

2007 IEEE NSREC
Short Course



Section IV

Radiation Hardening at the System Level

Ray Ladbury
NASA Goddard Space Flight Center

Radiation Hardening at the System Level

Ray Ladbury
NASA Goddard Space Flight Center
Radiation Effects and Analysis Group

Table of Contents

1	INTRODUCTION: SYSTEMS AND SYSTEM HARDENING.....	3
1.1	SYSTEM-LEVEL MITIGATION: WHY NOW?	6
1.2	OVERVIEW OF THE SYSTEM-LEVEL HARDENING PROCESS	8
2	WHEN DO WE MITIGATE?	13
2.1	REQUIREMENTS GENERATION AND FLOWDOWN	13
2.2	RADIATION TESTING AND ANALYSIS: ERRORS AND INFERENCE	15
2.2.1	<i>Errors and inference for TID</i>	15
2.2.2	<i>Errors and Inference for Displacement Damage</i>	19
2.2.3	<i>Errors and Inference for SEE</i>	20
2.3	FAILURE COST AND SEVERITY	22
2.4	FAILURE RISK AND CRITICALITY	24
3	RADIATION EFFECTS AND THEIR MITIGATION.....	28
3.1	RADIATION EFFECTS AND THEIR CONSEQUENCES.....	30
3.2	DESTRUCTIVE SINGLE-EVENT EFFECTS	30
3.2.1	<i>Single-event latchup (SEL)</i>	31
3.2.2	<i>Single-Event Burnout (SEB)</i>	34
3.2.3	<i>Single-Event Gate Rupture (SEGR)</i>	35
3.2.4	<i>Other Destructive SEE</i>	36
3.3	NONDESTRUCTIVE SEE.....	38
3.3.1	<i>Single-Event Transients (SET)</i>	38
3.3.2	<i>Single-event Upset (SEU)</i>	40
3.3.3	<i>Multi-Cell and Multi-Bit Upsets (MCU and MBU)</i>	40
3.3.4	<i>Single-Event Functional Interrupt (SEFI)</i>	42
3.4	DEGRADATION MECHANISMS	43
3.4.1	<i>Total-Ionizing Dose (TID)</i>	44
3.4.2	<i>Displacement Damage (DD)</i>	45

4	MITIGATION STRATEGIES	48
4.1	MITIGATION FOR DESTRUCTIVE SEE	50
4.1.1	<i>Threat Reduction</i>	50
4.1.2	<i>Event Detection and Protection</i>	51
4.1.3	<i>Cold Sparing</i>	52
4.2	HARDENING TECHNIQUES FOR NONDESTRUCTIVE SEE.....	53
4.2.1	<i>Hardening for SEFIs</i>	53
4.2.2	<i>Mitigation for SEUs, MBUs and other data loss mechanisms</i>	54
4.2.3	<i>Mitigating Single-Event Transients</i>	60
4.3	MITIGATION FOR DEGRADATION MECHANISMS.....	62
4.4	VALIDATION	65
5	PUTTING HARDENING INTO PRACTICE: AN EXAMPLE	66
5.1	DESTRUCTIVE SEL AND LATENT DAMAGE.....	67
5.2	TID AND RLAT	68
5.3	STUCK BITS	68
5.4	INFORMATION LOSS: NONDESTRUCTIVE SEL, SEFI, MBU AND SEU	68
5.5	SUMMARY	71
6	WHEN MITIGATION BREAKS DOWN.....	73
6.1	BREAKDOWN OF ASSUMPTIONS.....	74
6.2	TEST FIDELITY	76
6.3	ENSURING REPRESENTATIVE TEST SAMPLES	77
6.4	SYNERGISTIC EFFECTS	79
6.5	COST-EFFECTIVE RISK REDUCTION	79
6.6	SUMMARY	80
7	CONCLUSION: CASE STUDIES	82
7.1	SOLAR DYNAMICS OBSERVATORY (SDO).....	82
7.2	EXPLORATION ROVERS: SPIRIT, OPPORTUNITY AND BEYOND.....	84
7.3	SUMMARY	86
8	REFERENCES.....	88

1 INTRODUCTION: SYSTEMS AND SYSTEM HARDENING

When I sat down to put this section of the course together, I discovered that no fewer than 15 past short-course sessions—dating all the way back to the first short course in 1980—have dealt with system-level approaches to radiation hardening and radiation hardness assurance (RHA). (See Table I). This remarkable—and frankly rather daunting fact—raises a number of questions, such as: “Why such an emphasis on the system level?” and “What is there left to cover?”

Part of the answer to the first of these questions is that the attention paid to the system perspective really does reflect the breadth of the topic. It has never been possible to fulfill all mission requirements solely using radiation hardened components—whether those components were vacuum tubes or the latest Radiation-Hardened-by-Design (RHBD) ASICs and Field-Programmable Gate Array (FPGA) designs. If the mission has a challenging radiation environment and requirements, and there is no radiation-hardened or RHBD solution, system-level hardening provides one last line of defense whereby a critical part with marginal radiation performance may still meet its system requirements. (See Figure 1-1.)

Another reason why system-level strategies are popular in short-courses is because system-level solutions tend to be rather involved and tailored to a particular system. As such, they often do not lend themselves to treatment in RHA literature for a general audience. System-level hardening becomes necessary only when a critical component has failed to satisfy requirements in traditional piece-part hardness assurance [1]. The process of mitigating the failure of a critical component can be complicated and expensive—both in terms of resources and system performance, giving rise to the conflicting demands of system-level hardening: The designer tells the radiation engineer, “Just make it work,” while the program manager says, “Just make it cheap.” Faced with such conflicting demands, the most efficient solution often winds up being application specific, providing just enough mitigation to meet requirements, and application specific solutions often do not lend themselves to publication in radiation-effects literature.

Table I. NSREC Short Courses Dealing with System-Level Hardening

Year	Author	Title	Emphasis
2004	Lum	Hardness Assurance for Space Systems	Overview of system hardness assurance from requirements to environments, device effects, mitigation and production
2002	Poivey	Radiation Hardness Assurance for Space Systems	Overview of the usual hardness assurance program emphasizing requirements and parts qualification
1999	Heidergott	System Level Mitigation Strategies	Satellite constellations as a system; environmental tradeoffs vs. constellation complexity; system-level consequences; commercial part qualification issues; mitigations and fault-tolerant systems
1998	LaBel & Cohn	Applying State of the art Commercial and Emerging Technologies to Space Systems	Roles and Limitations of state-of-the-art technologies to space systems; technology selection
1998	Kinnison	Achieving Reliable, Affordable Systems	Characterization and mitigation of radiation effects
1994	Haddad & Scott	Adapting Commercial Electronics to the Naturally Occurring Radiation Environment	Commercial part qualification strategies; mitigation and verification strategies; RHBD on commercial designs.
1994	Normand	Single-Event Effects in Systems Using Commercial Electronics in Harsh Environments	Environment, testing and statistical treatment of SEE reflecting impact of application conditions.
1989	Peden, Josephson, Ritter, Rudie, Swant, Abare and Price	Survivability and Hardening of Space Systems (all 5 sections of the course deal with system level hardening)	strategic and natural environments with historical perspectives, environmental characteristics and specifics sections devoted to history (Josephson), general techniques (Ritter), power (Rudie), Communications (Swant), Signal Data Processing (Abare) and Payload (Price)
1987	Robinson	Packaging, Testing, and Hardness Assurance	Outline of RHA process, especially requirements, procurement, testing, analysis and oversight.
1986	Raymond	Subsystem Electronics Testing	Approaches to testing and verification at the subsystem level; geared to strategic environments, but good insights for subsystem-level proton or TID tests.
1985	Sievers & Gilley	Fault-Tolerant Computing: An Architectural Approach to Tolerating Radiation Effects	Overview of fault-tolerant computing and soft-error mitigation techniques
1984	Tasca, O'Donnell & Klisch	SGEMP Hardening of Spacecraft	Requirements, mitigation and verification of EMP hardening in electronic systems
1982	Allen	System Aspects of Radiation Hardening	Hardening process from requirements to design to production--strategic
1981	Messenger	Hardened Systems Development and Hardness Assurance	Hardening systems for strategic environments
1980	Johnston	Radiation Effects on Components and Circuit Hardening	Device and circuit level effects, along with mitigations and component selection
1980	Halpin	System Effects and Systems Validation	Weapons effects and system-level validation
1980	Rudie	System Hardening Methodology	Nuclear hardness requirements, guidelines and planning

Similarly, journals, texts and handbooks on system engineering often give radiation issues a light treatment if they threat them at all. The short course offers a forum for in-depth consideration of system-level hardening and for capturing the collected wisdom of the community when it comes to surmounting radiation threats.

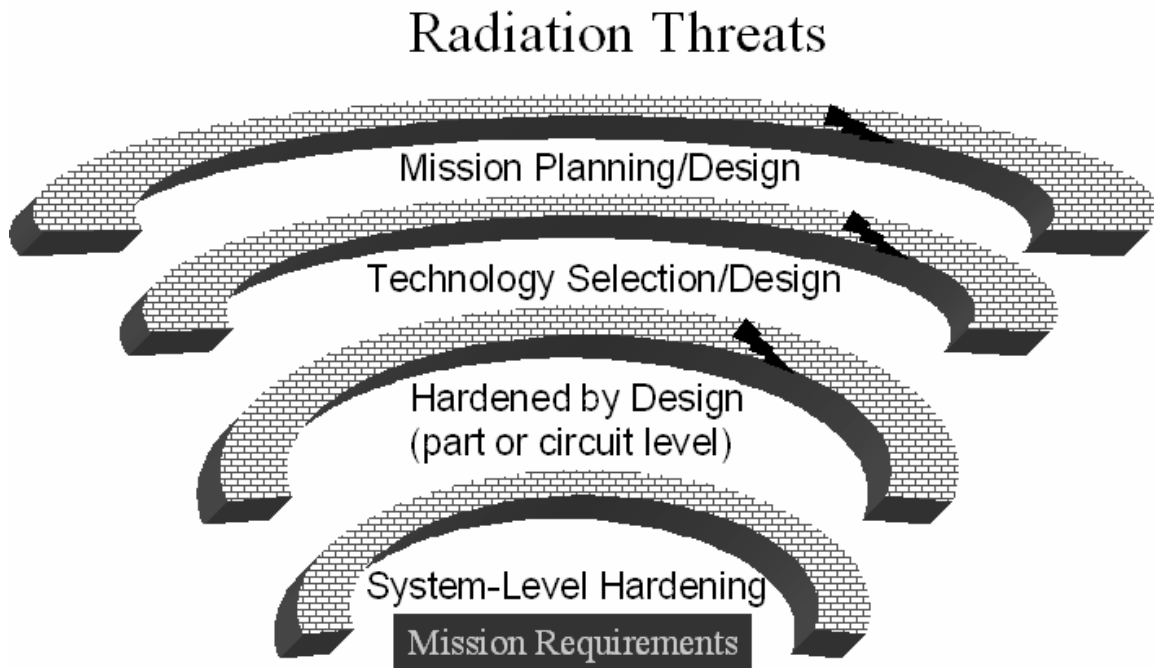


Figure 1-1. Radiation hardening can be viewed as a multi-tiered defense against radiation threat. The first line of defense uses mission planning and design to avoid unnecessary radiation risks. The second tier selects technology and designs the system so that components meet both mission and radiation requirements. Where there is no pre-existing radiation hardened solution, hardness can be designed into the system at the part or circuit level. When a critical component appears questionable despite these measures, radiation hardening at the system level can help the component meet its system-level requirements.

