

Cosmic RADIATION COMES TO ASIC AND SOC DESIGN

AS IC-PROCESS GEOMETRIES SHRINK, SINGLE-EVENT EFFECTS, SUCH AS SOFT ERRORS AND LATCH-UPS, WILL SOON BECOME PRIMARY CONCERNS FOR DESIGNERS OF ASICs AND SOC.

At a glance.....48
How wrong is an error?50
For more information56

SEEs (SINGLE-EVENT EFFECTS), such as soft errors, have since the early 1980s appeared in commercial electronics, but they are now becoming the dominant reliability-failure mechanism in modern CMOS technologies. These types of errors forced Sun to recall workstations in the late '90s and, *EDN* has learned, also caused failures in memories and ASICs controlling a Cisco router. Experts say that these effects will soon become more commonplace in logic as processes head into 65- and 45-nm nodes. Fortunately, military and aerospace designers offer several techniques, and EDA vendors offer new tools to address SEEs as they make their way into mainstream design. SEEs can cause temporary glitches or, in the worst case, disable a system, so users must evaluate design techniques and cost trade-offs for their target applications when deciding whether to tackle SEEs that may adversely affect their designs.

Researchers in the mid-1970s first observed SEEs in orbiting satellites and have since tracked

SEEs as growing problems at lower altitudes in lock step with IC-process-geometry reductions, lower voltages, and increasing clock speeds. SEEs have moved from a problem first affecting a range of space and aeronautical devices to one affecting memory ICs, FPGAs, and devices using combinatorial logic.

Memories have led the way when it comes to speed, low power, and density, and thus have become the proverbial canary in the coal mine when it comes to SEEs, according to Christopher Nicklaw, senior staff engineer at Silvaco. But the canary is no longer singing. "This is one of those unique times in which military electronics' needs are coexisting with the commercial requirements," says Nicklaw. "There is an agreement in both camps that this is a problem that we need to look into."

Robert Baumann, a member of the Silicon Technology Development Component Reliability Group at Texas Instruments, agrees: "Years ago, if you compared one logic flip-flop with one



ZAPPED!

coverstory *By Michael Santarini, Senior Editor*

6T [six-transistor] SRAM cell, there would be five or six orders of magnitude difference in sensitivity to SEEs,” says Baumann, one of the foremost researchers of SEEs and chairman of Sematech’s (Semiconductor Manufacturing Technology’s) JEDEC (Joint Electron Device Engineering Council) JESD89 de facto standard for commercial radiation testing (Reference 1). “Going toward 65 nm, we are almost getting parity with SRAM and flip-flop sensitivity.”

Baumann and Paul Dodd, another foremost expert on SEEs and acting manager for the radiation-effects department at Sandia National Laboratories (Albuquerque, NM), says that commercial designs are also more frequently encountering SEEs but that designers are commonly missing or misidentifying them as other failures. “It could be happening on everyone’s PC, but instead everyone curses Microsoft,” says Dodd. “Software bugs probably cause a lot of those blue-screen problems, but you can trace some of them back to radiation effects.” And designers cannot yet quantify the breadth of the problem because, as IC-design and EDA consultant Pallab Chatterjee points out, “It is something companies don’t brag about.”

Sun Microsystems encountered a public-relations nightmare when SEEs came to light, causing Sun server workstations to require occasional resets, which led to an embarrassing 1000-unit recall (Reference 2). But Sun is not alone, and the problem endures. Cisco Systems also encountered SEE failures with its 12000 series router line cards, reporting failures of memory and ASICs and subsequent debugging attempts for soft errors in that router, which sells for approximately \$200,000. Cards are showing memory-parity or ASIC errors that may have resulted in a card’s reloading with a two- or three-minute recovery. Data passes normally after the card reloads, according to a field note (Reference 3).

Baumann says that TI does extensive work in this area because its off-the-shelf DSPs find use in a broad range of applications. He says that some of TI’s customers use the DSPs in air-traffic-avoidance radar, and others can use the same part in cell phones. “Mitigation is ex-

AT A GLANCE

- ▶ The largest EDA vendors offer Spice simulation and analysis technologies for SEE (single-event-effect) detection but don’t yet offer tools that explicitly target the task.

