By Ritesh Mastipuram and Edwin C Wee, Cypress Semiconductor

SOFT ERRORS IN MEMORY DEVICES ARE BECOMING INCREASINGLY IMPORTANT CONSIDERATIONS IN SYSTEM DESIGN.

Soft errors' impact on system reliability

In the field of high-performance communication memory devices, it is critical for designs to be immune to soft errors or single-event upsets. During the SER (soft-error-rate) session of the 2003 IRPS (International Reliability Physics Symposium), Texas Instruments reliability scientist Robert Baumann stated that, "Soft errors induced the highest failure rate of all other reliability mechanisms combined." As device technology scales, the number of stages in the processor pipeline increases, the area efficiency of memory devices decreases, and a device's natural resistances against SEUs (single-event upsets) decreases.

It is important to first define the causes and effects of SEUs (**Figure 1**). SEEs (single-event effects) are associated with the change of states or transients in a device that energetic external radiation particles induce. You can classify SEEs into soft and hard errors. Soft errors are nondestructive, because resetting or rewriting the device restores normal behavior thereafter; hard errors are permanent. A common example of a hard error is an SEL (single-event latch-up).

Soft errors that are predominantly caused by external radiation are also known as SEUs. They result in transient, inconsistent errors in data that are unrelated to components or manufacturing failures. Intrinsic noise and interference can also cause SEUs; however, design engineers can accommodate these causes. SEUs manifest themselves as either SBUs (single-bit upsets) or MBUs (multiple-bit upsets). "SBU" refers to the flipping of one bit due to the passage of a single energetic radiation particle, where the physical separation from any other flipped bit is at least two memory cells. MBU refers to the flipping of several elements due to the passage of one or more radiation particles.

SEUs are random and rarely catastrophic, and they do not normally destroy a device. Many systems can tolerate some level of soft errors. For example, if you are designing a precompression capture buffer or a postdecompression playback buffer for an audio-, video-, or still-imaging system, an occasional bad bit may be unnoticeable and unimportant to the user. However, when you use memory elements in mission-critical applications to control system functions, soft errors can have a more serious impact and lead to not only corrupt data, but also a loss of function and system-critical failures.

GETTING WORSE, NOT BETTER

The SER problem first gained widespread attention as a memory-data issue in the late 1970s, because DRAMs began to show signs of random failures. As process technologies continue to shrink, the critical charge required to cause an upset is decreasing faster than the charge-collection area in the memory cell. Therefore, with smaller geometries, such as 90 nm, soft errors are more of a concern, and designers must take steps to control SER levels.

Technology scaling has been the primary engine for industry survival and is the driving factor for higher density, improved performance, and cost reduction. As device technology scales to deep-submicron gate lengths (0.25 microns to 90 nm and beyond), the cell size of memory products continues to decrease, thus driving the voltage lower (5V to 3.3V to 1.8V and smaller) and reducing the capacitance inside the cell (10 to 5 fF and smaller). Due to the lower capacitance, the critical charge (the minimum charge required for a cell to retain data) in

Figure 1

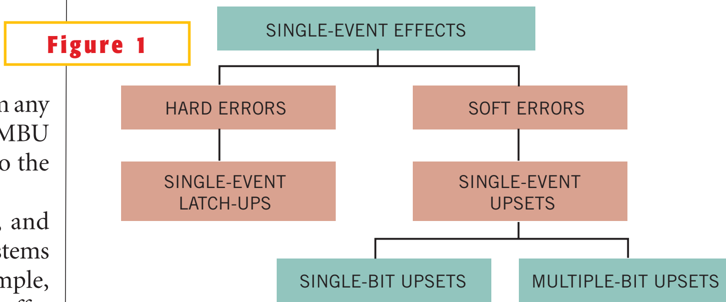

Single-event errors have many causes and effects.

memory devices continues to shrink, thereby decreasing their natural resistance to SEUs. Therefore, a low-energy alpha particle or a cosmic ray can disturb the cell.

The rate at which SEUs occur is given as SER, and you measure it in FITs (failures in time), which is the number of failures in 1 billion device-operation hours. A measurement of 1000 FITs corresponds to a MTTF (mean time to failure) of approximately 114 years.

