

Soft Errors in Advanced Computer Systems

Robert Baumann
Texas Instruments

Editor's note:

This article comprehensively analyzes soft-error sensitivity in modern systems and shows it to be application dependent. The discussion covers ground-level radiation mechanisms that have the most serious impact on circuit operation, along with the effect of technology scaling on soft-error rates in memory and logic.

—Dimitris Gizopoulos, University of Piraeus

■ **AS THE DIMENSIONS** and operating voltages of computer electronics shrink to satisfy consumers' insatiable demand for higher density, greater functionality, and lower power consumption, sensitivity to radiation increases dramatically. In terrestrial applications, the predominant radiation issue is the soft error, whereby a single radiation event causes a data bit stored in a device to be corrupted until new data is written to that device.

In the past two decades, researchers have discovered three key radiation mechanisms that cause soft errors in semiconductor devices at terrestrial altitudes. In the late 1970s, alpha particles emitted by trace uranium and thorium impurities in packaging materials proved to be the dominant cause of soft errors in DRAM devices.¹ During the same era, researchers showed that high-energy (more than 1 million electron volts, or MeV) neutrons from cosmic radiation could induce soft errors in semiconductor devices² via the secondary ions produced by the neutron reactions with silicon nuclei. In the mid-1990s, high-energy cosmic radiation proved to be the dominant source of soft errors in DRAM devices.³ Finally, researchers identified a third soft-error mechanism, induced by low-energy cosmic neutron interactions with the isotope boron-10 (¹⁰B) in IC materials, specifically in borophosphosilicate glass (BPSG), used widely to form insulator layers in IC manufacturing.⁴ This recently proved to be the dominant source of soft errors in 0.25- and 0.18- μm SRAM fabricated with BPSG.⁵

In advanced processes that produce digital signal

processors and microprocessors and use only the most highly purified chip and packaging materials with extremely low levels of uranium and thorium impurities, and where a dielectric material free of ¹⁰B has replaced BPSG layers, high-energy cosmic neutrons are responsible for the majority of soft errors observed.

In terrestrial environments, three mechanisms generate (either directly or as secondary reaction products) energetic ions responsible for actually inducing soft errors. The magnitude of the disturbance an ion causes depends on its linear energy transfer (LET), typically reported in units of MeVcm^2/mg . In a silicon substrate, every 3.6 eV of energy the ion loses produces one electron-hole pair. A simple conversion lets us plot LET in units of charge loss per distance, as Figure 1 illustrates. LET depends on the mass and energy of the particle and the material in which it is traveling.

Typically, more massive and energetic particles in denser materials have higher LET. For example, the LET of a magnesium ion (one of the ions commonly produced when a high-energy neutron reacts with a silicon nucleus) is significantly higher than that of alpha particles (helium ion) or lithium ions. Charge collection generally occurs within a micron or two of the junction; therefore, the collected charge (Q_{coll}) for these events ranges from 0 to several hundred femtocoulombs, depending on the type of ion, its trajectory, and its energy over the path through or near the junction.

The reverse-biased junction is the most charge-sensitive part of circuits, particularly if the junction is floating or weakly driven (with only a small drive transistor or a high-resistance load sourcing the current required to keep the node in its state). As Figure 2 shows, at the onset of an ionizing radiation event, a cylindrical track of electron-hole pairs with a submicron radius and very high carrier concentration forms in the wake of the energetic

ion's passage (Figure 2a). When the resultant ionization track traverses or comes close to the depletion region, the electric field rapidly collects carriers, creating a current/voltage glitch at that node. A notable feature of the event is the concurrent distortion of the potential into a funnel shape.⁶ This funnel greatly enhances the drift collection's efficiency by extending the high-field depletion region deeper into the substrate (Figure 2b). The funnel's size is a function of substrate doping—the funnel distortion increasing as substrate doping decreases. This collection phase completes within tens of picoseconds, and another phase follows in which diffusion begins to dominate the collection process (Figure 2c). An additional charge collects as electrons diffuse into the depletion region on a longer time scale (hundreds of nanoseconds) until all excess carriers have been collected, recombined, or diffused away from the junction area. Figure 2d shows the corresponding current pulse resulting from these three phases.

In general, the farther away from the junction the event occurs, the smaller the charge that will be collected and the less likely that the event will cause a soft error. In actual circuits, a node is never isolated but is actually part of a complex sea of nodes close to one another; thus, charge sharing among nodes and parasitic bipolar action (the formation of an unintentional bipolar transistor between junctions and wells) can greatly influence the amount of charge collected and the size and location of voltage/current glitches in the circuit.

