be obtained. Modeling may be required in order to convert the bulk material flux on the detector into a particle emission rate for the source as actually fabricated (SEMM itself may be used for this modeling).

Summary

SEMM is a computer program at the heart of a state-of-the-art soft-error modeling procedure whose goal is to predict the SER of semiconductor devices in circuits due to on-chip and cosmic ionizing radiation. It accepts a detailed 3D chip description and models the transient collection of induced holes or electrons at semiconductor junctions and contacts. Then, from user-defined waveform-dependent Q_{crit} values, it calculates the probabilities of soft failures due to the radiation sources chosen and evaluates the chip SER. The SEMM program has been associated with the IBM chip design process for more than a decade.

Acknowledgments

The predecessor of SEMM was coded by B. V. Gokhale, who also created helpful documentation of the methodology and contributed to the preceding description of charge collection. The code which accepts an on-chip source description and models random emission from the source toward a chip was written by F. J. Astudillo. Many IBM chip designers have contributed to the development of SEMM through their suggestions and comments resulting from their use of the procedure.




Figure 7

Example of calculated failure probability curve; failures occur when collected charge Q exceeds critical charge Q_{crit} (hatched region).

References

1. G. R. Srinivasan, P. C. Murley, and H. H. K. Tang, "Accurate, Predictive Modeling of Soft Error Rate due to Cosmic Rays and Chip Alpha Radiation," *Proceedings of the 32nd Annual IEEE International Reliability Physics Symposium*, San Jose, CA, April 12, 1994, pp. 12-16.
2. G. R. Srinivasan, H. H. K. Tang, and P. C. Murley, "Parameter-Free, Predictive Modeling of Single Event Upsets due to Protons, Neutrons, and Pions in Terrestrial Cosmic Rays," *IEEE Trans. Nucl. Sci.* **41**, No. 6, 2063-2070 (December 1994).
3. S. Kirkpatrick, "Modeling Diffusion and Collection of Charge from Ionizing Radiation in Silicon Devices," *IEEE Trans. Electron Devices* **ED-26**, No. 11, 1742-1753 (November 1979).
4. C. M. Hsieh, P. C. Murley, and R. R. O'Brien, "Collection of Charge from Alpha-Particle Tracks in Silicon Devices," *IEEE Trans. Electron Devices* **ED-30**, 686-693 (June 1983).
5. C. M. Hsieh, P. C. Murley, and R. R. O'Brien, "A Field-Funneling Effect on the Collection of Alpha-Particle-Generated Carriers in Silicon Devices," *IEEE Electron Device Lett.* **EDL-2**, 103-105 (April 1981).

6. G. A. Sai-Halasz, M. R. Wordeman, and R. H. Dennard, "Alpha Particle-Induced Soft Error Rate in VLSI Circuits," *IEEE Trans. Electron Devices* **29**, 725 (1982).
7. G. A. Sai-Halasz and M. R. Wordeman, "Monte Carlo Modeling of the Transport of Ionizing Radiation-Created Carriers in Integrated Circuits," *IEEE Electron Device Lett.* **EDL-10**, 211-213 (1980).
8. H. H. K. Tang, G. R. Srinivasan, and N. Azziz, "Cascade Statistical Model for Nucleon-Induced Reactions on Light Nuclei in the Energy Range 50 MeV-1 GeV," *Phys. Rev. C* **42**, No. 4, 1598-1622 (October 1990).
9. H. H. K. Tang, "Nuclear Physics of Cosmic Ray Interaction with Semiconductor Materials: Particle-Induced Soft Errors from a Physicist's Perspective," *IBM J. Res. Develop.* **40**, No. 1, 91-108 (1996, this issue).
10. J. F. Ziegler, "Terrestrial Cosmic Rays," *IBM J. Res. Develop.* **40**, No. 1, 19-39 (1996, this issue).
11. H. W. Bertini, *Report No. 3383*, Oak Ridge National Laboratory, Oak Ridge, TN, 1963.
12. J. P. Biersack and L. G. Haggmark, "A Monte Carlo Computer Program for the Transport of Energetic Ions in Amorphous Targets," *Nucl. Instr. & Meth.* **174**, 257-269 (1980).
13. J. F. Ziegler, J. P. Biersack, and U. Littmark, *The Stopping and Range of Ions in Solids*, Pergamon Press, New York, 1985, p. 41.
14. P. G. Shewmon, *Diffusion in Solids*, McGraw-Hill Book Co., Inc., New York, 1963, p. 52.
15. C. M. Hsieh, P. C. Murley, and R. R. O'Brien, "Dynamics of Charge Collection from Alpha-Particle Tracks in Integrated Circuits," *Proceedings of the 19th Annual IEEE International Reliability Physics Symposium*, Orlando, FL, April 7, 1981, pp. 686-693.
16. L. B. Freeman, "Critical Charge Calculations for a Bipolar Array Cell," *IBM J. Res. Develop.* **40**, No. 1, 119-129 (1996, this issue).

Received July 15, 1994; accepted for publication February 18, 1995

Philip C. Murley *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (MURLEY at FSHVMFK1, murley@fshvmfk1.vnet.ibm.com).* Mr. Murley is a senior engineer in the Technology Modeling Department. He received the B.A. degree in mathematics from Northwestern University, Evanston, Illinois, and the M.S. degree in applied mathematics from the University of Illinois, Urbana, where he worked on cylindrical shell-buckling problems, using the Illiac-II computer. In 1957 he joined IBM at the product development laboratory in Poughkeepsie, where he helped develop magnetic character recognition logic and also software for large-scale linearized network analysis. Since 1961, Mr. Murley has concentrated on mathematical and numerical analysis of semiconductor devices and processing. His current work includes semiconductor process modeling and soft-error simulation.

G. R. Srinivasan *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (srinivas@fshvmfk1.vnet.ibm.com).* Dr. Srinivasan (Ph.D. physical metallurgy, University of Illinois, Urbana) worked at Cornell University as a member of the research faculty, and at the Catholic University of America, where he was an associate professor of materials science. He has conducted and directed research in such diverse fields as phase transitions, epitaxy, device physics, materials theory, and electron microscopy. He joined IBM in 1974, and managed a theoretical modeling department where he built a group to do research in ion-stopping theories, ion channeling, diffusion, atomistic simulation of defects in semiconductors, soft-error modeling, and advanced device theory. In 1985, he was put in charge of a theoretical effort for modeling computer-chip soft errors due to high-energy particles. Dr. Srinivasan's contributions have been published in more than a hundred scientific papers, and he holds many patents. His formulation of the event-by-event simulation for cosmic ray effects forms the basis of the present soft-error model. He has received five IBM Invention Achievement Awards and Publication Awards, and an IBM Outstanding Technical Achievement Award. Dr. Srinivasan is active in the Electrochemical Society as a member of the Executive Committee and the Technical Planning Committee. He organizes an international symposium on process physics and modeling in semiconductor technology every three years, and is the principal editor of its proceedings volumes. He has also served as the divisional editor for the *Journal of the Electrochemical Society*.