The potential impact on typical memory applications illustrates the importance of considering soft errors. A cell phone with one 4-Mbit, low-power memory with an SER of 1000 FITs per megabit will likely have a soft error every 28 years. A high-end router with 10 Gbits of SRAM and an SER of 600 FITs per megabit can experience an error every 170 hours. For a router farm that uses 100 Gbits of memory, a potential networking error interrupting its proper operation could occur every 17 hours. Finally, consider a person on an airplane over the Atlantic at 35,000 ft working on a laptop with 256 Mbytes (2 Gbits) of memory. At this altitude, the SER of 600 FITs per megabit becomes 100,000 FITs per megabit, resulting in a potential error every five hours. The FIT rate of soft errors is more than 10 times the typical FIT rate for a hard reliability failure. Soft errors are not the same concern for cell phones as they can be for systems using a large amount of memory.

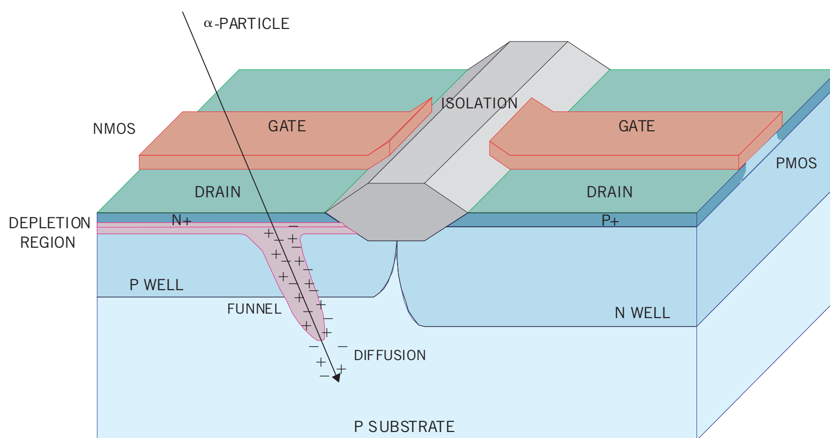

Figure 2

The alpha particle generates electron-hole pairs in its wake, causing the charges to drift so that the transistor sees a current disturbance.

TROUBLE SPOTS

The four common sources of SEUs are low-energy alpha particles, high-energy cosmic particles, thermal neutrons, and poor system design (**Table 1**). Low-energy alpha particles are generated by the radioactive decay of trace uranium-238 and thorium-232 in quartz filler used in mold compounds, or from polonium-210 in lead bumps that flip-chips use. These impurities release alpha particles with energy of 2 to 9 MeV (million electron volts). The energy required to form an electron-hole pair in silicon is 3.6 eV. Therefore, alpha particles can generate approximately 10^6 electron-hole pairs.

The electric field in the depletion region directly generates electron-hole pairs in its wake, causing the charges to drift so that the transistor sees a current disturbance (**Figure 2**). The depletion regions under the effect of the electric field collect free electrons. A fraction of the excess charge drifts to device nodes and, if it exceeds a certain critical charge, Qcrit, may flip the state of the memory cell. A lower Qcrit results in a higher SER. Alpha particles normally cause SBUs because they have lower energies, but they can cause MBUs in devices with low supply voltages.

High-energy cosmic particles react with the upper atmosphere of the Earth, and their collisions, modulated by solar flares and intergalactic cosmic rays, generate high-energy protons and neutrons. High-energy neutrons have energies of 10 to 800 MeV; in contrast, protons have energies greater than 30 MeV. High-energy neutrons have no charge; therefore, they do not coulombically interact with the semiconductor material, so their interaction with silicon differs from that of an alpha particle. For a high-energy neutron to cause a soft error, it must produce ionized particles by colliding with the silicon nucleus and undergoing impact ionization with the silicon nuclei. This collision can generate alpha particles and other heavier ions, thus producing electron-hole pairs but with higher energies than a typical alpha particle from mold compounds.