As to what remains to be covered, the answer is: plenty. Many of the basic strategies of system-level hardening predate the space age, having been borrowed from communications and fault-tolerant computing. Still, although the basic strategies have not changed drastically since the first short course in 1980 (see the sessions by Alan Johnston, Joseph Halpin and Norman Rudie), every change in technology requires adjustments in how these strategies are applied. Each new project poses its own challenges, as demands for increased performance bend and then break formerly immutable rules (e.g. bans on commercial parts or parts susceptible to single-event

latchup (SEL)). Moreover, the subject of system-level hardening is sufficiently rich and can be looked at from so many angles that it could probably become a yearly segment with little overlap year to year. Given the breadth of the topic we are grateful for the efforts of past presenters, which allow us to concentrate on some aspects of system-level hardening while directing the interested reader to past short courses for a more detailed treatment.

1.1 System-Level Mitigation: Why Now?

In the vast majority of applications, the normal process of piece-part hardness assurance is sufficient to demonstrate that a part will meet its application radiation requirements. It is only when the piece-part process indicates that the part's radiation hardness is doubtful and when there is no good substitute part that system-level hardening may be necessary. In some cases, the mitigation required can be as simple as adding spot shielding to reduce TID or displacement damage dose (DDD), or adding capacitive filtering against single-event transients (SET). In other cases, it may involve an analysis to provide guidance to designers so that the system can accommodate parametric degradation or errors that are likely to occur. For very important enabling technologies, hardening may involve designing the entire system around the radiation vulnerabilities of a critical component. However, system-level hardening can be very expensive—in terms of project resources and also in terms of performance. As a result, most projects will not be able to afford more than a few such components, and the enhanced capability provided by the component will have to be worth the effort.

Usually, the components that provide such enhancements are commercial devices or other cutting (some would say bleeding) edge devices. Examples where commercial technologies provide sufficient performance to be tempting include:

- 1) Memories: The largest available single-die, radiation-hardened, volatile memories are the family of 16 Mbit SRAM from Aeroflex [2], while a variety of vendors offer 1Gbit DDR and DDR2 SDRAMs—a factor of >60 more density with comparable savings in weight, power and board area. For nonvolatile memory, the potential benefits are even greater, with FLASH memories having up to 4 Gbit of storage and radiation options for the foreseeable future being 4 Mbit or less. In terms of radiation performance, the price for this savings is a risk of destructive single-event latchup (SEL) and other destructive failure modes, a high rate of single-event functional interrupt (SEFI), the possibility of part-to-part and lot-to-lot variability in both TID and

single-event effect (SEE) performance and the headaches of a short product lifecycle.

- 2) Processors The fastest radiation-hardened microprocessors execute a few hundred MIPS (for lack of a better accepted metric), while their commercial counterparts have execution speeds in the range of tens of thousands of MIPS. The main problems with commercial processors are that they tend to be susceptible to SEFI and for bulk CMOS, SEL. However, some missions and vendors have begun to look at system-level hardening approaches using commercial processors, in part taking advantage of the fact that some devices are manufactured in Silicon-on-Insulator (SOI) technology. [3], [4], [5]
- 3) Data Converters: Analog-to-Digital (A-to-D) and Digital-to-Analog (D-to-A) converters (ADC and DAC, respectively) are another area where radiation-hardened products lag the commercial sector. The highest resolution and highest speed radiation-hardened ADCs have 12-bit resolution and operate at 20 megasamples per second (MSPS), [6] while some commercial parts achieve 24-bit resolution, with 16-bit resolution being common. [7] Other commercial parts operate at 1 gigasample per second (GSPS) with 8-bit resolution. [8]. While most space applications do not require 16-bit resolution or GSPS speed, those that need it really need it, and have little alternative but to pursue a RHBD ASIC solution [9] or use a commercial offering. ASIC solutions are expensive, while commercial converters carry the risks of SEL, SEFI, TID degradation and high SET susceptibility.
- 4) Field-Programmable Gate Arrays: Radiation-hardened FPGAs have a long history of use in satellites, and have done a reasonable job pacing their commercial counterparts in terms of gate counts. The temptation of commercial SRAM-based FPGAs over even the largest radiation tolerant devices is that they can be reprogrammed, perhaps decreasing development costs and increasing mission flexibility. The disadvantage of course is that they can be reprogrammed randomly by SEE. In addition, however, powerful proprietary (so-called intellectual property, or IP) cores have been integrated with the basic FPGA architecture—e.g. Processor cores in the Xilinx Virtex IV FX-60. [10], [11]
- 5) Optoelectronics, Sensors and Detectors: At the boundary where electronics meets materials science, band-gap engineered semiconductors and other detector materials play a crucial role for some missions. Their radiation response however is not always well understood. [12], [13], [14] However, even relatively mundane imagers, light sources can pose risks from both single-event effects and degradation mechanisms. The most advanced parts available in this category are either commercial or custom parts. For the former, no thought is usually given to radiation performance, while for the latter, the radiation performance may not be known in advance.

In many cases, it will be difficult to develop a system-level hardening solution without compromising the very advantages that make the component or technology desirable in the first place, and it may be better to pursue an RHBD approach like those

discussed previously by Mike Alles. However, for some components and some applications a hardened ASIC, FPGA or even circuit approach will not be practical, and there will be no solution but to tailor the system so that the part can meet its requirements. Indeed, system-level hardening can be quite effective, in some cases yielding a system harder than would be possible fabricated purely from radiation-hardened components with no system-level hardening.

1.2 Overview of the system-level hardening process

Since system-level radiation hardening involves working closely with both system engineers and designers to achieve requirements, it is useful to understand the system-engineering perspective. The NASA Systems Engineering handbook defines a system as “a set of interrelated components that interact with one another in an organized fashion toward a common purpose” [15]. The common purpose of the system is captured in the mission requirements. In a space environment, a system faces a variety of radiation threats that could disrupt the organized interaction of the components and prevent the system from fulfilling its purpose. To prevent this, we must first carefully select components that meet their application requirements. For those components that are in danger of failing to meet requirements, we must tailor their interactions with the environment and within the system so that radiation threats are mitigated.

The system-engineering approach to dealing with hazards can be broken down into a deceptively simple three-step process [15] whose steps are:

- 1) Identify the threat
- 2) Evaluate the threat.
- 3) Mitigate the threat.

Achieving adequate reliability may require several iterations of steps 2 and 3, along with evaluation of the unintended consequences of any mitigation. Moreover, steps 2 and 3 are interdependent: A cursory evaluation of the threat may indicate that the application requires extensive mitigation that would be seen unnecessary given a thorough evaluation. Likewise, if our system incorporates mitigation from the beginning, we may be able to get by with a relatively simple validation test, rather than an expensive, extended threat evaluation. Since there may be no unique solution to this interdependent

problem, selection will usually be governed by the need to meet application requirements at minimum cost.

The NASA approach to radiation hardness assurance (see Figure 1-2) is compatible with this system approach. It begins by assessing the radiation threats and their severity for the mission environment. It then examines how the technologies needed for the mission will be affected by these radiation threats. Finally, the radiation engineer, designers, system engineers and other affected parties cooperate to develop mitigation strategies to ensure that the radiation threats do not impact mission requirements.[16] NASA's RHA strategy explicitly emphasizes the iterative and cooperative nature of the process—especially when it comes to mitigating undesirable radiation performance. For system-level hardening, we can assume that we already understand the radiation hazard and that a critical component is at risk of failing its radiation requirements. This means that we will be spending most of our time in the shaded boxes—that is, the boxes titled “Evaluate Device Usage,” “Engineer with Designers,” and “Iterate As Needed.”

As an example, in the TID hardness assurance methodology of MIL-STD 814 (see Figure 1-3), we consider system-level mitigation only when a part fails requirements and winds up in the box labeled: “Device not Acceptable Corrective Action Required”. Under normal circumstances, we might discard the marginal part and substitute one that is more likely to succeed.

However, sometimes a part is mission critical and there is no appropriate substitute, or if radiation analysis gets a late start, a problem part may already be integrated to the design and replacement would involve a very expensive and disruptive redesign effort. In such circumstances, the mission can:

- 1) Accept the risk of failure and use the part as is
- 2) Refine the testing and/or analysis to see if the part is likely to succeed under more realistic conditions.
- 3) Mitigate the risk of failure—in effect changing the requirements the part must meet while enabling the system to meet its requirements.

Testing, analysis and mitigation are time-consuming and expensive—in terms of project resources, schedule and—for mitigation—system performance. Even a few such parts can quickly break the bank for most projects.

