

Soft Errors in Electronic Memory – A White Paper

INTRODUCTION

In the early days of computers, “glitches” were an accepted way of life. Since then, as computers have become more reliable (and more relied upon), glitches are no longer acceptable – yet they still occur. One of the most intractable sources of glitches has been the transient “bit-flip”, or soft memory error: a random event that corrupts the value stored in a memory cell without damaging the cell itself.

Soft errors* in electronic memory were first traced to alpha particle* emissions from chip packaging materials. Since then, memory manufacturers have eliminated most alpha particle sources from their materials, changed their designs to make them less susceptible (e.g., moved ball-grid bumps farther away from memory cells), and even added shielding (usually internal die coatings). Tests and standards have been developed to measure and improve the resistance of memory chips to alpha particles – but soft errors have not disappeared.

Further testing, mostly performed by avionics and space organizations, pinpointed a more pernicious source of soft errors: cosmic rays*. At ground level, cosmic radiation is about 95% neutrons and 5% protons. These particles can cause soft errors directly; they can also interact with atomic nuclei to produce troublesome short-range heavy ions. Cosmic rays cannot be eliminated at their source, and effective shielding would require meters of concrete or rock. To eliminate the soft memory errors that are induced by cosmic rays, memory manufacturers must either produce designs that can resist cosmic ray effects or else invent mechanisms to detect and correct the errors.

MEMORY TECHNOLOGY VS. ERROR RATES

Changes in technology have significant impacts on error rates, but not always in predictable ways. For example, DRAM* error rates were widely expected to increase as devices became smaller; instead, small-scale DRAMs demonstrate a much better error resistance. One reason for this is that their smaller size allows less charge collection [16][32]; another reason is that cell size has scaled faster than storage capacitance [16], so the capacitance ratio has actually increased [17][32]. On the other hand, SOI (Silicon on Insulator) technology was expected to resist errors [16][16][30]; however, it demonstrates an unexpected tendency toward large charge collection, which may dramatically increase error rates [9].

Some general trends have been observed:

- Alpha particle problems have been largely eliminated. [3][7][21][26]
- DRAM soft error rates have dropped. [2][10][16][17][19][21]
- SRAM* error rates are growing. [2][11][13][17][19][25][32]
- SRAM has become more susceptible to soft errors than DRAM. [5][6][13][19][25]
- Soft error rates for both DRAM and SRAM are expected to rise. [5][6][9][11][21][23][27]

Tests have identified several specific design factors which impact error rates:

- Increased complexity raises the error rate. [13][16][17][24][27]
- Higher-density (higher-capacity) chips are more likely to have errors. [3][11][13][17][19][21][27][32]
- Lower-voltage devices are more likely to have errors. [5][10][11][14][16][21][32]
- Higher speeds (lower latencies) contribute to higher error rates. [5][11][13][14][16][17]
- Lower cell capacitance (less stored charge) causes higher error rates. [5][11][13][17][27][30][31]

* Terms are defined in Appendix A.

SOFT ERRORS IN ELECTRONIC MEMORY – A WHITE PAPER

- Shorter bit-lines result in fewer errors. [5][17]
- Wafer thinning improves error tolerance (especially with backside contacts). [9]
- “Radiation hardening” can decrease error rates by several orders of magnitude [13][24][27], but these techniques cost more, reduce performance, use more power, and/or increase area [9][11][14][30].

RECENTLY REPORTED DATA ON SOFT ERROR RATES (SERS)

SER* modeling and simulation are highly complex and frustratingly inaccurate [2][14][16][32]. In addition, there are no standard testing or reporting mechanisms for cosmic ray induced SERs [3][5][19][22] and most manufacturers don’t reveal their test results [5][23], so comparisons and conclusions are difficult to come by. The measurements and estimates in the following table were drawn from a variety of documents. (For calculations, see Appendix B.)