- ▶ ICs can undergo accelerated testing in extreme environments or real-time testing.

- ▶ Commercial tools are available to help users detect whether their designs are encountering SEEs, and SOC (system-on-chip) designers can borrow several techniques from rad-hard and memory design to help fix them.

- ▶ The new trend in military and aerospace applications is to use a hardened-by-design approach.

- ▶ Memory and commercial-IC vendors target SEEs using error detection and correction.

pensive. No one wants to do it if they don’t have to, but sometimes ... you have to put in enough of a safety net that if someone uses your part in an application you weren’t expecting, it has at least some minimal reliability performance.”

The subject of SEEs is old hat to designers of radiation-hardened devices and DRAMs and to companies such as International Rectifier, which has for years been SEE-hardening designs for aerospace applications in which radiation exists and reliability is a must (Reference 4). Researchers in the late ‘70s observed the problem in terrestrial applications, and it crept in the 1980s into commercial applications running terrestrially. Alpha particles from package materials caused SEEs that first affected DRAM, and, through process refinement, these particles have become less common, but they still account for roughly 30% of all SEEs, Baumann estimates. Thermal neutrons from cosmic

radiation of energy less than 15 eV (electron volts) or terrestrial high-energy cosmic particles, such as neutrons, protons, and muons, also cause SEEs (Figure 1). These particles cause reactions with silicon and oxygen nuclei and break them apart, leaving ionizing fragments, which in turn generate a charge—not a desirable effect, according to Baumann.

Manufacturers measure SEEs in FITs (failures in time); one FIT equals one failure in 10⁹ device hours. Typical hard failures, such as electromigration, have a FIT of one to 50, and the aggregate failure rate is 200 FITs. Unchecked soft errors can have FITs of 50,000 per IC, however. The saturation of neutrons, which is the radiation event that causes soft errors, increases at higher altitudes. A typical flight altitude of 30,000 feet increases by 300 times the chance of encountering an SEE, according to Baumann. According to Steve Wender, group leader at the Los Alamos (NM) Neutron Science Center testing facility, placing six inches of steel shield around a device achieves only a twofold improvement in FIT. At least five classes of SEE exist. If any of these effects occurs in a toy or a cell phone, it may be trivial and might not preclude a further purchase, but if they occur in a mission-critical device, such as a pacemaker, a drive-by-wire brake system, or even a handheld device for transferring funds between bank accounts, a three-second glitch can become a catastrophe (see sidebar “How wrong is an error?”).

Michael Buehler-Garcia, director of marketing at EDA start-up iRoC, points out that SEE problems become complex as the industry moves to SOC (system-on-chip) design because no one team typically owns all the areas—on-chip memory, logic, IP (intellectual property), and software—that SEEs can affect (Reference 5). Bruce Takala, a researcher at the Los Alamos laboratory, agrees: “As they put more and more functions, such as built-in SRAM, on processors, those devices become susceptible, as well,” he says.

Bel Lazar, vice president of High Reliability Operations at International Rectifier, says that, even on relatively straightforward designs, SEE prevention starts at knowing the process el-

TABLE 1—LOGIC-PROTECTION TRADE-OFFS


| | Effort | Protection efficiency | Timing impact | Area impact |
|---------------------------------|--------|-----------------------|---------------|-------------|
| Triple-mode redundancy | Low | Good | Medium | Three times |
| Detection and system correction | Medium | Good | Medium | Two times |
| Hardened library | High | Fair | Low | Medium |
| Other techniques | High | Fair | Low | Low |

ements and the design's sensitivity to SEEs. "When you are designing a dc/dc converter, you have to worry that your component is SEE-immune and then that the circuit is SEE-immune as a circuit or as the sum of all components," he says.

TOOLS STEP IN TO DETECT SEEs

Although the three largest EDA vendors—Cadence, Synopsys, and Mentor Graphics—offer Spice simulation and analysis technologies for SEE detection, they don't yet offer tools that explicitly target the task. This lack of tools is largely a result of the fact that SEE has not yet become a mainstream problem or an eminently large revenue stream. Soft errors are beginning to appear on the radar screen as problems. However, these problems are not yet among those that most of these companies' customers experi-

Figure 1



Neutron concentration and the likelihood of encountering an SEE is greatest at an altitude of 50,000 feet. At 30,000 feet, the typical cruising altitude for commercial aircraft, ICs are 300 times more likely to encounter an SEE than at sea level.