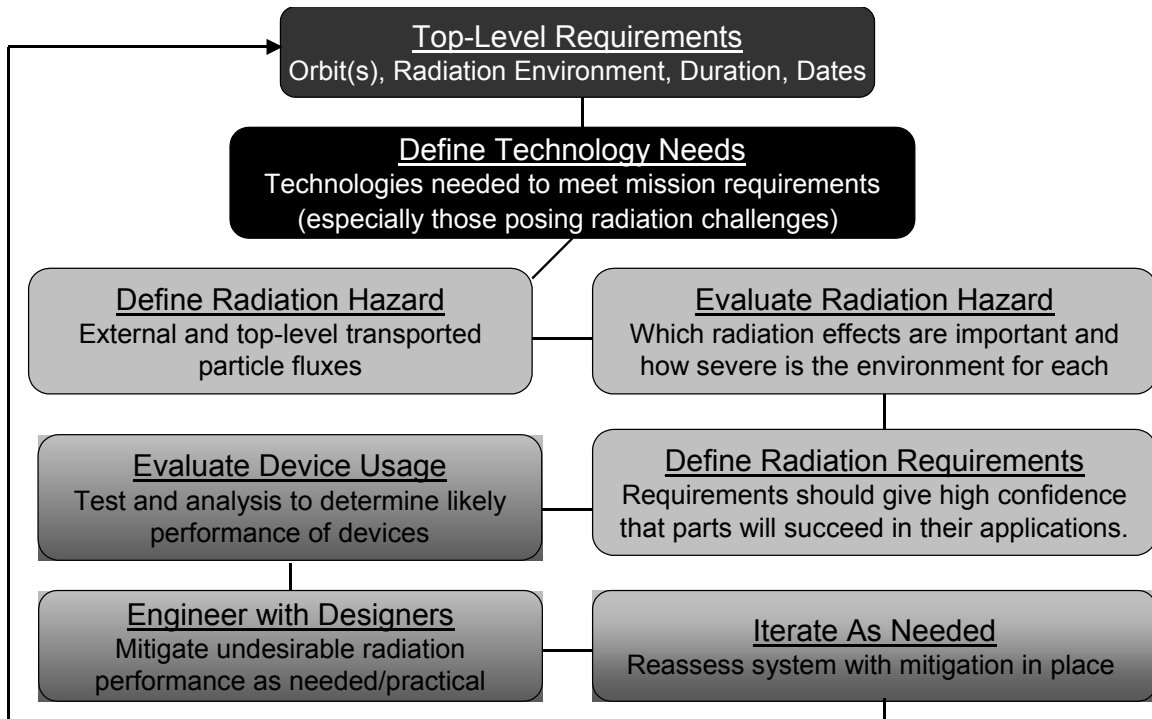


Figure 1-2. NASA’s approach to radiation hardness assurance (gray boxes) once top-level mission requirements and the critical technologies needed to meet them have been defined (black boxes). System-level hardening activities occur mainly in the shaded boxes. (Adapted from reference 16.)

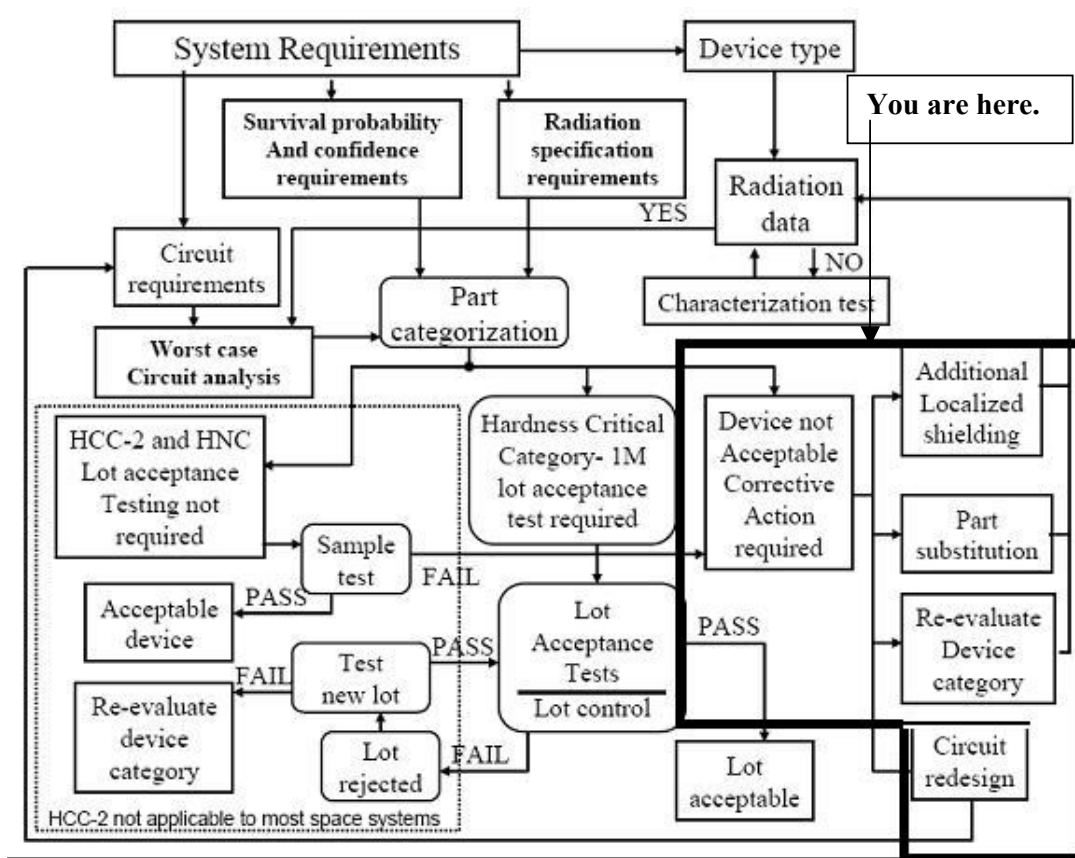


Figure 1-3. Most RHA issues can be resolved using Piece-Part Hardness Assurance Methodology outlined in figure 2 of MIL-HDBK 814[17] and adapted here. When a part falls into the Unacceptable category (block-bordered region) system-level hardening may be needed.

NASA’s RHA approach also explicitly emphasizes development of radiation requirements. Top-level mission requirements may be difficult to apply directly to radiation threats, and so may require interpretation to guide designers, parts engineers and system engineers. Clear, verifiable and relevant radiation requirements are essential to determining whether system-level hardening is needed. We will deal with radiation-requirement generation in section 2 below.

The application environment—whether it is the orbit and mission duration for a satellite system or the near-interaction region for a particle physics experiment—determines the radiation environment. Performance requirements determine the technologies required for the mission. Knowing the technologies and the environment, we can identify the radiation threats to which the mission is susceptible. Indeed, many

severe threats can be avoided by screening out parts that pose *unnecessary* risk so that system-hardening efforts focus on the challenges that really lead to improved performance in achieving mission goals. References [1], [18], and Poivey [19] discuss the role of parts selection and screening in hardness assurance.

Most hardening is done by selecting parts that meet radiation requirements, by developing designs for ASIC or FPGA using hardening techniques, or by circuit designs to harden, for example, an analog system—solutions discussed in the sections of this short courses by Mike Alles, Fernanda Kastensmidt and Ben Blalock. However, when these techniques are not practical or fall short of the hardening required by the system, system-level hardening is a last line of defense to ensure system requirements are met..

Because the designation of the part as unacceptable is usually determined by testing or analysis at the piece-part level, piece-part hardness assurance serves as a departure point for our discussion of system-level hardening. A designation of “unacceptable” at the piece-part level is not necessarily a prediction of failure, but may instead indicate that we have insufficient information to predict with confidence that the part would fulfill its performance requirements. Since system-level mitigation can be costly, additional testing or analysis may be advisable if there is a reasonable chance that it will show that mitigation is not needed.

If the refined testing and analysis still show that mitigation is needed, then one must select a mitigation strategy that will meet requirements without bankrupting the project. Details of mitigation schemes have been covered in several short courses, including [20] for general discussions, and [21] and [22] for fault tolerance. In some cases, the difficulty may be selecting the best among several viable strategies. In other cases, there may be no viable strategy consistent with available resources, and the project will need to decide whether to relax requirements or make more resources available.

2 WHEN DO WE MITIGATE?

Here we elaborate on a few concepts that will be useful in our discussions of system-level hardening. We first consider radiation requirements, since these define goals for hardening and the radiation environment. Next, since a part's acceptability is usually determined by radiation testing, we look at the uncertainties involved in radiation testing. Then, we look at the failure cost or severity, which scales the level of effort to be directed toward mitigating the radiation threat. Finally, we define radiation risk and related concepts—ways of combining the severity of a failure's consequences with the probability it will actually occur.

2.1 Requirements generation and flowdown

As mentioned above, requirements drive the cost and difficulty of any hardening effort. Clear, relevant, and verifiable requirements are essential to development of cost-effective systems. Unfortunately, top-level requirements may be too vague to be verified without significant elaboration at lower levels. For instance, take the example of the Solar Dynamics Observatory (SDO), which is scheduled to begin a 5 year mission observing the Sun in 2009 (see Appendix in section 7). SDO's top-level survivability requirement for the Solar Dynamics Observatory might state:

The satellite shall operate for at least 5 years in a geostationary orbit.

While this requirement is acceptable at the top level, it is not clear what it implies for the radiation performance of components and systems on the spacecraft, and it cannot be verified until the spacecraft has in fact completed its five-year primary mission in the intended orbit. As such, one must derive from the top-level requirements, radiation requirements that support these requirements and that are verifiable in terms of radiation test data and analysis. At the level of the Mission Assurance Requirements (MAR) or the Parts Assurance Requirements (PAR), one derived requirement might state:

Builders shall use only parts that are immune to destructive single-event effects (SEE).

This gives clear direction in the selection of components for the mission. However, given that this requirement must be verified with test data, and that test data always have systematic and statistical uncertainties (e.g. part-to-part variation, Poisson errors in SEE error counts, etc.), we need to construct a verification requirement that gives us an operational definition of immunity to destructive SEE. For example:

For the purposes of part qualification, immunity to single-event effects shall be demonstrated if all of at least two samples of the part do not exhibit the effect when each is exposed to a fluence of 10^7 ions with an equivalent LET (where valid) of at least 75 MeVcm²/mg.

Note that the process moves from a very general requirement with clear relevance to achievement of mission goals via several steps to a verification requirement that can be clearly interpreted in terms of test data. At each step, there is an interpretation that must be justified as relevant. At the MAR/PAR level, the mission duration/survivability requirement is interpreted as a ban on the use of parts susceptible to destructive SEE. This clearly supports the intent of the top-level requirement. Although it could be challenged as too stringent for a 5-year mission, the fact that destructive SEE can occur any time makes a ban on their use prudent. The verification requirement interprets the term “immunity” in a way that can be demonstrated easily with SEE test data. Both the required fluence and minimum LET requirement could be challenged, but are defensible in terms of historic testing efforts for geostationary missions. Thus, the verification requirement, while stringent, is clear, verifiable and relevant to top-level requirements due to its traceability to top-level mission requirements.