The magnitude of Q_{coll} depends on a complex combination of factors, including the device's size, biasing

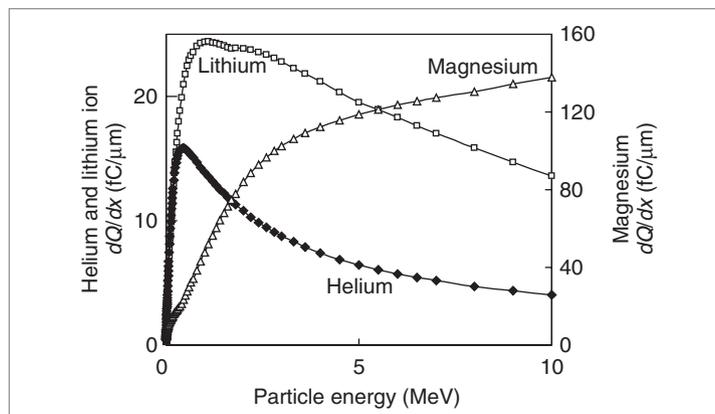


Figure 1. Linear energy transfer converted into charge generation per linear distance for various ions in silicon. On the y-axis, dQ/dx is the differential charge per distance traveled in a medium, measured in femtocoulombs per micron.

of the various circuit nodes, the substrate structure, device doping, the type of ion (as well as its energy and trajectory), where the ion occurs within the device, and the device's state. However, Q_{coll} is only half the story; the device's sensitivity to this excess charge also requires consideration. This sensitivity depends primarily on the node capacitance, the operating voltage, and the strength of feedback transistors; all these factors define the amount of charge, or critical charge (Q_{crit}), required to trigger a change in the data state. The device's response to the charge injection is dynamic and depends on the pulse's magnitude and temporal characteristics; therefore, Q_{crit} is not constant but

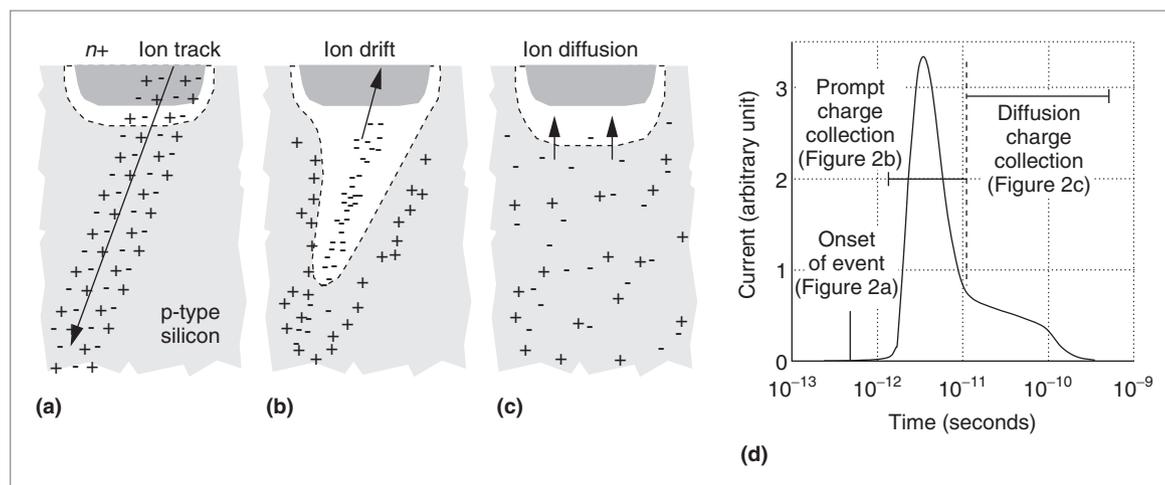


Figure 2. Charge generation and collection in a reverse-biased junction: formation of a cylindrical track of electron-hole pairs (a), funnel shape extending high field depletion region deeper into substrate (b), diffusion beginning to dominate collection process (c), and the resultant current pulse caused by the passage of a high-energy ion (d).