Neutrons are particularly troublesome, because they can penetrate most manmade construction; for example, a neutron can pass through five feet of concrete. The flux rate is geoposition-dependent and increases at higher altitudes due to a lower shielding effect of the atmosphere. In London, the effect is 1.2 times worse than at the equator. In Denver, with its high altitude, the effect is three times worse than at sea level in San Francisco. In an airplane, the effect can be 100 to 800 times worse than on the ground.

TABLE 1–COMPARISON OF SOFT-ERROR SOURCES

SER type	Source	Mechanism	Trend
Alpha	Thorium and uranium contamination in-mold compound, silicon, or lead bumps	2- to 9-MeV alpha particle creating electron-hole funnel traveling 25 microns in silicon	Exponential increase with scaling
Cosmic	Intergalactic sources modulated by solar flares	High-energy neutrons/protons (10 MeV to 1 GeV) colliding with silicon nuclei	Decrease in failures in time per megabit
Thermal neutron	Boron present in BPSG25-meV neutrons	Collision with B10 in BPSG	Highest, always dominates if present

Thermal neutrons are major contributors to soft failures and typically have energies of approximately 25 meV. The Boron-10 isotope that occurs in large quantity in BPSG (boron-phosphor-silicate-glass) dielectric layers easily captures these low-energy neutrons. Capturing a neutron results in a fission that produces lithium, an alpha particle, and a gamma ray, which may lead to potential bit-flips. Thermal neutrons are primarily an SEU issue only if BPSG is present; eliminating the use of B-10 isotopes effectively addresses the problem.

Poor system design is the final common source of SEUs. High-performance memory devices normally comprise SRAM cells, combinational logic, and latches. In high-performance communication-memory products, the area efficiency is usually low. Past academic research shows that combinational logic is less susceptible to soft errors than memory cells because of natural resistances set up by masking. However, these natural resistances could diminish as devices scale and the number of stages in the processor pipeline increases.

MEASURING SEUs

Accelerated SER measurement and system-level SER testing are two ways to measure a device's susceptibility to soft errors. Accelerated SER testing involves exposing the chip to various types of radiation, usually alpha or cosmic particles, at high intensities and other conditions that conform to JEDEC (Joint Electron Device Engineering Council) standards. Placing a thorium or uranium source on a decapped chip, measuring the total number of upsets within a certain time, and extrapolating the FITs per megabit can reveal the susceptibility of a device to alpha particles. Accelerated high-energy neutron (cosmic) testing is more complicated and is typically performed at research laboratories such as Los Alamos National Laboratory (Los Alamos, CA), where neutron sources are available. These two methods of accelerated data measurement are fair approximations of the FIT rate but typically overstate the actual failure rate. Using accelerated data can provide a good approximation to calculate the total time you need to perform a system SER measurement.

System-level SER testing is the evaluation of chip exposure to alpha and cosmic particles' SEUs by using a tester environment containing hundreds or thousands of chips and evaluating their failure rate under nominal conditions. A good way of filtering the effect of alpha particles and cosmic rays in a system could be to measure the data by placing the system several meters underground, where the effect of cosmic rays is negligible and then monitoring the system at high altitudes, where the effect of alpha particles is negligible. Testing could take as long as a year to increase the reliability of the results. System-level soft-error-rate testing is fairly expensive, so memory vendors do it on a per-technology rather than per-device basis to keep costs down.

System-level SER testing is cumulative of the alpha as well as cosmic SER, and the data depends greatly on the system's geographic location. Companies adjust SER FIT rates to the geographic location of New York at sea level to minimize discrepancies in measured values between companies; to maintain a common reference point of data between products; and to account for the cosmic flux rate's variation with latitude, altitude, geomagnetic, and shielding effects.

COMBATING SEUs

Approaches to reduce the SER include system-level changes (error-correction checking), process changes (buried layers, triple wells), circuit hardening (resistive feedback, higher capacitance at storage nodes, higher drive), and design hardening (redundancy). At the system level, designers can mitigate the increase in SER for SRAMs by using ECC (error-correction-code) detection and correction so that every addressable word of data stored in memory includes check information (**Figure 3**). The combination of data and check information is often called a check word.