Similarly, for degradation mechanisms we may derive from the mission duration requirement above a requirement in the MAR:

All components shall perform their required functions with no system-level degradation after exposure to two times the mission total ionizing dose (TID).

Verification (and therefore any decision about the need for mitigation) relies on interpretation of test data—which we discuss next.

2.2 Radiation Testing and Analysis: Errors and Inference

The complicated nature of radiation testing justifies a brief review of the types of errors that occur in different radiation tests, how these errors are handled and their implications for inferences about likely component performance. The radiation effects of concern here can be divided into two categories:

- 1) Single-event effects (SEE) result from the prompt effects of a strike by a single ionizing particle in a sensitive portion of the component.
- 2) Cumulative effects such as displacement damage (DD) or damage due to total ionizing dose (TID) result from the gradual accumulation of damage by non-ionizing or ionizing radiation, respectively.

Our goal in radiation testing is to bound the component's susceptibility to the effect so that we can determine whether the effect is likely to impact achievement of requirements. The bound will be a function of the mean radiation performance observed in the test and the errors on that result. As such, we need to understand the systematic and random errors that occur in radiation testing. We look first at TID and displacement damage, and then at SEE.

2.2.1 Errors and inference for TID

Because, TID hardness usually varies more than SEE from part-to-part and wafer diffusion lot to wafer diffusion lot and because lot qualification is done by destructive testing on a sample drawn from the lot, errors on TID results are normally dominated by sampling errors. The most general scheme for dealing with such errors—binomial sampling—makes no assumptions about the failure distribution, but requires sample sizes that are uneconomically large. Instead of taking this general approach, most TID methodologies assume that parts from a single wafer diffusion lot will have similar TID performance [17] and so the failure distribution will be well behaved—that is, that the probability of having parts perform much worse than the mean performance will be small. This means that by increasing radiation design margin (RDM)—the ratio of the mean failure dose to that seen in the application—we diminish the probability the parts will fail to acceptably small levels. Unfortunately, not all failure distributions are well behaved (see figure 2-1), and if we merely assume a distribution is well behaved without requiring supporting evidence, we risk introducing systematic errors into the analysis. Still, this scheme usually works quite well.

TID data can be analyzed in two complementary ways to provide guidance to designers. The first way looks at the distribution of dose levels where different parts in the test sample meet some parametric or functional failure criterion. This type of distribution is useful for inferring what percentage of parts will fail at a given dose and for setting design margins. The second way of analyzing data looks at the distribution of parametric values or shifts for different parts in the test sample at a particular dose. This type of distribution gives designers information about how much degradation to design for in the system if we want more than a certain percentage of parts (say 99%) to be successful in that application. Both types of distributions are useful for system-level hardening.





Figure 2-1. Distribution of Ibias for 11 lots of Analog Devices OP484 op amps shows considerable variation lot-to-lot and even within lots. Two of the lots even show evidence of bimodality, with most parts in one lot and one part (extra-large “plus” sign) straying into the other mode. This suggests that using design margin to increase hardness assurance may be a risky technique for these parts. (Adapted from reference 23.)

If we assume a particular form for the distribution (e.g. normal or more often lognormal) we can use our test sample to infer the parameters of the distribution. However, the smaller the test sample size, the more likely it is that our inferred


distribution will be in error. For instance, figure 2-2 shows the distributions of how much the sample mean and standard deviation differ from the actual normal population mean and standard deviation as a function of sample size.

One solution to this dilemma is to use the one-sided tolerance limits, symbolized by K_{TL} to achieve a bounding estimate of the allowable dose or allowed degradation that guarantees success at the desired probability. In general the K_{TL} values are a function of the sample size n , the desired success probability P_s and the desired confidence level, CL—a measure of the assurance we require that our failure distribution will not be too optimistic. Figure 2-3 shows the K_{TL} values as a function of sample size for several desired success probabilities for the 90% CL. Values of K_{TL} are available in MIL-STD 814, or they can be calculated using Monte Carlo techniques. The advantage of the latter technique is that it can be used for any distribution form. The shape of the curves in Figure 2-3 is typical of how random errors often behave as data are added: When data are scarce, adding data has a large effect. However, for still more data, the curve approaches an asymptotic value.

As mentioned above, use of design margin for ill behaved failure distributions introduces systematic error into the analysis. While such systematic errors are not common, other types of systematic errors can affect TID testing. TID degradation can be quite application dependent. The fact that degradation can vary not just quantitatively, but qualitatively with dose rate—so-called enhanced low dose rate sensitivity (ELDRS) is well known. However bias conditions, temperature and other operating conditions can also bias experimental results. While it is impossible to have 100% fidelity to the application, the greater the fidelity of test conditions, the less the likelihood of systematic errors in the data.



(a)



(b)

Figure2-2. Distributions of the sample mean a) and sample standard deviation b).

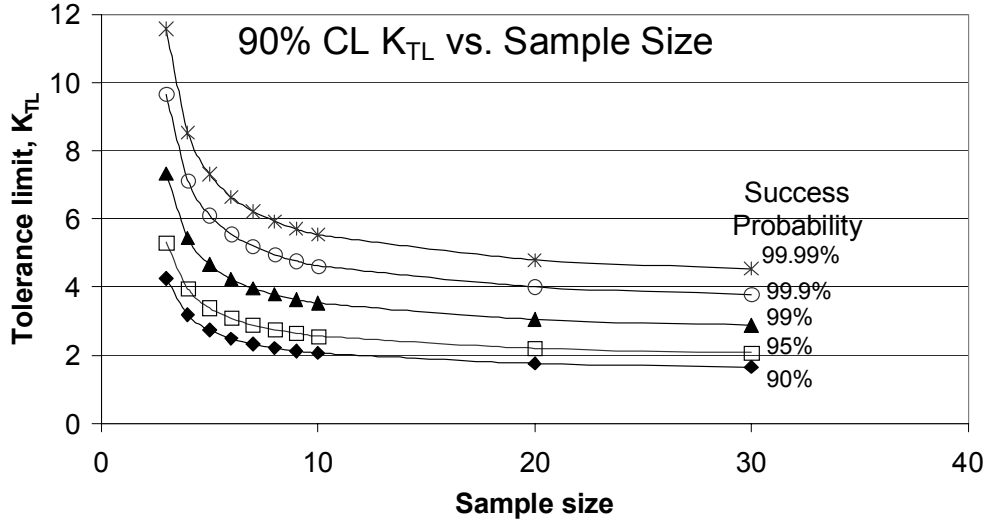


Figure 2-3. 90% CL one-sided tolerance limits, K_{TL} vs. sample size for various P_s .

In assessing whether mitigation is needed for a particular application, it is important to consider both random and systematic errors on TID data. In some cases, by testing more parts, one can reduce sampling errors sufficiently that mitigation may not be needed. In other cases, one may discover that data suggesting a part will succeed in an application may be systematically flawed, and it may be easier and more economical to mitigate TID degradation—e.g. by spot shielding—than to correct the error.

TID mitigation decisions can also be influenced by errors and biases in the analyses that estimate the dose a part is exposed to or circuit analyses that estimate the effect of a degradation on the application. Top-level TID estimates are often based on equivalent spherical shielding distributions, so TID estimates can be lowered significantly by more detailed sector analyses or Monte Carlo programs such as NOVICE.[24] On the other hand, if the spacecraft model—usually a sophisticated CAD file—given as input to a radiation transport program is flawed, the estimated doses will be also.

2.2.2 Errors and Inference for Displacement Damage

Although displacement damage performance generally varies less from lot to lot than does TID performance, the same statistical considerations and analyses apply. Moreover, for most part types, test results are also less application dependent. One exception is for optoelectronic devices, where annealing of damage is accelerated and increased in amphoterically doped LEDs by application of a forward voltage [25]. As with TID, estimates of displacement damage dose can be reduced significantly using more

sophisticated radiation transport techniques. However, because protons dominate displacement damage for electronics, adding sufficient shielding to significantly reduce displacement damage may add a lot of mass to the system.


2.2.3 Errors and Inference for SEE

Although SEE rates can vary from lot to lot for commercial devices, this variability is much less important than for TID. Instead, random errors on SEE cross sections fluctuate about a mean value μ according to the Poisson distribution:

$$P(n, \mu) = \frac{\mu^n}{n!} \exp(-\mu). \quad (2-1)$$

Percent errors due to Poisson fluctuations decrease as $\mu^{-1/2}$. Thus, as long as one can gather enough data, random errors can be reduced to the point where they are insignificant.

Systematic errors for SEE on the other hand can be pernicious, because they can bias experimental results in a particular direction and can be difficult to estimate. (See Figure 2-4.) In part, this is because the charge collection volume in the device under test (DUT) may deviate from the assumed rectangular parallelepiped. Also, the SEE cross section vs. LET curve may deviate systematically from the assumed Weibull form. Biases in radiation environment models constitute a third possible source of systematic errors.



I

Figure 2-4. The shallow trench isolation of SiGe HBTs coupled with the importance of charge collection by diffusion causes their charge collection volume to deviate significantly from a rectangular parallelepiped, resulting in violations of effective LET assumed for standard rate calculation methodologies. (Adapted from reference [26])

Reference [27] examined this problem, looking at the distributions of predicted to actual ratios of SEE rates by a variety of analysts for a variety of parts. The author determined that the canonical 2x uncertainty usually cited for SEE rates is a reasonable estimate of the systematic errors in rate calculation methods, but that estimates may be significantly worse depending on the part, the adequacy of the data and the procedure followed by the analyst. Unfortunately, it is likely that the development of technology in the past decade has resulted in significantly larger systematic errors, especially for some technologies [26].

Usually these systematic uncertainties dominate the errors on SEE rates. The only exception occurs when SEE cross sections are based on very low event counts, as is often the case for destructive SEE. In this situation, a variety of fits may be consistent with data including error bars. For the SEL data in Figure 2-5, the 99% worst-case fits to the data yield error rates about 5 times worse than those for the best fit. Whether one chooses a conservative or a tight fit to the data may depend not just on one's understanding of the random and systematic errors on the data, but also on the consequences of the SEE if it occurred.