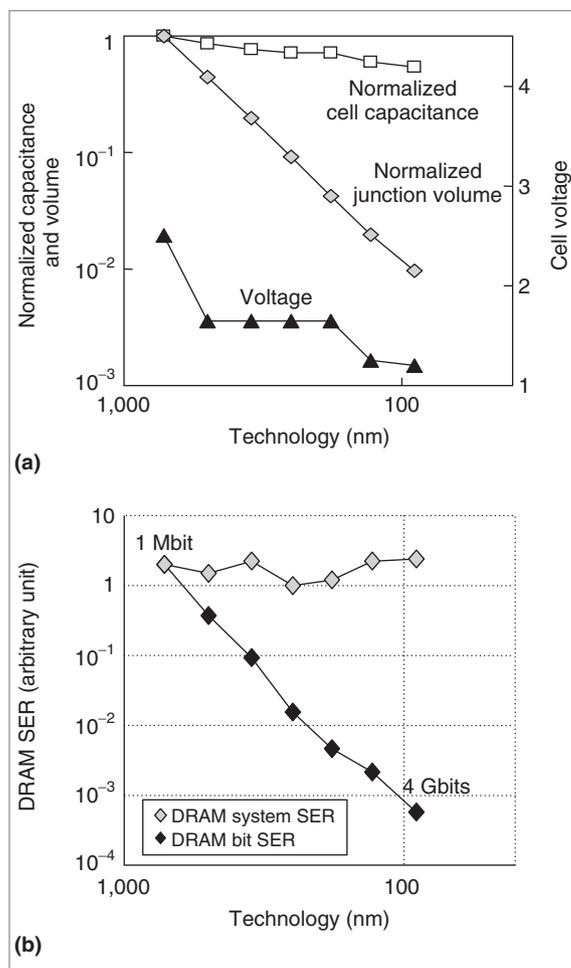


Figure 3. DRAM reliability has remained fairly constant: DRAM scaling parameters, normalized cell capacitance, normalized junction volume, and cell voltage as a function of technology node (a), and DRAM single-bit soft-error rate and system soft-error rate as a function of technology node (b).

depends on the radiation pulse characteristics and the dynamic response of the circuit itself, making the effect extremely difficult to model.⁷

For simple isolated junctions (such as DRAM cells in storage mode), a soft error will be induced when a radiation event occurs close enough to a sensitive node such that Q_{coll} is greater than Q_{crit} . Conversely, if the event produces a Q_{coll} less than Q_{crit} , the circuit will survive the event and no soft error will occur. In SRAM or other logic circuits with active feedback, there is an additional component of critical charge related to the magnitude of the compensating current and the switching time of the device. Increasing the feedback so that more current is

available for compensating the current charge produced by the radiation event increases the Q_{crit} . Similarly, if the time to switch the device is increased, more compensation current can be collected, thus also increasing the Q_{crit} .

The commonly used unit of measure for the soft-error rate (SER) and other hard-reliability mechanisms is the FIT (failure in time). A FIT is equivalent to one failure in 10^9 device hours. Soft errors have become a huge concern in advanced computer chips because, uncorrected, they produce a failure rate exceeding that of all other reliability mechanisms combined. For example, a typical failure rate for a hard-reliability mechanism (gate-oxide breakdown, metal electromigration, and so on) is about 1 to 50 FITs. Half a dozen critical-reliability mechanisms degrade IC performance, but in general the overall failure rate is typically between 5 and 150 FITs. In stark contrast, without mitigation, the SER can easily exceed 50,000 FITs per chip.

Memory SER sensitivity versus technology scaling

Creating the functionality of today's electronic systems and appliances requires integrating several distinct components. At the core of each system is a microprocessor or a digital signal processor with large embedded memories (usually SRAMs) interconnected with extensive peripheral logic. Larger systems also use discrete main memory (usually DRAM). Finally, all systems have some analog or digital I/O components that let the device respond to and interact with the outside world. The SER of these different components varies as the technologies scale.

That soft errors first became a noticeable problem in DRAM is somewhat ironic, because after many generations, a DRAM is currently one of the more robust electronic devices. The DRAM bit SER was high when manufacturers used planar capacitor cells that stored the signal charge in 2D, large-area junctions, because these cells were very efficient at collecting radiation-induced charge. To address pause-refresh and soft-error problems while increasing packing density, DRAM manufacturers developed 3D capacitor designs that significantly increase Q_{crit} while greatly reducing junction collection efficiency by eliminating the large storage junction in silicon. Collection efficiency decreases with the junction's decreasing volume (junction/well doping also plays a role), whereas the cell capacitance remains relatively constant with scaling because it is dominated by the external 3D capacitor cell.

Figure 3a illustrates these DRAM device-scaling trends,

along with DRAM cell voltage scaling. Voltage reduction has reduced Q_{crit} , but with concurrent aggressive junction volume scaling, there's a much more significant reduction in collected charge. The net result to the DRAM SER appears in Figure 3b: The DRAM SER of a single bit is shrinking about $4\times$ to $5\times$ per generation. Although the DRAM bit SER has decreased by more than $1,000\times$ over seven generations, the DRAM system SER has remained essentially unchanged. System requirements have increased memory density (bits per system) almost as fast as the SER reduction that technology scaling provided. Thus, DRAM system reliability has remained roughly constant over many generations. So, contrary to the popular misconception that the DRAM SER is problematic—undoubtedly left over from the days when DRAM designs used planar cells—a DRAM is one of the more robust devices in terms of soft-error immunity.

In contrast, early SRAM was more robust against the SER because of high operating voltages and because an SRAM stores data as an active state of a bistable circuit made up of two cross-coupled inverters, each strongly driving the other to keep the SRAM bit in its programmed state. As with DRAM, the Q_{crit} for the SRAM cell is largely defined by the charge on the node capacitance, but the SRAM cell has a dynamic second term related to the load transistor's drive capability. The load transistor keeps the node voltage at the proper value: The stronger the transistor, the more robust the node, because the node must collect more charge to induce switching. With technology scaling, designers have deliberately minimized the SRAM junction area to reduce capacitance, leakage, and cell area while aggressively scaling down the SRAM operating voltage to minimize power consumption. Figure 4a shows these device-scaling trends.