HOW WRONG IS AN ERROR?

Military and aerospace devices, commercial memory, and logic have so far been susceptible, to varying degrees, to four classes of SEEs (single-event effects): soft errors, SELs (single-event latch-ups), SEFIs (single-event functional interrupts), and SEBs (single-event burnouts). A fifth type, SETs (single-event transients), will likely become a problem in commercial designs at the 65- and 45-nm process nodes. They are already problems at the 90-nm process, according to EDA start-up iRoC Technologies, which tests for SEE effects (Table A).

Soft errors, or SEUs (single-event upsets), by far the most common SEEs, are transient faults. External radiation—typically from cosmic rays—cause SEUs. They occur when, for example, a radiation event switches a state from a one to a zero or a zero to a one. When soft errors affect a design, the system typically requires a reset and then runs normally until another radiation event occurs.

SELs occur when a radiation event causes a sudden voltage drop and essentially creates a

dual bipolar circuit that locks the circuit into a high-current state, which eventually overheats—and may even destroy—the device.

Robert Baumann, a member of the Silicon Technology Development Component Reliability Group at Texas Instruments, says that an SEL is similar to an electrical latch-up and that any circuit that is susceptible to electrical latch-ups is also sensitive to SELs. Paul Dodd, acting manager of the radiation-effects department at Sandia National Laboratories, says that rebooting can usually correct latch-ups, but that the original latch-up may have caused immediate or latent damage, such as electromigration, to the device.

An SEFI occurs when a microcomputer program counter encounters a radiation event and 10 lines flip the counter, so that the part no longer understands what it is doing. A SEFI can trigger a device to run its power-off self-test program. "The device can just go into perpetual power-off, self-test mode

until it is reset," Dodd says.

SEBs, also known as single-event gate ruptures, occur when a large radiation event causes a leakage path in the gate oxide and destroys gates in a design, rendering the chip and, most likely, the system nonfunctional. SEBs are rare and dramatic and have occurred to date only in high-voltage power FETs.

SETs can affect standard-cell logic, memories, and programmable logic. SETs occur when a radiation event causes a false signal that propagates through logic and, because the clocks are so fast on new process geometries, the logic clocks in as a signal on the latching edge of the clock.

"Voltage or current transients have always been present," says Dodd. "But the difference is: As we go to logic at higher speeds, those transients can propagate through the circuit. The fact that the node returns to its original value doesn't matter, because you've already clocked down the line a false transient that looks like a signal. And, once that signal propagates through logic gates and finds a memory element, it can flip a bit and become an SEU." Michael Buehler-Garcia, director of marketing at iRoC Technologies, claims that the company is seeing SETs in control logic at the 90-nm process node.

TABLE A—SOFT ERRORS AT VARIOUS PROCESS NODES

| Process node | Application | Sensitivity to soft errors |
|---------------------|--------------------|----------------------------|
| 0.25 micron | Consumer | None |
| | Networking/storage | None |
| | Aerospace/military | Memory and logic |
| 0.18 to 0.13 micron | Consumer | None |
| | Networking/storage | Memory |
| | Aerospace/military | Memory and logic |
| 90 nm | Consumer | Memory and logic |
| | Networking/storage | Memory and logic |
| | Aerospace/military | Memory and logic |
| 65 nm and below | Consumer | Memory and logic |
| | Networking/storage | Memory and logic |
| | Aerospace/military | Memory and logic |

ence, according to Rohit Kapur, a scientist at Synopsys. As a result, he says, the company has received no request for soft-error support. Although Synopsys offers no tools specifically for SEEs, International Rectifier now uses tools from ISE (Integrated Systems Engineering), which Synopsys now owns, for device-level SEE simulation. “ISE has some capabilities for simulating SEEs, although you need to know a lot about what radiation is going to be doing,” says Milt Boden, director of radiation-hardened and high-reliability silicon R&D at International Rectifier. “At a Spice-model level, you put in current pulses at different nodes in your circuit and see how different nodes behave during different parts of your timing cycle.” The difficulty with the approach is knowing what pulse to put in, he says, noting that the company is still researching that topic.