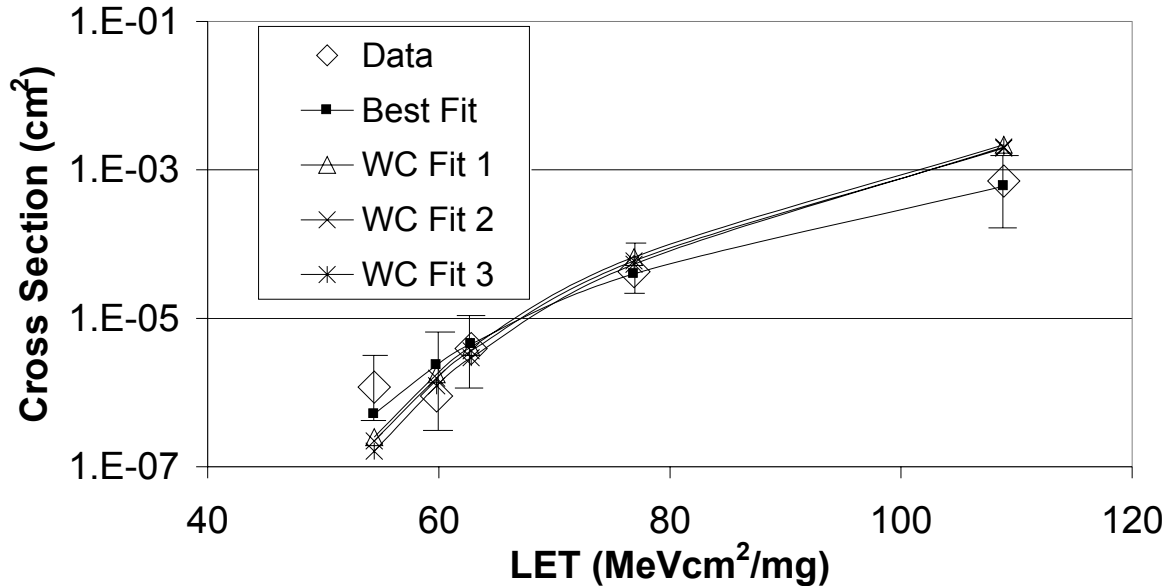


Figure 2-5, When SEE cross sections (σ) are based on low event counts, the σ vs. LET curve can be fit by a broad range of Weibull curves if one includes error bars. For this SEL σ vs. LET curve for a 256 Mbit SDRAM at 85 °C, the resulting curves make a difference of about 5 \times for the SEL rate.

2.3 Failure Cost and Severity

While engineers seek to avoid failure in any application, not all failures are of equal concern. Failures that result in a loss of mission objectives clearly merit greater level of effort than nuisance failures, and failures that could result in loss of life are of grave concern. One way of quantifying this difference is to look at the cost of the failure. Failure costs can be defined in terms of lost observing time for a science mission, dropped cell phone calls, data lost or fatalities. In terms of directing level of effort and comparing risks, it is easiest to define the failure cost monetarily. It should also be remembered that failures—especially those that make the news—have costs beyond the immediate loss of mission objectives, assets, etc. These may include loss of prestige, future business, and so on.

Estimating failure costs is a fairly involved effort. However, it can be very useful for efficient allocation of resources. It is here that the collaborative nature of radiation hardening is best illustrated, since the only way for a radiation specialist to understand the costs associated with an error or failure is by working with designers and system

engineers. The radiation engineer has an understanding of what a failure or error mode does at the device or component level. Designers usually have the best understanding of a system's functionality and will be able trace how the failure will affect it. Finally, system engineers will be best placed to understand how an error or failure mode will affect requirements. For this exercise it is essential that all involved understand the nature of the error or failure mode, its consequences from the part to the mission level and options for possible mitigation and recovery. One useful question to ask is: "If you knew with 100% certainty that this error/failure would occur the first day of the mission, what would you spend to stop it from happening." Then one can modify the question with the fault occurring midway through the mission, three-quarters of the way and so on. This gives an idea not only of the failure cost, but how it changes with time.

An alternative way of prioritizing hardening efforts is the concept of failure severity. MIL-STD 1629A defines severity as a relative measure of the consequences of a failure mode. Like cost, severity is defined in terms of the consequences of a failure. The advantage of severity is that it has a small number of broad categories, and reliability analyses often determine it as a matter of course. The Failure Modes Effects and Criticality Analysis (FMECA) is a very useful tool for identifying threats to system reliability, and the single-event effects criticality analysis (SEECA) can function similarly in SEE analysis.[28]

The NASA System Engineering Handbook defines 4 levels of criticality:

- Category I Catastrophic failure (possible death or system loss)
- Category II Critical failure (possible major injury or system damage)
- Category III Major failure (possible minor injury or mission effectiveness degradation)
- Category IV Minor failure (requires system maintenance, but does not pose a hazard to personnel or mission effectiveness).

MIL-STD 1629A defines severity levels as:

- Category I catastrophic
- Category II critical
- Category III marginal
- Category IV minor.

As with failure cost, failure severity must be determined in collaboration with designers and system engineers. Failure severity and failure cost can be related by assigning a dollar amount to each severity category—effectively scaling the level of effort to be directed toward mitigating a failure of a given severity category.

2.4 Failure Risk and Criticality

The failures of greatest concern are those with high severity/cost that are likely to happen. Failure criticality couples a failure severity with a probability of occurrence (See Figure 2-6.). While criticality is very useful for tracking system vulnerabilities, it can be difficult to compare different criticality categories: for example, is a Category II/high-probability failure a greater or lesser concern than a Category I/Low probability failure? The concept of failure risk, R_f combines failure cost— C_f —and failure probability— P_f —but in a quantitative (see Figure 2-7),

$$R_f = C_f P_f \quad (2-2)$$

More generally, both the probability and cost of failure may be functions of time, and the mission-integrated risk will be

$$R_f = \int C_f(t) x P_f(t) dt \quad (2-3)$$

Severity \ Probability	Category I (Catastrophic)	Category II (Critical)	Category III (Major)	Category IV (Minor)
High Probability	Very Critical	Critical	Moderately Critical	Moderately Critical
Moderate Probability	Very Critical	Critical	Moderately Critical	Acceptable
Low Probability	Critical	Moderately Critical	Acceptable	Acceptable

Figure 2-6 Criticality reflects the level of concern posed by a failure (and thus scales the level of effort appropriate for its mitigation) by qualitatively combining failure probability P_f and failure severity.

The quantitative nature of the risk metric facilitates comparison of risks and prioritization of resources. Moreover, since we can assign a cost to each severity

category and severity is often determined for FMECA, much of the work is often done for us.

The risk metric can also be used to direct risk mitigation efforts by optimizing risk reduction per unit cost. One way to do this is to adopt the strategy that minimizes a new combined risk metric R_c that also reflects testing costs C_t and mitigation costs C_m :

$$R_c = R_f + C_t + C_m \quad (2-4)$$

and adopt the strategy that minimizes this metric. Just as the cost of failure may include intangible elements, the cost of mitigation should include reduced performance, reduced reliability and other factors as well as the actual cost of implementing the mitigation—all of which are determined collaboratively with designers and system engineers.

Figure 2-8 represents the decision process for mitigation schematically. Initially, we have little information about our probability of failure or its consequences, and so we have a high, though not unbounded, risk. (If nothing else, the risk is bounded by the cost of the mission.) As we begin to test and analyze our critical component, our estimated risk falls rapidly. At some point, however, risk reduction vs. testing/analysis curve begins to flatten, and the reduction in risk we get from additional testing or analysis is less than the cost of that effort. We can, however, still lower the risk by mitigation, and we apply mitigations M1, M2 and so on in decreasing order of cost-effectiveness (that is, risk reduction minus mitigation cost) until we meet requirements or until there are no more cost-effective mitigation strategies. In the next section, we look at the radiation effects in our system and how we can mitigate them.

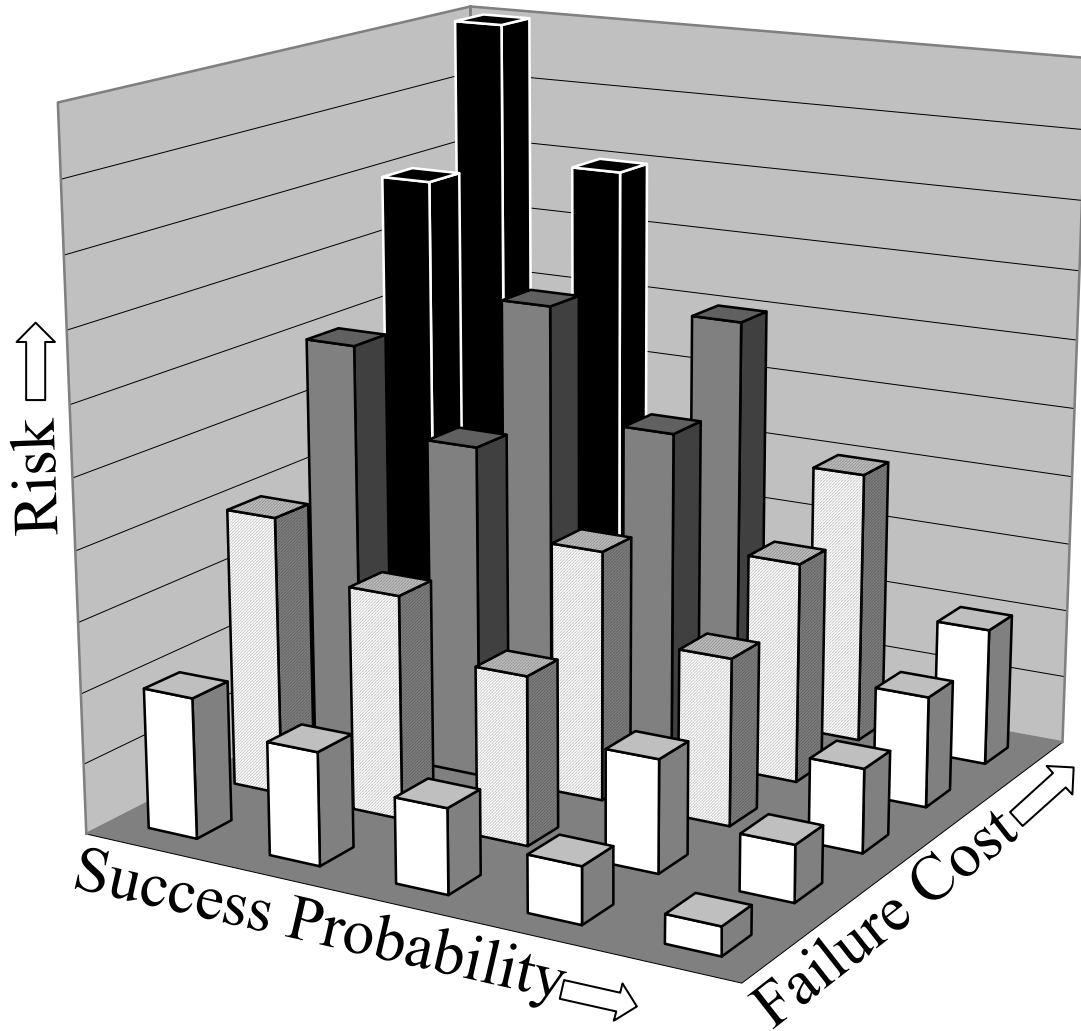


Figure 2-7. Risk is the product of failure cost, C_f and failure probability P_f , and so reflects a quantitative measure the level of concern posed by a failure.