With each successive SRAM generation, a reduction in operating voltage and node capacitance has cancelled out the reduction in cell collection efficiency caused by shrinking cell depletion volume. Initially, the SRAM single-bit SER was increasing with each successive generation, particularly in products using BPSG, as Figure 4b illustrates. Most recently, as feature sizes have shrunk into the deep-submicron range (less than 250 nm), the SRAM bit SER has reached saturation and might even be decreasing. This saturation is primarily due to the saturation in voltage scaling, reductions in junction collection efficiency, and increased charge sharing caused by short-channel effects with neighboring nodes.

Ultimately, because scaling also implies increased memory density, saturation in the SRAM bit SER doesn't translate to saturation in the SRAM system SER.

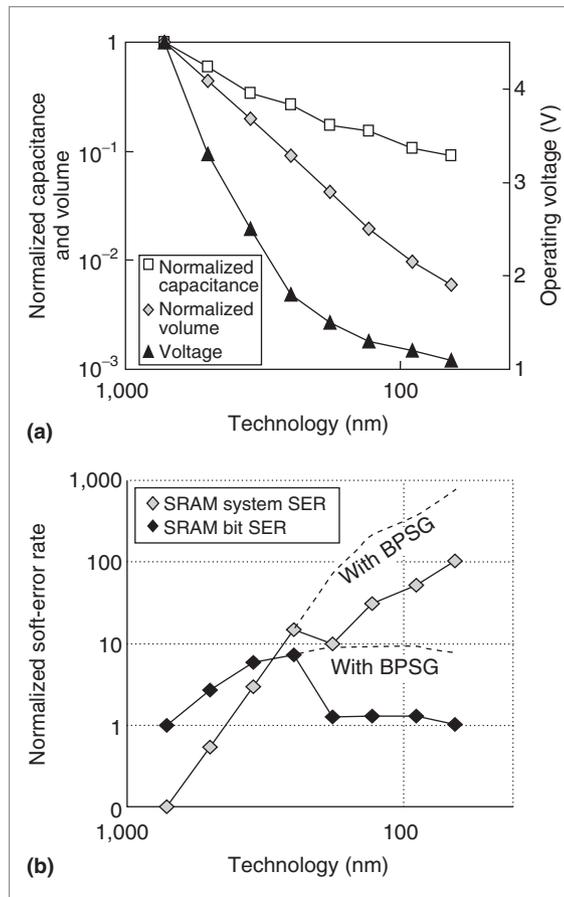


Figure 4. SRAM scaling trends: device parameters, normalized storage node capacitance, normalized junction volume, and voltage as a function of technology node (a), and SRAM single-bit and system soft-error rate as a function of technology node (b). The discontinuous reduction in soft-error rate following the 250-nm node is due to borophosphosilicate glass (BPSG) elimination (dotted lines estimate the soft-error rate with BPSG present in the technology).

Exponential growth in the amount of SRAM in microprocessors and digital signal processors has led the SER to increase with each generation, with no end in sight. This trend is of great concern to chip manufacturers because SRAM constitutes a large part of all advanced ICs today.

Sequential and combinatorial logic sensitivity

A computer's discrete and embedded SRAM and DRAM would be useless without the peripheral logic that

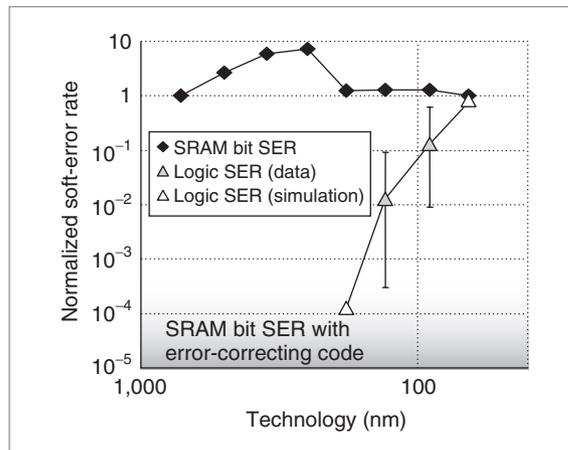


Figure 5. Comparison of the SRAM bit soft-error rate (SER) with the flip-flop/latch SER, obtained from test structure and product characterizations and simulations. The shaded region at the bottom of the plot represents the effective bit failure rate of SRAM with error correction. The large variation in logic SER is due to the dozens of logic types tested, not experimental error.

interconnects them. Although less sensitive than SRAM, logic devices can also experience soft errors.⁸ Sequential logic elements include latches and flip-flops that hold system event signals, buffer data before it goes in or out of the microprocessor, and interface to combinatorial elements that perform logical operations based on multiple inputs. The SER of these devices and its impact on the system are much harder to quantify because their period of vulnerability (when they are actually doing something critical in the system rather than simply waiting) varies widely, depending on the circuit design, frequency of operation, and the actual algorithm executing.