Consultant Chaterjee notes that, in military and aerospace applications, contractors usually supply their own models and tolerance specifications, such as pulse strength. And small EDA vendors, such as iRoC are just now starting to offer those models. The company got its start putting commercial ICs through rigorous radiation-event testing and developed tools for standard-cell and custom-designed soft-error detection and analysis. The company based the soft-error models of those tools on SEE testing it conducted for numerous IC vendors. Michael Nicolaidis, PhD, chief technology officer and co-founder of iRoC, says that ICs can undergo accelerated testing at labs with radiation equipment—in this case, neutron beams—such as the Los Alamos lab, TRIUMF (Tri Universi-

TABLE 2—LOGIC-PROTECTION TRADE-OFFS

| Protection method | Effort | Timing impact | Area impact |
|-------------------|--------|---------------|--|
| Standard ECC | Low | High | High for short words, low for long words |
| Low-area ECC | Medium | High | Low for short words |
| High-speed ECC | Medium | Medium | High |
| Embedded | High | Low | Low |
| System level | High | Low | Low |

ties Meson Facility) Laboratory (Vancouver, BC, Canada), and TSL (The Svedberg Laboratory, Uppsala, Sweden). Alternatively, ICs can undergo real-time testing in labs in extreme environments. One such real-time-testing facility, which resembles a villain’s lair in a *James Bond* movie, is the International Foundation High Altitude Research Station (Jungfraujoch, Switzerland) at an elevation in the Alps of more than 11,000 feet, which represents a neutron flux 11 times that of sea level. On the other extreme, the Underground Laboratory of Modane in the Frejus Tunnel, lies under 6000 feet of rock in a tunnel through the base of the Alps between France and Italy. It represents a neutron flux that is 100,000 feet lower than sea level.

Manufacturers put most devices, which are subject to time-to-market constraints, through accelerated testing, in which the testers direct protons into a tungsten-eradiating target product. According to the Los Alamos Labs’ Wender, the lab’s testing spectrum closely matches the atmospheric spectrum that cosmic rays cause but is a million times more intense than environmental rates. “We’ve seen a wide span of results—from devices that fail before the shutter is fully open to devices that are quite resistant,” he says, noting the results are proprietary. Such testing has allowed iRoC to derive tools and underlying models that allow users

to simulate radiation strikes, taking into account both the strength of a strike and all possible angles of that strike on a part of a device. The company’s SoCFIT full-chip-analysis tool determines which areas of a design are susceptible to soft errors (Figure 2). The company’s TFIT tool provides detailed, Spice-level analysis of the problem areas once SoCFIT finds them. These models find use in military, aerospace, and commercial applications.

Silvaco International is one of only a few vendors offering commercial analysis with SEE-analysis features. Silvaco a few years ago expanded its SmartSpice circuit simulator to accommodate SEE-aware Spice models. The company also recently completed a Defense Advanced Research Projects Agency contract to enhance its Harmony AMS mixed-signal simulator with soft-error-rate and other radiation-effect capabilities. That tool is commercially available, but some of Silvaco’s other technologies are available only to military contractors.

In addition to iRoC and Silvaco, Denali Software also offers C-level memory models with built-in soft-error tolerance. And EDA start-up Alternative System Concepts is trying to raise money to further develop its Virtual TMR (triple-modular-redundancy) tool. Programmable-logic vendors also offer tools to aid in radiation-hardening designs. Xilinx, for example, has for years offered a TMR tool, and its parts also come with built-in error-correction and -detection code that can flag and correct single- and double-bit—but not multibit—upsets. Fortunately, commercial tools are available to help users detect whether their de-