Type of Memory	Reported SER	Error per bit-hour	FIT*/Mbit	Source
Goal for new Cypress products	200 FIT*	?	?	[11]
SRAM (quoted by vendors)	200 to 2,000 FIT	?	?	[7]
“typical”	1,000 FIT	?	?	[11]
DRAM at full speed	Few hundred to few thousand FIT	?	?	[19]
SRAMs at 0.25 micron* and below	10,000 to 100,000 FIT	?	?	[19]
Commercial CMOS* memory	>1E-5 to 1E-7 per bit-day*	>4E-7 – 4.2E-9	4 million – >400 million	[11] [27]
“some” 0.13-micron technologies	10,000 or 100,000 FIT/Mbit*	1E-11 – 1E-10	10,000 – 100,000	[11]
1Gbit* memory in 0.25µm*	One error per week	6E-12	6,000	[14]
4M SRAM	<1E-10 upset*/bit-day	<4.2E-12	<4,200	[24]
1 Gbit of DRAM (Nite Hawk)	2.3E-12 upset/bit-hour*	2.3E-12	2,300	[23]
SRAM and DRAM	1 – 2 E-12 upset/bit-hour	1 – 2E-12	1,000 – 2,000	[23]
~8.2 Gbits of SRAM (CRAY YMP-8)	1.3E-12 upset/bit-hour	1.3E-12	1,300	[23]
SRAM	1,000 FIT/Mbit	1E-12	1,000	[13]
256 MBytes*	One error per month	7E-13	700	[3][19]
160 Gbits of DRAM (Fermilab)	2.5 errors per day	7E-13	700	[23]
32 Gbits of DRAM (CRAY YMP-8)	6E-13 upset/bit-hour	6E-13	600	[23]
MoSys 1T-SRAM (no ECC*)	500 FIT/Mbit	5E-13	500	[17]
Micron estimate, 256 MBytes	2 – 4 errors per year	1.2 – 2.4E-13	120 - 240	[19]
“ultra-low” failure rates	50 to 100 FIT per Mbit	5E-14 – 1E-13	50 – 100	[4]

Judging from these reports, 1000 to 5000 FIT per Mbit seems to be a reasonable SER for modern memory devices. Obviously, wide variances exist. One reputable source [31] reported anomalously low rates; these results were not included in the table above, pending further research.

* Terms are defined in Appendix A.

OTHER ISSUES

Measuring Error Rates

True “life-testing” for SER is difficult and expensive; “accelerated” testing is cheaper and more common, but it does not translate directly to real-life performance [32].

Memory errors occur mostly during read/write activity, so the SER rises with memory speed [5] and with the intensity of memory use [17]; “memory cycling at 100 nanoseconds can give soft error rates 100 times that of memory idling in refresh mode (15 microseconds) [19].”

Error rates rise with altitude: SER is 5 times as high at 2600 feet as at sea level, and 10 times as high in Denver (5280 feet) as at sea level [26]. “SRAM tested at 10,000 feet above sea level will record SERs that are 14 times the rate tested at sea level [11].”

Implementing Error Checking and Correction (ECC)

The most commonly used system of error recovery, ECC, adds extra bits (check bits) to each data item. These bits are re-computed and compared whenever the data item is accessed. ECC is credited with a “90% correction rate [21]” or better:

“The use of ECC will decrease the soft FIT* rate to a value at least as small as the hard FIT rates... [7]”

Most ECC algorithms can correct single-bit errors and detect, but not correct, double-bit errors. This type of algorithm is called SEC/DED (Single Error Correct / Double Error Detect). The vast majority of errors are correctable by this method:

“For SRAMs, almost 100% of the errors were single-bit errors ... 94% of the errors in ... DRAMs were single-bit, and the remaining errors were adjacent-bit ... a 1-bit ECC circuit can be used for multi-bit errors since adjacent bits have different address data sets [10].”

ECC is usually implemented by adding an extra chip to each multi-chip memory module (to supply the extra bits) and incorporating ECC logic and circuitry at the system or board level. In embedded “system-on-a-chip” designs, both the circuitry and the extra bits must be implemented at the chip level. Either way, ECC imposes a cost in dollars, performance, and/or size. Several sources comment on the impact of incorporating ECC:

ECC memory solutions generally cost more: “10% to 20% [1]” or “more than 25% [17]”, or “considerably [3]”.

ECC memory solutions are larger, with a “20% die-area penalty typically associated with error checking and correction (ECC) [25].”