Flip-flops and latches are fundamentally similar to the SRAM cell in that they use cross-coupled inverters to store the data state. However, they tend to be more robust because they are usually designed with more and larger transistors, which can more easily compensate for any spurious charge collected during radiation events. Ultimately, the reliability concern with sequential and combinatorial logic circuits is that, as with SRAM, their SER sensitivity is increasing with scaling, as Figure 5 illustrates. Soft errors in logic are especially worrisome in high-reliability systems whose memory has been protected by error correction, because in such systems the peripheral logic failure rate can be the dominant reliability failure mechanism.

In a combinatorial circuit, in which the output is based on a logical relation to the inputs (with no capability for retention), collection of a sufficient radiation-induced charge will generate a short-lived transient in the output (a single-event transient, or SET).⁹ If this radiation-induced glitch actually propagates to the input of a latch or flip-flop during a latching clock signal, the erroneous input will be latched and stored. In older technologies, the SET could not propagate because it usually couldn't produce a full output swing or was quickly attenuated because of large load capacitances and large propagation delays. In advanced technologies, where the propagation delay is reduced and the clock frequency is high, the SET can more easily traverse many logic gates, and the probability that it is latched increases. SET-induced soft errors are not expected to become an issue until the technology reaches or surpasses the 65-nm node. Once an SET can propagate freely, synchronous—and especially asynchronous (self-clocked)—circuits would be extremely sensitive to such events. In technology nodes beyond 90 nm and at high product operating frequencies, there is increased risk that a large fraction of observed soft failures will be related to latched SET events.

Mitigating soft errors

The most obvious way to eliminate soft errors is to remove the radiation sources that cause them. To mitigate the dominant SER threat that the reaction of low-energy neutrons and ¹⁰B poses, manufacturers have removed BPSG from virtually all advanced technologies. To reduce alpha particle emissions, semiconductor manufacturers use extremely high purity materials and processes, production screening all materials having low background alpha emission measurements. Another way to reduce alpha particles is to design chips on which the materials with the highest alpha emission are physically separated from sensitive circuit components. One last solution frequently used to shield the high alpha emission from packaging materials is to coat the chip with a thick polyimide layer before packaging. Although large SER reductions are possible either by removing the sources of or shielding the ¹⁰B reaction products and alpha particles, a large portion of the high-energy cosmic neutrons will always reach the devices and cause soft errors. Ultimately, high-energy cosmic neutron radiation defines the SER limit.

Mitigation by process and technology tweaks

To some extent, process and technology choices can address the remaining SER. Substrate structures or dop-

ing profiles that minimize the depth from which carriers can be collected can do much to reduce Q_{coll} , thereby reducing SER. In DRAM, designers have used multiple-well isolation to reduce charge collection. Researchers have also suggested well-based mitigation technologies for CMOS logic.¹⁰ Guard ring structures around sensitive junctions in SRAM devices have also provided SER robustness, at the expense of SRAM density. Although bulk silicon substrates are the mainstay of commercial microelectronics, for specialized applications some manufacturers use substrates incorporating a very thin silicon layer on a thicker layer of buried oxide (silicon on insulator, or SOI). Researchers have shown that SOI substrates reduce SER sensitivity.¹¹⁻¹³

Ultimately, though, improvements garnered by substrate engineering provide a limited path for mitigating soft errors. Most process solutions reduce the SER by less than an order of magnitude, at the expense of additional process complexity, yield loss, and substrate cost. An exception to this trend is a recently reported process solution using additional capacitance provided by an embedded DRAM capacitor attached to all sensitive nodes to increase the Q_{crit} of SRAM and logic devices.¹⁴ This approach doesn't use extra area but does add the expense of several additional process steps needed to define the embedded DRAM structures. Additionally, although increasing the Q_{crit} does reduce the SER by up to 250×, this reduction isn't enough for many high-reliability applications.

Mitigation by simple design or layout tweaks

Design and layout changes can reduce radiation sensitivity significantly. Any change that increases Q_{crit} while maintaining or reducing Q_{coll} will improve a device's SER performance. For example, a typical high-density SRAM cell consists of six transistors: two allow data to be read and written to and from the cell, and four make up the two cross-coupled inverters responsible for maintaining the data state. Q_{crit} is a function of the storage node capacitance and voltage, and of an additional term for the restoring charge supplied by the pull-up/pull-down transistor. This restoring term is proportional to the cell's switching time and the current from the load transistor. Increasing the load transistor's current drive or increasing the SRAM cell's switching time increases the cell's robustness against corruption. Thus, Q_{crit} can increase significantly with additional or larger drive transistors, so that a larger restoring current can be provided during a radiation-induced transient. Adding resistance between the two inverters can increase the time to flip the cell,¹⁵

effectively allowing the pull-up/pull-down transistor more time to restore the data state. (This approach affects the cell's write time and in high-speed technologies is not a realistic solution.)