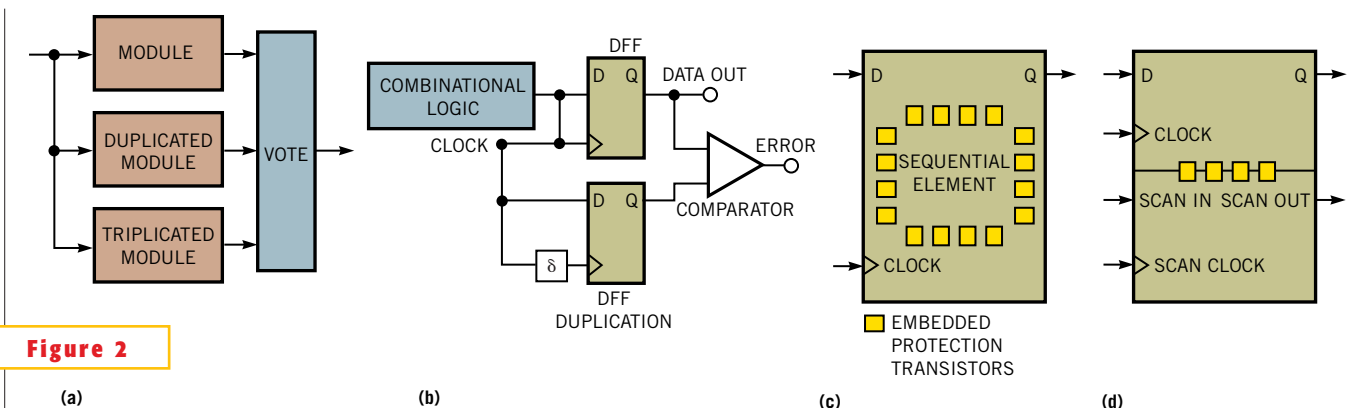


Figure 2

Popular protection techniques for logic include redundancy (a and b), shielding (c), and scan (d) (courtesy iRoC Technologies).

signs are encountering SEEs, and SOC designers can borrow several techniques from rad-hard and memory design to help fix them. The techniques have their advantages and disadvantages, and the tools are expensive (Table 1).

BORROWING FROM MILITARY AND AEROSPACE

When it comes to soft errors, “speed kills,” says Sandia National Labs’ Dodd. “The faster you are, the more likely you are to be susceptible to soft errors and transients.” Thus, traditional techniques that military and aerospace designers employ to attack SEEs—especially soft errors—typically involve slowing an IC’s clock to a point at which it doesn’t clock in false signals that cause an upset. The mainstay feedback technique in radiation-hardened military and aerospace adds cross-coupled feedback resistors inside SRAM. This method adds an RC delay between the legs of a standard six-transistor SRAM, says Dodd. “It basically slows the feedback process. And if you slow it down enough, the system recovers from the transient fault before the signal feeds back to the other side and locks in the error.” Employing the technique reduces overall system performance and thus makes the technique prohibitive for most commercial applications.

The new trend in military and aerospace applications is to use a hardened-by-design approach. “The more you get behind the speed curve, the more you want to use the stuff people are using on the ground,” says Dodd. This technique gives designers of traditional rad-hard-

ened equipment the best bang for their buck, he says. “Instead of using process techniques to slow things down, you use internal redundancy in the circuits themselves, where the circuit is inherently hardened to radiation coming in,” he notes. Creating internally redundant circuits means implementing more transistors to get the same functions. A six-transistor SRAM cell turns into a 10- or 12-transistor cell, which adds cost, latency, and space.

TI’s Baumann says that ICs for aerospace applications commonly employ TMR. “TMR is not feasible for most commercial applications, so the trick is determining which logic elements you need to make more robust and which ones you can skip,” he says. For example, a designer would want to make a register file holding a critical address more robust. And, likewise, the designer would typically omit branching instructions, which, if an SEE hit them, would require more CPU cycles. “In advanced geometries, even when you are using hardened-by-design methods, you may have to put a large separation between the redundant halves of a circuit, and that approach creates a wiring nightmare, says Dodd. “The farther apart you space them, the more wiring or longer traces are needed, which increases timing delay.”

Designers like the hardened-by-design technique because it doesn’t require a rad-hard foundry but instead obtains the necessary hardness from the circuit itself. “There is a limiting factor on the hardened-by-design techniques because dual-

node upsets are starting to occur that defeat the added redundancy,” says Dodd. He explains that, although dual-node redundancies don’t mean that designers can’t use the circuit, they make the system more complex. He says that designers must make trade-offs in performance, area, or power consumption in using these hardened-by-design configurations. “It is easier for us to take in the space-radiation-effects community than the commercial guys who have to come out with a next-generation product that has a 20% higher clock speed and a lower price.” He says that designers need to make some hard choices when trading off reliability versus performance.