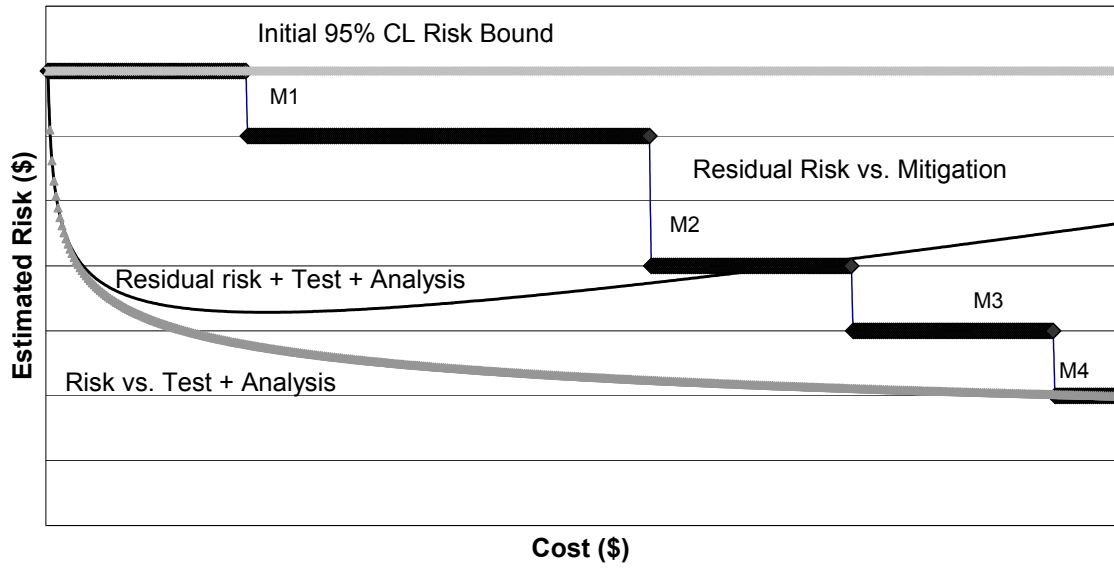


Figure 2-8. Schematic representation of tradeoffs between testing and analysis and mitigation.

3 RADIATION EFFECTS AND THEIR MITIGATION

At the system level, we are mainly interested in the consequences of radiation effects and their mitigation. While the basic mechanisms responsible for those effects are very important for hardening at the device or circuit level, they matter at the system level only to the extent that they affect our design or suggest mitigation strategies. Discussions of basic mechanisms can be found in past Short Courses and the references therein.

For this very high-level approach, we first look at the types of effects we need to mitigate—dividing the universe of radiation effects into prompt—or single-event effects—and cumulative effects, or TID and displacement damage. We further divide SEE between destructive and nondestructive SEE. Destructive SEE tend to affect particular technologies—such as SEL in CMOS devices or single-event gate rupture (SEGR) in power MOSFETs. On the other hand, nondestructive SEE tend to affect a particular cell type, regardless of its technology—for example, SEU occur in bistable cells, while if the cell is not bistable, it will be prone to SET. TID and displacement damage susceptibility depend on details of the technology, but also on application conditions. Moreover, while rapid destructive—or functional—failures due to TID and displacement damage do occur in some devices, more usually, failure occurs after a period of gradual parametric and functional degradation.

The consequences of radiation effects also depend on the application to varying degrees. Consequences of destructive SEE are usually unambiguous—loss of functionality or if the system contains redundant functionality, decreased reliability. Consequences of nondestructive SEE tend to be much more application dependent. An SET may have no effect if it is not captured by a bistable device downstream. The consequences of an SEU depend on which bit flopped and when. Usually there is little ambiguity when a SEFI occurs, since by its nature, a SEFI disrupts normal device functionality. Still, the consequences may be very different depending on when the SEFI occurs. TID and DD degradation may initially have no system-level consequences, as long as parametric degradation does not exceed system tolerances. Eventually, however, parametric degradation results in performance degradation and eventual functional failure.

In order to understand how various system mitigation techniques work, it is helpful to understand how the system responds to an error, failure or degradation. Regardless of the cause, at the system level, what is observed is anomalous function. Figure 3-1 illustrates the steps that occur in response to an anomaly (grey rectangles) and the mitigation strategies that can be helpful for each stage (white rectangles). Of course the most effective strategy is to keep the error from occurring by reducing its probability of occurrence. However, four out of five stages have to do with anomaly recovery, and the effective mitigation involves limiting the consequences or speeding that recovery. Error consequences can be limited by correction (e.g. EDAC or voting) or by isolation of the error (e.g. by preventing execution of a command prior to its validation). Error detection can be facilitated by such techniques as overcurrent sensing for SEL and watchdog timers and error counting to discover SEFIs. Error recovery and restoration of normal operations are facilitated by having a sufficient understanding of system error modes that a recovery plan can be implemented automatically by software or quickly by ground operations. Finally, even after normal operations are restored, mitigation of anomaly consequences can continue.

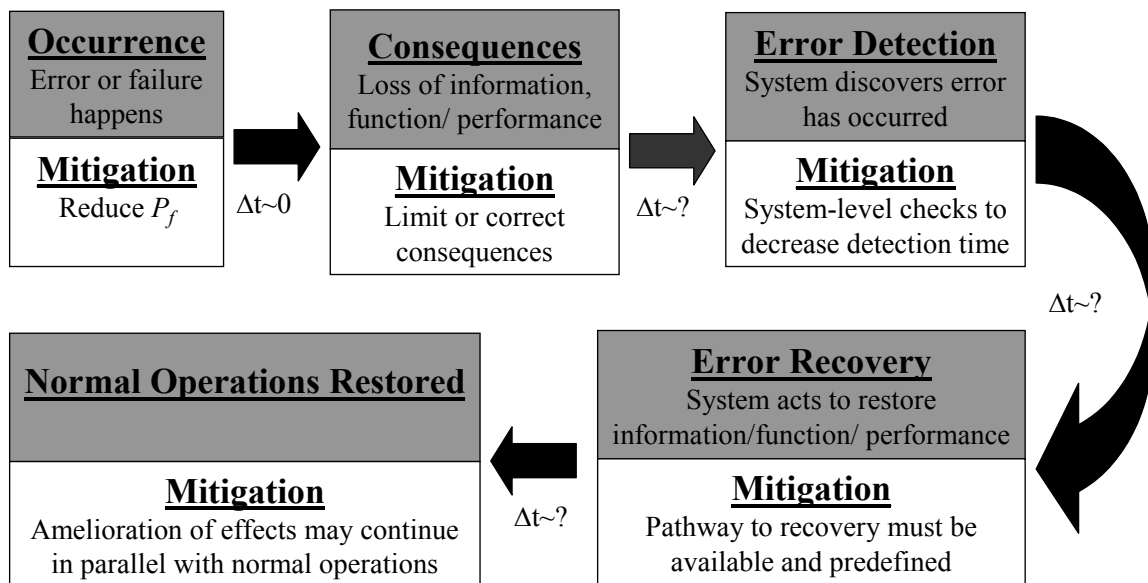


Figure 3-1. Steps in the process of recovery from an anomaly.

3.1 Radiation Effects and their Consequences

Here we discuss how radiation effects impact system operations—including reliability, functionality and performance. The purpose of the discussion is to give an idea of the types of consequences that will require mitigation. We begin by discussing destructive SEE, then nondestructive SEE and finally degradation mechanisms.

3.2 Destructive Single-Event Effects

The mechanisms of destructive SEE have been dealt with in previous short courses [29] and in the pages of the Transactions on Nuclear Science, most comprehensively and recently in reference [30]. Because destructive SEE susceptibility tends to be technology dependent, any device fabricated in a potentially vulnerable technology should be assumed vulnerable to the failure mechanism unless test data and/or design and validation demonstrate otherwise. Moreover, the degree of vulnerability cannot be determined without test data. Unfortunately, destructive SEE testing is difficult to carry out. Every data point may represent the destruction of an expensive microcircuit. Even if the destructive mechanism can be quenched (e.g. by detection and current limiting) prior to damaging the part, destructive mechanisms are at the very least very disruptive to the testing process. The result is that destructive SEE cross section estimations are often based on more limited data and subject to greater systematic errors than those for nondestructive SEE. In addition, there may be other complicating factors. A device that exhibits SEL may in fact be susceptible to multiple SEL modes—some of which are destructive while others are not. A common nondestructive mode can even mask a rare destructive mode. Such was the case for the XA-1 ASIC used on NASA's SWIFT gamma-ray burst telescope, which exhibited nondestructive SEL at LET~8 MeVcm²/mg with a saturated cross section of 5×10^{-3} cm². [31] However, at high LET, roughly one out of 30 SELs was found to be destructive. Given the 256 (128 primary and 128 redundant) XA1 ASICs being flown, the SEL threat was real, and the project had to implement mitigation (See below.)

Rate calculation for destructive SEE may also be subject to greater uncertainties than for nondestructive SEE. Cross sections for SEL and SEGR have been shown to depend on ion energy/range as well as LET, and effective LET is invalid for single-event gate rupture (SEGR) and single-event burnout (SEB) and related phenomena. As we will

discuss below, all of these factors argue for a conservative approach when deciding whether to use a device susceptible to destructive SEE and whether system-level mitigation is required for such a device.

Most destructive SEE render a single part inoperable, although for stuck bits or single-event dielectric rupture, only a portion of the part is rendered inoperable. Specific radiation-induced failure modes encountered in microelectronic components are discussed below.