Adding extra transistors also adds additional sensitive area and therefore appears self-defeating. However, the range of radiation events encountered in the terrestrial radiation environment deposits the charge within a few microns of the struck node. Therefore, if the components are physically separated to ensure that these "typical" single events cannot affect both transistors driving the same node, then a device with robust data states can form.¹⁶ (This approach is based on the exceedingly small probability of having multiple events in the same device node at the same time.) This is an effective approach in sequential logic, but it's very expensive for embedded memories, where it incurs a large area penalty and moderate power and speed penalties.

Mitigation by error detection and correction

By far, the most effective method of dealing with soft errors in memory components is to use additional circuitry for error detection and correction. In its simplest form, error detection consists of adding a single bit to store the parity (odd or even) of each data word, regardless of word length. Upon retrieval of data, a check compares the parity of the stored data with that of its parity bit. If a single error has occurred, the data parity won't match the bit parity. Thus, the parity system enables soft-error detection for a minimal cost in terms of circuit complexity and memory width (each word increases by only a single bit). The two disadvantages of this system are that the detected error cannot be corrected, and the check won't reveal a double error because the parity will match. This is true for any even number of errors: If the data was stored with odd parity, the first error changes the odd parity to even parity (a detectable error), but the second error changes the parity back to odd (a nondetectable error).

Error detection and correction (EDAC) or error-correcting codes (ECC) address these shortcomings. Designers typically achieve error correction by adding extra bits to each data vector and encoding the data so that the *information distance* between any two possible data vectors is at least three. More parity bits and additional circuitry can produce larger information distances, but in general most designers favor single-error-correction, double-error-detection (SECD-ED) schemes. In these systems, a single error (a change of plus or minus 1 in information space) leaves no

chance of mistaking the corrupted vector for its nearest neighbors, because the information distance is 3. In fact, two errors in the same *correction word* will also produce a valid error vector. The only limitation is that with two errors the error vector will not be unique to a single data value; therefore, the system supports only detection (not correction) of double-bit errors. A 64-bit-wide memory needs eight correction bits to detect two errors and correct a single error. Because most soft-error events are single-bit errors, EDAC/ECC protection provides a significant reduction in failure rates (typically, more than a 10,000× reduction in effective error rates). However, this reduction comes at a higher cost in terms of design complexity, additional memory required, and inherent latency introduced during access, parity check, and correction.

Design and layout tricks

Design and layout tricks can harden sequential and combinatorial logic, analogous to SRAM hardening. Because designers use fewer logic gates than in the high-density SRAM cells of most chips, and bit density isn't as crucial as in large memory arrays, logic design solutions can be more comprehensive. Most design approaches rely on node redundancy (using multiple storage nodes for each data state). The layout of these nodes typically minimizes the probability of an event having enough energy to disrupt two or more nodes in the system.¹⁷ Since the charge transients from radiation then affect only a single node, node redundancy ensures that the data state of the logic device is preserved.

The analog of error correction in sequential logic involves multiple identical logic paths feeding into a majority voting (two out of three) circuit. Basically, this architecture allows a soft error in a single logic path to be ignored because the other two logic paths constitute a majority; thus, the correct data “wins” the vote. This method uses three times the chip area and requires specialized simulation tools to identify the critical logic paths. (The high cost induces designers to protect only the most sensitive paths.) Time-multiplexed designs can also offer robustness against soft errors and SET because the input is sampled at several different times and a voter circuit sets the output on the basis of matching inputs (for example, two out of three). This approach works because the probability of two independent errors occurring in the same circuit path within a small time interval is exceedingly low. Using time- and spatial-multiplexed designs can build in even more robustness, at an increased cost.^{18,19}

Mitigation by system redundancy

The final and most ambitious form of redundancy is the use of duplicate or redundant systems—that is, multiple identical components running in lockstep (executing the same code at the same time). In a dual-component system, detection of a mismatch between devices triggers a restart. In systems with more than two units, a majority voting strategy can make restarting unnecessary. This is the most expensive redundancy scheme, but it does reduce soft failure rates to near 0, providing the necessary reliability for certain long-term remote or mission-critical computer applications.

Product perspectives

As mentioned earlier, 1 FIT is one error in a billion device hours, and advanced processors with large multi-megabit embedded SRAM can easily have soft-failure rates exceeding 50,000 FITs per chip. An SER of 50,000 FITs is equivalent to about one soft failure every two years, assuming the component is in service 24 hours a day. Will a digital signal processor failure rate of 50,000 FITs in a cell phone application affect a customer's perception of the cell phone's reliability? Probably not. Given that the phone won't be operating all the time and that the soft failure can occur anywhere in the chip (only if the error occurs in one of a few critical bits crucial to the phone's operation will the error be perceived), the cell phone probably won't fail once in its lifetime due to soft errors. Therefore, for single-user applications, it doesn't make sense to implement costly error correction or redundancy, even when the SER is very high.