COMMERCIAL METHODS

Memory and commercial-IC vendors take a different approach to targeting SEEs—error detection and correction—from manufacturers targeting aerospace and military applications. In error detection, users add extra bits and encoding and decoding, so if an error occurs in a data word, the extra bits reflect that error. For example, a single bit may read 1,1, and its error-correction bit reads 1,1 if no error occurs. If an SEU (single-event upset) occurs, the bit reads 0,1 or 1,0. And, if both switch from 1,1 to 0,0, a double error has occurred. The system doesn’t detect the double error because it considers 0,0 to be a valid data state. Baumann points out that the method has two shortcomings: Error detection cannot do correction and it does not detect whether a double error has occurred

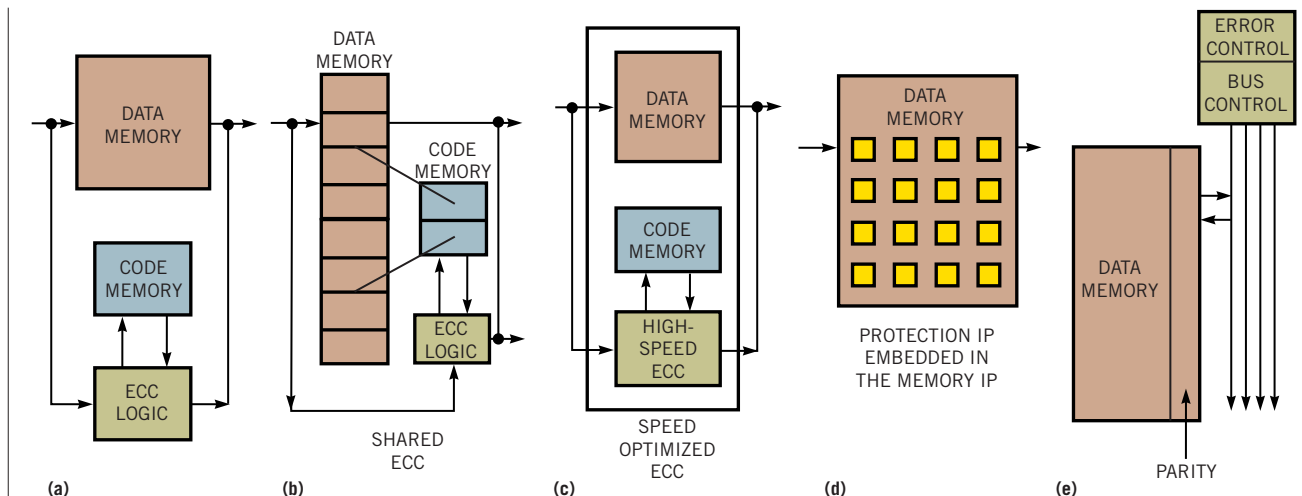


Figure 3

ECC (a, b, and c), shielding (d), and error management (e) protect memory from SEEs (courtesy iRoC Technologies).

if there is parity between two signals.

To get past the shortcoming, Baumann says, users can upgrade to correction by using ECC (error-correction code) (Figure 3). In this method, designers add more bits onto the extra bit for error detection. The typical method is to use a total of three bits—1,1,1 or 0,0,0—so that any variation, such as 1,1,0 or 1,0,1, signals that an upset has occurred. Because the approach can isolate the single-bit error, the technique can also correct the event with ECC. The system can also detect the occurrence of—but cannot correct—double-bit errors. And the ECC cannot detect a multiple-bit error, such as changing 0,0,0 to 1,1,1 (Table 2).