ECC memory solutions are generally slower – e.g., a performance reduction of “3-4% on PC133 CAS2 ECC SDRAM [1]”. For the very fastest SRAMs*, building in full-fledged error correction might degrade performance as much as 33% [5].

Other algorithms for error detection and correction (collectively known as EDAC) can correct multi-bit errors, but they are rarely implemented in memory. Another error correction technique, “scrubbing” [11][24][27], uses a background task to inspect memory periodically for single-bit errors. Because scrubbing does not wait for a memory access, it reduces the likelihood that single-bit errors will accumulate, and thus reduces the risk of uncorrectable (or undetectable) multi-bit errors.

* Terms are defined in Appendix A.

Soft Errors* and Logic

Soft error rates in processors and logic are also increasing as devices use smaller scales. Not only are these errors harder to track than memory errors, their rate is expected to increase as speeds rise: “the key parameters are gate oxide thickness, relative speed and density, and power supply voltage... [16]” “The logic cross-section is very close to the memory cross-section. ... In addition, the logic FIT* rate is expected to increase due to other phenomenon such as cross coupling, ground bounce and delay faults. Finally, the FIT rate in logic is likely to be comparable to the FIT rate in memory [8].” “The physical phenomenon is exactly the same: There is no reason why logic transistors are less impacted by transient errors [4].”

Hard Errors*

Quite aside from soft errors, particles with high energies can cause permanent damage to memory cells. These “hard” errors exhibit error rates that are strongly related to soft error rates [29], variously estimated at 2% of total errors [26] or “one or two orders of magnitude less than soft error rates – often in the range of 5 to 20 FIT [7]”. A one-bit hard error is correctable with ECC*, just as if it were a soft error; however, the error will recur each time the bad cell is used. As hard errors accumulate, they eventually render the memory device unusable. Recently, a very few state-of-the-art memory devices have incorporated new self-healing technologies to repair hard errors; these technologies are outside the scope of this paper.

Conclusions

Soft errors are a matter of increasing concern as memories get larger and memory technologies get smaller. Even using a relatively conservative error rate (500 FIT/Mbit), a system with 1 GByte of RAM can expect an error every two weeks; a hypothetical Terabyte system would experience a soft error every few minutes. Existing ECC technologies can greatly reduce the error rate, but they may have unacceptable tradeoffs in power, speed, price, or size.

Soft errors can be disastrous for systems with large memories, critical applications, or high altitude locations. Some type of error detection/correction is mandatory in these cases, in spite of the cost in price and/or performance.

Even with error correction, systems may be susceptible to multi-bit errors, hard errors, or errors in logic. These risks may need to be addressed with memory scrubbing, redundancy, self-healing mechanisms, shielding, new advances in hardware technologies, and/or new error correction methods.

* Terms are defined in Appendix A.

APPENDIX A – DEFINITIONS

- alpha particle** A helium nucleus (2 protons + 2 neutrons); commonly emitted by radioactive materials.
- bit-day** One memory bit active for one day.
- bit-hour** One memory bit active for one hour.
- CMOS** Complementary Metal Oxide Semiconductor: The fabrication technology currently used to build most electronic memories.
- cosmic ray** A high-energy particle originating in outer space.
- DRAM** Dynamic Random Access Memory: today's highest-density memory technology.
- E+n or E-n** Exponential notation for a power of ten; E+6 = million, E-3 = one thousandth
- EDAC** Error Detection And Correction: a generic term covering many schemes.
- FIT** Failures In Time: Errors per billion (10^9) hours of use. Usually reported as FIT per Mbit.
- Gb or Gbit** Gigabit: 1,073,741,824 bits (1024 Megabits), or approximately 10^9 bits.
- GB or GByte** Gigabyte: 1,073,741,842 bytes (8 Gigabits), or approximately 10^9 bytes.
- hard error** A permanent error caused by a physical defect in the system.
- Mb or Mbit** Megabit: 1,048,576 bits, or approximately 10^6 bits.
- MB or MByte** Megabyte: 1,048,576 bytes (8 megabits), or approximately 10^6 bytes.
- MBTF** Mean Time Between Failures
- micron** Micrometer; one millionth of a meter. Abbreviated as μm (sometimes just μ). Semiconductor manufacturing technologies are measured by the smallest features they can produce; 0.13 micron technology can build features about half as large as 0.25 micron technology.
- μm** Abbreviation for micrometer (see "micron" above).
- RAM** Random Access Memory: Any type of memory in which non-sequential locations can be read and written in any order.
- SEC/DED** Single Error Correct / Double Error Detect: any correction algorithm corrects single-bit errors and detects (but does not correct) double-bit errors.
- SER** Soft Error Rate
- SEU** Single Event Upset: NASA's term for a radiation-induced soft error.
- soft error** A transient error, which is not due to any defect in the physical system.
- SRAM** Static Random Access Memory: today's highest-speed memory technology.
- Tera-** Prefix meaning one trillion (10^{12}). One Terabyte equals 1,024 Gigabytes.
- upset** Another name for a soft error, frequently used in avionics.