However, using that same chip in a telecom base station, as a component in a mainframe server, or in a life-support system creates a different situation. Reliability requirements are much higher in such systems, and many chips operate in parallel, so that it's necessary to multiply the single-chip SER of one soft fail every two years by the number of chips in the system. Thus, one failure every two years for a single chip becomes a failure rate of once a week for a system with 100 chips. For such applications, error correction is mandatory.

Figure 6 shows the monthly number of soft errors as a function of the number of chips in a system and the amount of SRAM integrated in each chip. Logic SER is not included, and the failure rates are based on an uncorrected 1.6-kFIT-per-megabit SRAM SER. The key point is that the level of mitigation required to meet the customer's reliability expectations depends much more on the end application's reliability requirements than on the component's specific SER.

AT TERRESTRIAL ALTITUDES, three mechanisms are responsible for soft errors: the reaction of high-energy cosmic neutrons with silicon and other device materials, the reaction of low-energy cosmic neutrons with high concentrations of ^{10}B in the device, and alpha particles emitted from trace radioactive impurities in the device materials. While the DRAM system SER is relatively unchanged by scaling, SRAM and peripheral logic system SERs are increasing rapidly with each new technology node. In fact, soft errors are now the dominant failure mode of all reliability mechanisms combined (in qualified products). However, because the impact of soft errors on computer systems is extremely application dependent, in single-user commercial applications soft errors are often not a concern, while for larger (multichip) or high-reliability applications, error correction and/or redundancy techniques are mandatory to reduce soft errors to a tolerable level.

As the complexity of microelectronic components continues to increase, design and reliability engineers will need to address several key areas to enable advanced SoC products in the commercial sector. The first challenge will be to develop improved and more accurate system-level modeling of soft errors, including not just device and component failure rates but architectural and algorithmic dependencies as well. With these improved models, the next challenge will be to develop optimized memory and logic mitigation designs that offer a good balance between performance and robustness against soft errors and SETs. Finally, the last challenge will be the development of viable commercial fault-tolerant solutions that will render complex systems relatively insensitive to soft errors at the operational level. ■

■ References

1. T.C. May and M.H. Woods, "A New Physical Mechanism for Soft Error in Dynamic Memories," *Proc. 16th Int'l Reliability Physics Symp. (IRPS 78)*, IEEE Electron Devices Soc., 1978, pp. 33-40.
2. J.F. Ziegler and W.A. Lanford, "The Effect of Sea-Level Cosmic Rays on Electronic Devices," *J. Appl. Phys.*, vol. 52, no. 6, Jan. 1981, pp. 4305-4312.
3. E. Normand, "Single Event Upset at Ground Level," *IEEE Trans. Nuclear Science*, vol. 43, no. 6, Dec. 1996, pp. 2742-2750.
4. R.C. Baumann et al., "Boron Compounds as a Dominant Source of Alpha Particles in Semiconductor Devices," *Proc. 33rd Int'l Reliability Physics Symp. (IRPS 95)*, IEEE Electron Devices Soc., 1995, pp. 297-302.
5. R.C. Baumann and E.B. Smith, "Neutron-Induced ^{10}B

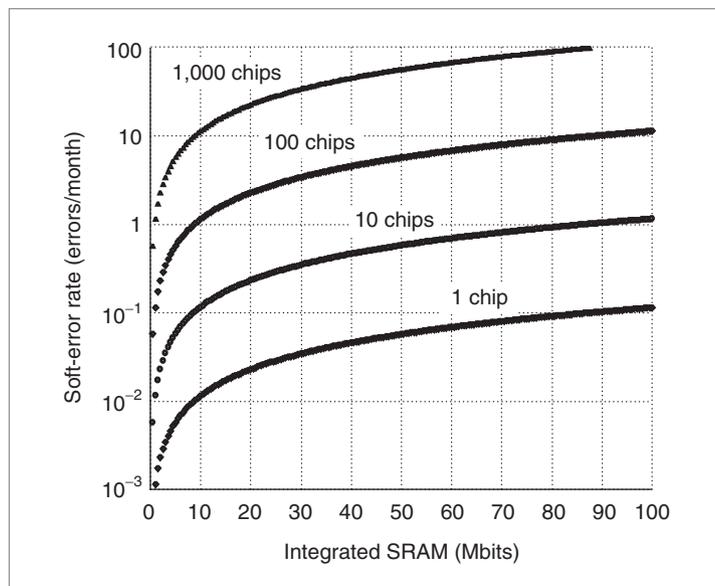


Figure 6. Monthly system soft-error rate as a function of the number of chips in the system and the amount of embedded SRAM per chip.