ECC, such as Hamming code, is also cumbersome, because each information word requires the addition of an extra byte or 8 bits. That approach is acceptable, says Baumann, if designers are building a big memory array, but if they have a bunch of small SRAM blocks throughout an SOC, this approach can cause as much as a 50% increase in area. Nicolaidis from iRoC adds: “A single neutron as it strikes can emit multiple secondary particles, typically ions, that can strike multiple memory cells,” says iRoC’s Nicolaidis. “If the affected memory cells belong to the same memory word, then corrective code can’t fix it.” Baumann says that designers often try to use bit spacing so that, if a multiple error occurs, the system has a better chance of identifying it. In finer process geo-

**IF YOU STORE MORE CHARGE,
YOU HAVE MORE SIGNAL,
SO IT TAKES A LOT MORE
OF A RADIATION EVENT
TO UPSET YOUR CELL.**

metries, however, the spaced bits are often so close together that a single error can simultaneously switch multiple bits. A triple-bit switch, though rare today, would defeat most error-correction and -detection code that memory and FPGA devices employ, according to Nicolaidis (Reference 6). Row or column multiplexing can use 4-, 8-, or even 16-bit spacing, thus reducing the probability of double- and triple-bit failures to essentially zero.

Experts point out that materials can also play a part in hardening designs against SEEs. Most military- and aerospace-device manufacturers claim to employ processes that prevent SEEs. Baumann says that some commercial materials can also help. For example, STMicroelectronics uses giant capacitors on SRAM for hardening against SEEs. “If you store more charge, you have more signal, so it takes a lot more of a radiation event to upset your cell,” Baumann says. The use of SOI (silicon-on-insulator) techniques also makes silicon five

times more resistant to SEEs. “If you are using SOI for other reasons, it’s an added plus but it wouldn’t make sense economically to switch to SOI just for soft errors,” says Baumann. “If you do error correction properly, you get four orders of magnitude improvement.” □

REFERENCES

1. “Measurement and Reporting of Alpha Particle and Terrestrial Cosmic Ray Induced Soft Errors in Semiconductor Devices,” and Addendums 1 and 2, JESD-89, August 2001, www.jedec.org.
2. Lyons, Daniel, “Sun Screen,” *Forbes* magazine, Nov 13, 2000, www.forbes.com/global/2000/1113/0323026a.html.
3. Cisco 12000 Single Event Upset Failures Overview and Work Around Summary, April 15, 2003, www.cisco.com/en/US/products/hw/routers/ps167/products_field_notice09186a00801b3df8.shtml.
4. Mastipuram, Ritesh and Edwin C Wee, “Soft errors’ impact on system reliability,” *EDN*, Sept 30, 2004, pg 69, www.edn.com/article/CA454636.html?spacedesc=contributedFeature.
5. Dipert, Brian, “Banish bad memories,” *EDN*, Nov 22, 2001, pg 61, www.edn.com/article/CA181882.html.
6. Holbert, Keith E, “Single event effects,” Arizona State University, www.eas.asu.edu/~holbert/eee460/see.html.

TALK TO US

Post comments via TalkBack at the online version of this article at www.edn.com.



FOR MORE INFORMATION...

For more information on the products and technologies mentioned in this article, please contact the companies and organizations directly, and please let them know you read about them in *EDN*.

Alternative Systems Concepts

www.ascinc.com

Cadence

www.cadence.com

Cisco Systems

www.cisco.com

Defense Advanced Research Projects Agency

www.darpa.mil

Denali Software

www.denali.com

International Foundation High Altitude Research Station

www.ifjungo.ch

International Rectifier

www.irf.com

iRoC Technologies

www.irotech.com

Joint Electron Device Engineering Council

www.jedec.org

JungfrauJoch and Gornergrat

www.ifjungo.ch/foundation/foundation.html

Los Alamos Neutron Science Center

http://lansce/lnl.gov

Mentor Graphics

www.mentor.com

Sandia National Laboratories

www.sandia.gov

Sematech (Semiconductor Manufacturing Technology)

www.sematech.org

Silvaco

www.silvaco.com

STMicroelectronics

www.st.com

Synopsys

www.synopsys.com

Texas Instruments

www.ti.com

TRIUMF (Tri Universities Meson Facility) Laboratory

http://tcmsm.ca/intro/triumf

TSL (The Svedberg Laboratory)

www4.tsl.uu.se/tsl/tsl/

Underground Laboratory of Modane

http://lyoinfo.in2p3.fr/manoir/lsm_eng.html

Xilinx

www.xilinx.com