SOFT ERRORS IN ELECTRONIC MEMORY – A WHITE PAPER

APPENDIX B – CALCULATIONS

1 week = 168 hrs; 1 upset/week = 6E-3 upset/hr

1 month = ~730 hrs; 1 upset/month = 1.4E-3 upset/hr

1 year = 8760 hrs; 1 upset/yr = 1.1E-4 upset/hr

1 Gbit = ~1E+9 bits; 1 error per Gbit-hour = ~1E-9 upset/bit-hour

256 MB = 2,048 Mbits; 1 error per 256 MB-hour = ~ 5E-10 upset/bit-hour

1 FIT/Mb = 1 upset per 10⁹ hrs per 10⁶ bits = 1E-15 upset/bit-hour

1 upset/bit hr. = 10¹⁵ upsets per 10⁹ hours per 10⁶ bits = 1E+15 FIT/Mb

To calculate system MTBF in hours: **$m = 1,000,000,000 / (f * s)$**

where: “m” = MTBF “f” = FIT/Mb rating “s” = size of system in Mb

Normal PC (512 MB = 4096 Mb) with normal DRAM (1000 FIT/Mb):

$$\text{MTBF} = 1\text{E}+9 / (1000 * 4096) = 244 \text{ hrs.}$$

APPENDIX C – REFERENCES

- [1] Anandtech FAQ, “What is ECC RAM?” 24 October 2001 <http://www.anandtech.com/guides/viewfaq.html?i=3>
- [2] Baumann (Texas Instruments), “Soft Error Characterization and Modeling Methodologies at TI” 2000 http://www.sematech.org/meetings/20001030/20_SER_Baumann.pdf
- [3] Cataldo, “IBM moves to protect DRAM from cosmic invaders” *EETimes*, 10 June 1998 <http://www.eetimes.com/news/98/1012news/ibm.html>
- [4] Cataldo, “MoSys, iRoC target IC error protection” *EETimes*, 6 February 2002 <http://www.eetimes.com/story/OEG20020206S0026>
- [5] Cataldo, “SRAM soft errors cause hard network problems” *EETimes*, 17 August 2001 <http://www.eetimes.com/story/OEG20010817S0073>
- [6] ChipCenter Networking Product Review, “DRAM-CAM – MOSAID’s 9 Mbit DRAM-Based Chip” 28 January 2002
- [7] Cisco Systems, “Increasing Network Availability” (Date unknown) <http://www.cisco.com/warp/public/779/largeent/learn/technologies/ina/IncreasingNetworkAvailability-WhitePaper.pdf>
- [8] Design And Reuse, “iRoC Releases Robust SPARC Test Report” 28 January 2002 <http://www.us.design-reuse.com/news/news65.html>
- [9] Dodd et al (Sandia Nat’l Labs), “Epi, Thinned, and SOI Substrates” December 2001 http://parts.jpl.nasa.gov/mrqw/mrqw_presentations/S4_dodd.ppt
- [10] Eto et al, “Impact of Neutron Flux on Soft Errors in MOS Memories” 15 February 1999 <http://rd49.web.cern.ch/RD49/Neutron.pdf>
- [11] Graham, “Soft errors a problem as SRAM geometries shrink” *ebn*, 28 Jan 2002 <http://www.ebnews.com/story/OEG20020128S0079>
- [12] Haddad (BAE Systems), “Advanced Microprocessor Technologies and Challenges for Space Applications” December 2001 http://parts.jpl.nasa.gov/mrqw/mrqw_presentations/Keynote1_haddad.pdf
- [13] Harling, “Embedded DRAM Has a Home in the Network Processing World” *Integrated System Design*, 3 August 2001 <http://www.eedesign.com/isd/OEG20010803S0026>