3.2.1 Single-event latchup (SEL)

A *pnpn* structure in CMOS is equivalent to a bipolar silicon-controlled rectifier (SCR) that exhibit a regenerative latchup mechanism. (See Figure 3-2) Once triggered, this regenerative mechanism can amplify currents to the point where the device fails due to thermal overstress. Although SEL has been observed only in CMOS technologies, SEL is inherently a bipolar phenomenon that depends critically on the bipolar gains of parasitic bipolar junctions implicit in *pnpn* structures. As a result, SEL is exacerbated at high temperatures (Figure 3-3a) [32], [33]. Also, the parasitic structures inherently depend on substrate characteristics, so cross sections depend on ion energy and range, not just on an ion's LET [34], [35]. (See Figure 3-3b)

Protons and neutrons can initiate SEL as well as heavy ions. Recent in-depth investigations have shown that at least for some technologies, proton-induced SEL susceptibilities can increase significantly with proton energies [36] and may be worst-case at grazing incidence (possibly due to nuclear inelastic scattering of heavy metals in the part itself) [37].



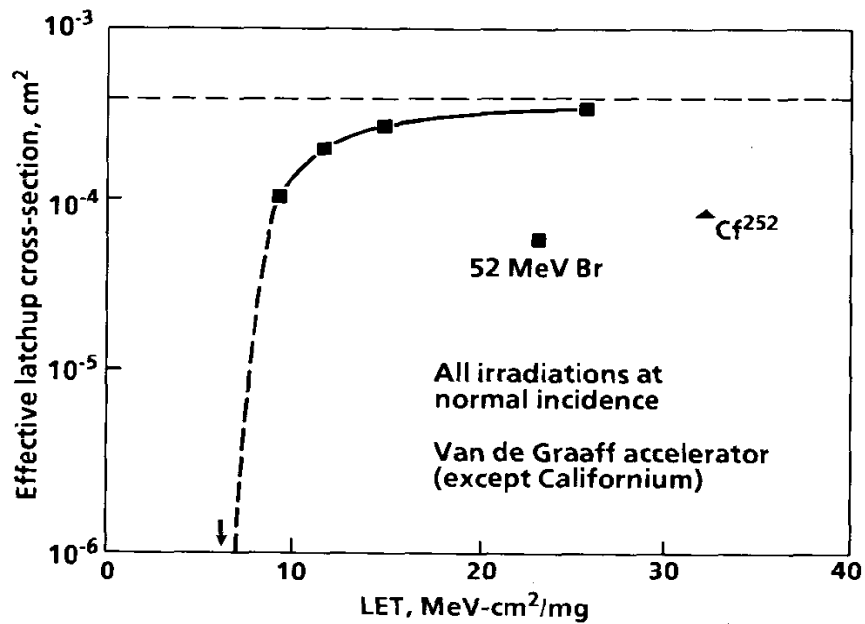


Figure 3-2. a) The parasitic bipolar SEL susceptible structure and its equivalent circuit. (Adapted from reference [29])

Understanding the implications of SEL for a system is complicated by the fact that SEL is not invariably destructive. In some cases—so-called micro-latch events resulting in localized high current—devices sometimes even remain functional. In some cases, the SEL may involve the entire device, but current densities may not be high enough to damage the device. However, even here, the situation is complicated by the fact that a device remaining functional after SEL is not a guarantee that it has not suffered damage that would compromise its reliability.[38] Such latent damage can be very difficult to detect, especially in devices with multiple metallization layers. NASA has issued an advisory on best practice for testing for and detecting latent damage in SEL susceptible circuitry [39]. Investigations of possible latent damage are also described in references [40] and [41].



(a)



(b)

Figure 3-3. a) The parasitic bipolar nature of SEL means susceptibility depends critically on temperature (after reference [32]). b) Similarly, because the bipolar structure involves the substrate, cross section depends on ion energy and range as well as LET (after reference [34]). Low-energy ions give lower cross sections whether they come from accelerators (Br) or the fissioning of Cf-252.

Given the variety of SEL responses a part can exhibit, testing for SEL susceptibility can be quite complicated. From this very brief summary, the important issues for system-level hardening are:

- 1) Understanding of SEL test data and its errors and limitations is essential to making appropriate hardening decisions.
- 2) Investigating latent damage is critical, because mitigations that are effective for destructive SEL are very different from those used for nondestructive mechanisms.
- 3) SEL is a high-current mechanism, and detecting this high current is the easiest way to determine that an SEL has occurred. However, it may not work for microlatch or if several devices are sharing a single power source as in a package of stacked die.
- 4) SEL is exacerbated at high temperature, so it is important to that testing envelops worst-case operating temperatures. In some cases, if the operating temperature can be maintained below a critical level SEL susceptibility can be lowered to the point where it is negligible.
- 5) SEL testing needs to take place at worst-case conditions, including the highest application voltage and temperature, and sufficiently high-energy heavy-ions and protons (if needed).
- 6) Destructive SEL renders the affected die inoperable. Because it is not possible to predict the electrical properties of the failed part, it is recommended that any mitigation strategy should electrically isolate the failed part from the rest of the system.

3.2.2 Single-Event Burnout (SEB)

Even when there is not a *pnpn* structure, an ion strike can turn on a real BJT or a parasitic BJT structure in a (usually) n-channel MOSFET. The resulting second-breakdown causes a high-current state and can cause thermal failure of the device [42]. One can initiate a SEB without necessarily destroying the device, so it is usually possible to gather sufficient statistics to reasonably bound SEB susceptibility for a transistor.

Discrete transistors are only vulnerable to SEB when they are in their nonconducting (or OFF) state, and only when applied voltages (VCE for a BJT or VDS and VGS for a power MOSFET) are outside of a so-called safe-operating region, which must be defined by testing. As a rule of thumb, radiation hardened MOSFETs with rated VDS < 200V have not been seen to fail at less than 30-35% of their rated VDS for VGS=0. However, commercial MOSFETs have failed as low as 22% of rated VDS.[43] Operating devices within their safe operating voltages is the only effective mitigation for SEB. Effective LET is invalid for SEB, so typical rate calculation methods do not work.

3.2.3 Single-Event Gate Rupture (SEGR)

As with SEB, SEGR only affects discrete transistors (in this case power MOSFETs) when they are in their nonconducting states ($V_{GS} \leq 0V$ for n-channel devices or $V_{GS} \geq 0V$ for p channel devices). Moreover, the effective mitigation strategy for SEGR is the same as that for SEB—derating of applied voltages to within empirically determined safe operating regions. The mechanism for SEGR is very different than that for SEB. In the case of SEGR, holes from the ion strike pile up under the gate, increasing the electric field across the MOSFET gate oxide to its dielectric breakdown point. (See Figure 3-4)




Figure 3-4. SEGR occurs when holes from the ion track aggregate under the gate oxide and increase the field across the oxide to the dielectric breakdown point. (Adapted from reference [44])

The resulting flow of current causes thermal failure of the gate oxide. As with SEL, SEGR testing is complicated by the fact a strip chart of gate current vs. ion fluence reveals so-called SEGR precursor events, where the leakage current through the gate increases in a stair-step fashion. These events represent localized breakdowns in the oxide and can result in latent damage. In addition, SEGR susceptibility depends on particle range/energy as well as LET (see figure 3-5). This means that commonly used SEE rate calculation methods will not work for estimating SEGR rates, although other methods have been investigated [46].




Figure 3-5 The complicated dependence of SEGR susceptibility on LET, ion energy and angle of incidence make it very difficult to calculate an accurate SEGR rate using conventional methods. (Adapted from reference [45])

3.2.4 Other Destructive SEE

Single-Event Dielectric Rupture (SEDR) was found to occur in one-time programmable antifuse-based FPGAs. The failure mode disables the affected gate and can cause partial loss of functionality while leaving most of the device functional. Rates for this failure mode were low, because the antifuse presents a small cross section and the

rupture follows an angular dependence similar to that for SEGR. Moreover, effective workarounds were found that mitigated reliability concerns and more recent generations of these devices seem to be more immune to such breakdowns [47].

Second Breakdown or burnout in bipolar devices has been known since 1994, when it was seen in a high-speed ECL devices [48]. More recently, it has been seen in several bipolar linear ICs [49],[50]. An angular dependence similar to SEGR and SEDR, and onset LET over $30 \text{ MeVcm}^2/\text{mg}$ ensure that rates for this failure mode seen to date are low. Moreover, the parts where this phenomenon has been seen are usually not so unique that another part immune to destructive failure could not be substituted if the rate were problematic.

Single-Event Snapback (SES), like SEL is a regenerative effect that occurs in CMOS, but unlike SEL, it does not require a 4-layer *pnpn* structure [51]. SES occurs when an ion injects sufficient current into the drain junction of an n-MOS transistor resulting in avalanche multiplication. If the event is sustained by an external current source, the overcurrent can damage the device. SES has not yet emerged as a serious reliability concern. However, since it does not require a 4 layer structure, it is potentially a concern for SOI technologies [52]. Most hardening against SES is done at the process level, although as indicated above, current limiting can be effective in limiting damage due to this effect.

Failures in FLASH Memories have been seen to occur when the parts are operated in their WRITE or ERASE modes, even as these parts have exhibited significant immunity to all SEE when operated in their READ mode or when unbiased [53]. The failures during WRITE and ERASE are not surprising, since these modes use a charge pump to achieve the high voltages needed to program or erase the FLASH memory cells, and the failures are probably due to SEGR. For some parts, the onset of these errors occurs at sufficiently high LET that limited reprogramming would pose a relatively low risk.

Stuck bits—When an ion passes through a gate oxide of a transistor in a bistable cell (or in a pseudobistable Flash or DRAM memory cell), it may leave a dense local deposit of trapped charge that renders the cell unprogrammable [54]. Such a stuck bit may have an indeterminate output, or it may reside at its bleed-down value (0 or 1). At the system

level, it may look like a permanent SEU, and it reduces the effectiveness of mitigation for nondestructive SEE. To date, stuck bits have not occurred at rates that cause concern for most applications. Moreover, stuck bits usually anneal with time, so they are unlikely to accumulate to levels where they pose a serious concern.