Fission as a Major Source of Soft Errors in High Density SRAMs," *Microelectronics Reliability*, vol. 41, no. 2, Feb. 2001, pp. 211-218.

6. C.M. Hsieh, P.C. Murley, and R. O'Brien, "Dynamics of Charge Collection from Alpha-Particle Tracks in Integrated Circuits," *Proc. 19th Int'l Reliability Physics Symp. (IRPS 81)*, IEEE Electron Devices Soc., 1981, pp. 2-9.
7. P.E. Dodd and L.W. Massengill, "Basic Mechanisms and Modeling of Single-Event Upset in Digital Microelectronics," *IEEE Trans. Nuclear Science*, vol. 50, no. 3, June 2003, pp. 583-602.
8. S. Buchner et al., "Comparison of Error Rates in Combinational and Sequential Logic," *IEEE Trans. Nuclear Science*, vol. 44, no. 6, Dec. 1997, pp. 2209-2216.
9. P.E. Dodd et al., "Production and Propagation of Single-Event Transients in High-Speed Digital Logic ICs," *IEEE Trans. Nuclear Science*, vol. 51, no. 6, Dec. 2004, pp. 3278-3284.
10. M.P. Baze, S.P. Buchner, and D. McMorow, "A Digital CMOS Design Technique for SEU Hardening," *IEEE Trans. Nuclear Science*, vol. 47, no. 6, Dec. 2000, pp. 2603-2608.
11. O. Musseau, "Single-Event Effects in SOI Technologies and Devices," *IEEE Trans. Nuclear Science*, vol. 43, no. 2, Apr. 1996, pp. 603-613.
12. P. Roche et al., "Comparisons of Soft Error Rate for SRAMs in Commercial SOI and Bulk Below the 130-nm Technology Node," *IEEE Trans. Nuclear Science*, vol.

- 50, no. 6, Dec. 2003, pp. 2046-2054.
13. E.H. Cannon et al., "SRAM SER in 90, 130, and 180nm Bulk and SOI Technologies," *Proc. 42nd Int'l Reliability Physics Symp. (IRPS 04)*, IEEE Electron Devices Soc., 2004, pp. 300-304.
 14. R. Wilson, "ST Tames Soft Errors in SRAM by Adding Capacitors," *EE Times*, 13 Jan. 2004, <http://www.eetimes.com/semi/news/showArticle.jhtml?articleID=18310682>.
 15. L.R. Rockett, "Simulated SEU Hardened Scaled CMOS SRAM Cell Design Using Gated Resistors," *IEEE Trans. Nuclear Science*, vol. 39, no. 5, Oct. 1992, pp. 1532-1541.
 16. D. Bessot and R. Velazco, "Design of SEU-hardened CMOS Memory Cells: The HIT Cell," *Proc. 2nd European Conf. Radiation and Its Effects on Components and Systems*, IEEE Press, 1993, pp. 563-570.
 17. T. Calin, M. Nicolaidis, and R. Velazco, "Upset Hardened Memory Design for Submicron CMOS Technology," *IEEE Trans. Nuclear Science*, vol. 43, no. 6, Dec. 1996, pp. 2874-2878.
 18. L. Anghel and M. Nicolaidis, "Cost Reduction and Evaluation of a Temporary Faults Detecting Technique," *Proc. Design, Automation and Test in Europe (DATE 00)*, IEEE CS Press, 2000, pp. 591-598.
 19. D.G. Mavis and P.H. Eaton, "Soft Error Rate Mitigation

Techniques for Modern Microcircuits," *Proc. 40th Int'l Reliability Physics Symp. (IRPS 02)*, IEEE Electron Devices Soc., 2002, pp. 216-225.



Robert Baumann is a Distinguished Member of the Technical Staff at Texas Instruments, where he is responsible for characterizing and mitigating radiation effects in advanced semiconductor technologies. His research interests include the impact of radiation effects on microelectronics reliability. Baumann has a BA in physics from Bowdoin College and a PhD in electrical engineering from Rice University. He is a Fellow of the IEEE.

■ Direct questions and comments about this article to Robert Baumann, Texas Instruments, Silicon Technology Development Component Reliability Group, 13560 North Central Expressway, MS 3740, Dallas, TX 75243; rbaum@ti.com.

For further information on this or any other computing topic, visit our Digital Library at <http://www.computer.org/publications/dlib>.

Join the IEEE Computer Society online at computer.org/join/

Complete the online application and

- Take Web-based training courses in technical areas for free
- Receive substantial discounts for our software development professional certification program
- Get immediate online access to *Computer*
- Subscribe to our new publication, *IEEE Security & Privacy*, or any of our 22 periodicals at discounted rates
- Attend leading conferences at member prices
- Sign up for a free e-mail alias—you@computer.org

THE WORLD'S COMPUTER SOCIETY