SOFT ERRORS IN ELECTRONIC MEMORY – A WHITE PAPER

- [14] Holbert (Arizona State U.), “Single Event Upsets” (Date unknown)
<http://www.eas.asu.edu/~holbert/eee460/see.html>
- [15] iRoC, “Why is the next generation of ICs more prone to transient errors?”
<http://www.iroctech.com/index.adml?l=111&search=transient&id=27941919401073327943148&r=60>
and:
“What is the current ratio for transient errors rates?”
<http://www.iroctech.com/index.adml?l=111&search=SER&id=27941919401073327750374&n=2&i=0&r=61>
- [16] Johnston (JPL), “Recent Work on Radiation Effects in Microelectronics at JPL” 27 September 2000
http://rd49.web.cern.ch/RD49/RD49News/Allan_Johnston.pdf
- [17] Johnston (JPL), “Scaling and Technology Issues for Soft Error Rates” October 2000
<http://nepp.nasa.gov/docuploads/40D7D6C9-D5AA-40FC-829DC2F6A71B02E9/Scal-00.pdf>
- [18] Jones (MoSys), “When Memories Forget” 2001 <http://www.mosys.com/news/dc21k.pdf>
- [19] Leung, Hsu, & Jones (MoSys), “The Ideal SoC Memory: 1T-SRAM™” (Date unknown)
<http://www.mosys.com/news/idsoc.pdf>
- [20] Lieberman of ASTCO, “ECC Whitepaper” (1998) <http://www.corsairdirect.com/ecc.html>
- [21] Locklear (Dell), “Chipkill Correct Memory Architecture” August 2000
<http://www.ece.umd.edu/courses/enee759h.S2003/references/chipkill.pdf>
- [22] Messer et al, “Susceptibility of Modern Systems and Software to Soft Errors” 7 March 2001
<http://www.hpl.hp.com/techreports/2001/HPL-2001-43.pdf>
- [23] Micron Technical Note, “DRAM Soft Error Rate Calculations” <http://download.micron.com/pdf/technotes/DT28.pdf>
- [24] Normand (Boeing), “Single Event Upset at Ground Level”
http://www.boeing.com/assocproducts/radiationlab/publications/SEU_at_Ground_Level.pdf
- [25] Rodgers et al (BAE Systems), “Advanced Memories for Space Applications” December 2001
http://parts.jpl.nasa.gov/mrqw/mrqw_presentations/S1_rodgers.pdf
- [26] Semiconductor Business News, “MoSys adds soft-error protection, correction” 28 Jan 2002
<http://www.siliconstrategies.com/article/showArticle.jhtml?articleId=10802390>
- [27] Smith (Computer & Information Science Dept., Linköpings Univ., Sweden), “RAM Reliability” April 1998
<http://www.ida.liu.se/~abdmo/SNDFI/docs/ram-soft.html>
- [28] Swift (NASA/JPL), “Space Radiation Effects, Part III: Recoverable SEU Effects” 2001
http://parts.jpl.nasa.gov/docs/Radcrs_Final.pdf
- [29] Takeuchi et al, “Origin and Characteristics of Alpha-Particle-Induced Permanent Junction Leakage” *IEEE Transactions on Electron Devices*, March 1990
- [30] Wall & Macdonald, editors (JPL) “The NASA ASIC Guide, Draft 0.6: Chapter 4” (Date unknown)
<http://nppp.jpl.nasa.gov/asic/Sect.3.4.html#A0>
- [31] Ziegler et al, “Cosmic Ray Soft Error Rates of 16-Mb DRAM Memory Chips” *IEEE Journal of Solid-State Circuits*, February 1998
- [32] Ziegler (U. S. Naval Acad.), “Review of Accelerated Testing of SRAMs” November 2000
<http://www.srim.org/SER/SERTrends.htm>