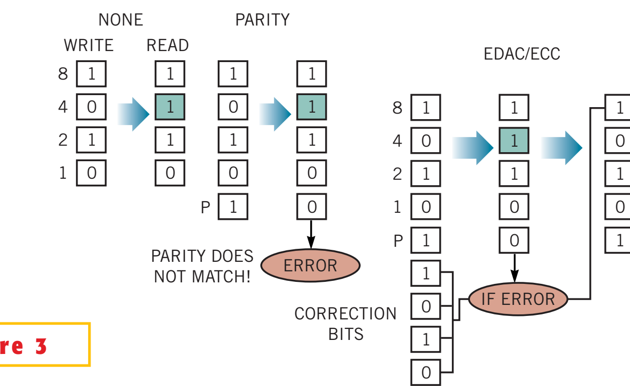

Figure 3 ECC detection and correction can mitigate the increase in SER for SRAMs.

Check information serves two purposes. First, when a check word is read from memory, the check information can help determine whether any of the data bits have changed. In ECC detection, the check information can help determine whether a single bit or more than one bit has changed. Second, if only a single bit has changed, ECC correction helps determine which bit changed and facilitate correcting the data by flipping the bit back to its complementary value.

An ECC-detection circuit detecting a change in one or more bits in a word of data is broadly categorized as an ECC error. You can further categorize such errors as functions of the number of bits in error. Currently available ECC circuits can correct single bits and report multibit errors. Designers can implement ECC detection and correction as hardware or software.

System designs can implement interleaving, wherein the bits in each word comprise cell addresses that are physically separate or interleaved on the memory device. This technique helps ensure that no upset of two nearest-neighbor memory cells resides in the same check word, which can make multicell cosmic-ray events into multiple ECC-correctable single-bit errors. Memory vendors are adopting interleaving in the memory design itself so that the memory topological bit map organizes memory bits pertaining to a byte from multiple blocks inside the memory array.

Bit flips that occur in memory go undetected until the affected word of data is read from memory and presented to the ECC-detection mechanism. Undetected, or "latent," errors. Strictly speaking, ECC correction applies only to the copy of the affected word of data. The

data as it resides in memory still contains the flipped bit. If this flipped bit in memory remains uncorrected, an exposure to another bit flipping in the same word of data can result in an uncorrectable error. It is important, therefore, that the system correct the flipped bit in memory.

Another way to increase the resistance of the device to SEUs is to increase the critical charge stored inside the memory cell to increase the critical threshold, Qcrit. Alternatively, using SOI (silicon-on-insulator) technology improves device reliability by reducing the charge depth. PMOS threshold voltage affects the recovery time of the cell, which indirectly affects the resistance to SEUs.

Designers can drift generated charge away from a soft error by using buried junctions (triple-well architecture) to increase the recombination far away from the active region. This process creates an opposite electric field with respect to the NMOS-depletion region and forces charges into the substrate. However, the triple well helps only when the radiation event occurs in the NMOS region.

At the process level, the use of purer mold compounds, such as those with alpha emission at the detection limit of 0.001 particles/cm/hr, lowers alpha emission. In advanced technologies, PSG can replace BPSG to remove effects from thermal neutrons.

As process technologies continue to shrink, the effect of soft errors on memory devices has gone from "insignificant" to an important consideration in system design. Depending on the application, SEUs strongly affect some systems and have no effect on others. However, SRAM is taking steps in both process development and product design to minimize the susceptibility to SEUs, thereby extending their use well beyond 90 nm. With proper care at the system- and product-design levels, SRAMs will remain a viable memory approach in numerous process generations.□

AUTHOR'S BIOGRAPHY

Ritesh Mastipuram is an applications engineer at the Memory Products Division of Cypress Semiconductor, where he is responsible for new-product definition, application support, application-system analysis, and board-level failure-analysis debugging. He has a master's degree in electrical engineering from Wright State University (Dayton, OH) and a bachelor's degree in engineering from Bangalore University (India).

Edwin C Wee is a product-design engineer at Cypress Semiconductor, where he is responsible for supporting new-product introduction and for evaluating product performance, reliability, and quality for volume production manufacturing. He has a master's degree in electrical engineering from Stanford University (Stanford, CA) and a bachlor's degree in electrical engineering and computer science from the University of California—Berkeley.

TALK TO US

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