3.3 Nondestructive SEE

In contrast to destructive SEE tests, those for nondestructive SEE face no obstacle to gathering enough events at each LET value to make Poisson errors negligible. For this reason, systematic errors often determine the uncertainties for nondestructive SEE rate calculations. Another important thing to emphasize about nondestructive SEE is that the adjective nondestructive applies to the device in which the effect occurred, not necessarily the system in which that device resides. It is small comfort to know that if an SET puts a satellite into a flat spin, the device that initiated the disaster is still in good operating condition. Thus, it is important not only to understand the nature of the SEE, but also to understand it in the context of the component application. Consequences of nondestructive SEE often depend on application conditions, when the error occurs in the mission and so on. Nowhere is the application dependence more evident than when we consider single-event transients.

3.3.1 Single-Event Transients (SET)

The effect of a transient is determined by its magnitude and duration (and possibly its waveform) and by whether a sensitive device downstream latches the transient and turns it into an upset at the system level. Transients are divided in analog SET (ASET)—usually associated with linear analog circuits—and digital SET (DSET), which are produced in digital circuits. ASET amplitudes and durations—especially in linear bipolar components—are often sensitive to application conditions including supply voltages, input voltages, loads and so on. (See Figure 3-6.) DSET are also produced with a range of durations and amplitudes, but tend to be shorter than ASET. Both types of transients tend to dissipate due to capacitive filtering, fanout and passage through several intermediate gates on the way to a bistable latch.

Whether the latch responds to the transient depends on how much the transient amplitude has decreased and when the transient occurs relative to clock edges[56]. For

this reason, it is important that the SET test setup have sufficient fidelity to the applications that circuit behavior can be accurately inferred. Because achieving such high test fidelity may require testing under many application conditions, SET testing at a heavy ion facility can be costly. Reference [57] discusses the more cost effective option of using a pulsed laser system to map out the application dependence of the SETs, with heavy-ion testing reserved for verification and determination and SET rate estimation.




Figure 3-6. SET durations for the LM124 depend not just on the gain, but on how the gain is achieved. Reference [55] showed that the width of transients from a particular transistor in the gain stage (Q9 specifically) of an LM124 in an inverting configuration depended on what resistance values were chosen to achieve the desired gain (here 10x) as well as the gain itself.

Since a transient's effect depends on the sensitivity of devices downstream, SETs remain a dynamic area of radiation effects research. The recent concern over DSET [58] and over-voltage conditions due to transients for deep submicron CMOS technologies are examples of new SET related threats [59]. As Ben Blalock discussed earlier in this volume, transients from a few linear bipolar devices last up to a millisecond [60]. In addition, what constitutes a long transient is relative. In very-high-speed technologies

such as SiGe and some high-speed A-to-D converters, transients may last only a few ns, but can corrupt dozens of data bits, giving rise to burst errors. Even ultra-fast devices like rf amplifiers may now need to be considered as a source of SET when they interface to very fast technologies. Each new generation of technology means we may need to reconsider our previous assumptions about potential SET related threats—and do so with no data from previous programs on which to base our conclusions.

For system mitigation of SETs, the important considerations are their short duration and localized effect, as well as the very application dependence that complicates analysis of transient effects. These characteristics suggest we can mitigate SETs by capacitive filtering, by resampling the output on timescales longer than transient durations, or by choosing our applied voltages, loads, etc. to minimize SET sensitivity.

3.3.2 Single-event Upset (SEU)

SEU analysis is more deterministic than that for SET. Still, consequences of an SEU depend as much on the state of the system when the SEU occurs as they do on which bit upsets. In some cases, an SEU will have serious consequences only if it occurs during a narrow time window. For some applications, SEU may be correctable or have acceptable consequences if they occur at less than a critical rate, while in other applications a single bit flip could have serious effects.

Increasingly, because SET within a microcircuit may manifest as SEU, SEU rates are becoming more dependent on application conditions such as frequency, which makes it important to test at worst-case operating frequencies [61].

For mitigation purposes, SEU corrupt information stored in the device in which they occurs. Its effects, while persistent, are both local and correctable. However, increasing density due to scaling along with the burst errors mentioned above result in ever higher proportions of events that corrupt many bits simultaneously. More and more SEUs have consequences of multi-cell upsets (MCU).

3.3.3 Multi-Cell and Multi-Bit Upsets (MCU and MBU)

Whether an SEU corrupts a single bit or multiple bits in a data word, the consequence is the same—corrupted data. MBU are mainly a concern because they are more difficult to mitigate. For example, MBU can undermine simple parity checks and other error

detection used in many command and control systems, just as they can defeat many error correction codes in system memory. Mitigating MBU requires much more sophisticated (and costly) mitigation than that used for SEU, so it is important to know whether MBU are possible in the system.

Although the continuing decrease in microcircuit feature sizes has resulted in increased occurrence of upsets that affect several adjacent cells—so-called MCU—many state-of-the-art memories now interleave bits from different words so that bits from the same word are never physically adjacent. To determine whether the separations used are sufficient, MBU tests must include worst-case conditions. For instance, unless one knows how the memory is organized, it is advisable to rotate the DUT about both the tilt and roll axes and to use angles as close to grazing incidence as possible. Moreover, since charge shared between adjacent cells increases with LET, the MBU cross section is even less likely to exhibit saturation than the SEU cross section.

As device sizes continue to shrink, the question of how much separation is sufficient is under debate. Most accelerators produce ions with energies ranging from 10 to 40 MeV/amu and ranges up to a few hundred microns in silicon, while the GCR flux peaks at energies of about 1 GeV/amu and ranges on the order of centimeters in silicon. The higher energies of GCR protons and light ions produce higher energy recoil particles as well. The longer ranges and broader ionization tracks of these ions raise concern that terrestrial SEE testing could underestimate MBU cross sections. (See Figure 3-7) This is an area where modeling may be helpful if bit maps can be obtained from vendors or reverse engineered. Reference [62] reports on using a pulsed laser to reverse engineer the bit layout for a SRAM, which had a particularly a simple (and as it turns out, flawed) pattern of interleaving. Reference [63] discusses algorithms for finding the bit layout using broad-beam data, although such algorithms are laborious and computationally intensive. Although MBU require more sophisticated error correction schemes they still involve only a relatively small number of localized and physically adjacent bits.

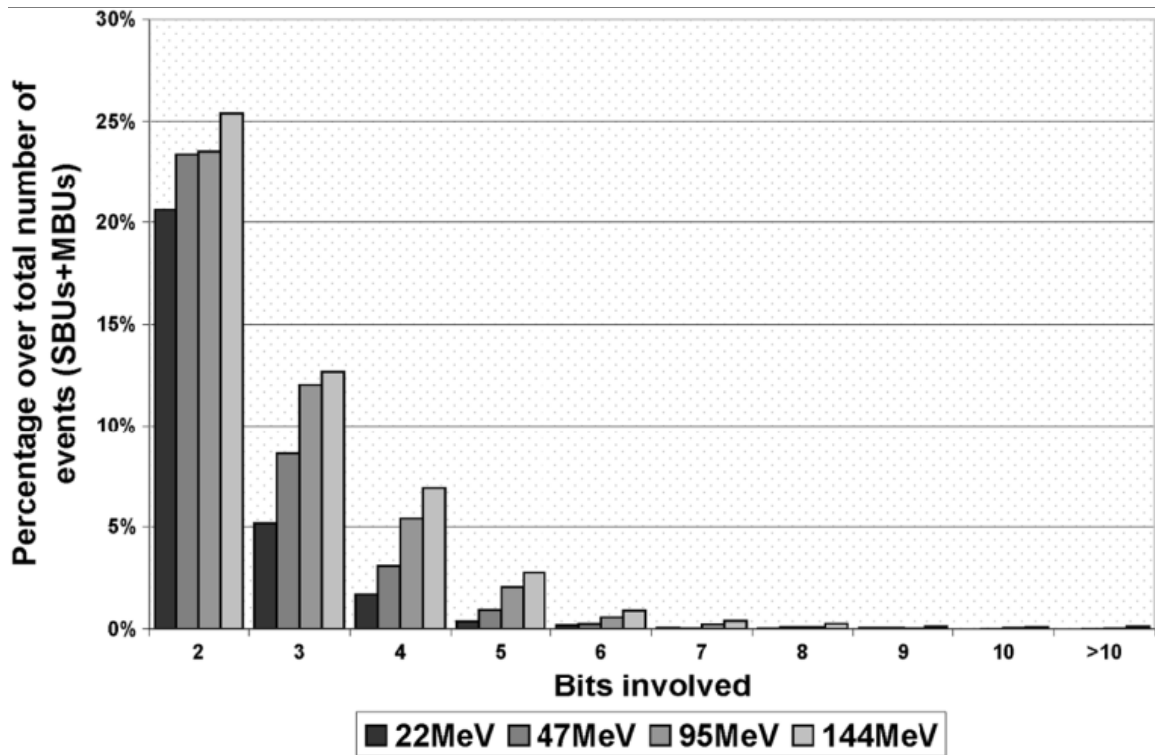


Figure 3-7. Proportion of MCU corrupting from 2 to >10 bits for different neutron energies incident on an SRAM fabricated in 150 nm feature size technology. Bits were interleaved with bits from the same word separated by at least 8 columns. Although some upsets involved more than 8 bits, they never spanned more than 6 columns. (Adapted from Reference [64])

3.3.4 Single-Event Functional Interrupt (SEFI)

SEFI are the most disruptive of nondestructive SEE because they render the affected device unable to carry out its normal functions and can corrupt large amounts of data. SEFI often require action at the system level for recovery—actions ranging from a refresh of corrupted control bits to cycling power for the device, reloading the pre-upset configuration and redoing the interrupted function. Since the first rule of ground/support operations is to do no harm, unless the SEFI can be diagnosed and corrected autonomously, it may result in an extended outage.

As complexity of microcircuits increases, so do the complexity and number of SEFI they exhibit. A recent analysis [65] of the time required to characterize fully a DDR SDRAM over 3 parts, 3 data patterns, 3 heavy ions and all of its 68 different modes of operation would take years! Indeed, since each different operating mode likely exhibits

its own SEFI types and that post analysis would likely be needed to separate the different SEFI types, the analysis effort would also be unmanageable. For this reason, if a system uses a device susceptible to complicated SEFI, it is usually necessary to coordinate testing and mitigation efforts—both to verify the effectiveness of mitigation and to limit the scope of the testing campaign to identifying and characterizing any SEFI modes that impact the system despite mitigation. Perhaps the only good news about SEFI is that they are usually obvious when they occur (see Figure 3-8), and worst-case recovery involves cycling power to the affected component, circuit or system. For purposes of mitigation the thing to remember is that although very disruptive, SEFI are recoverable and isolated to a single integrated circuit.

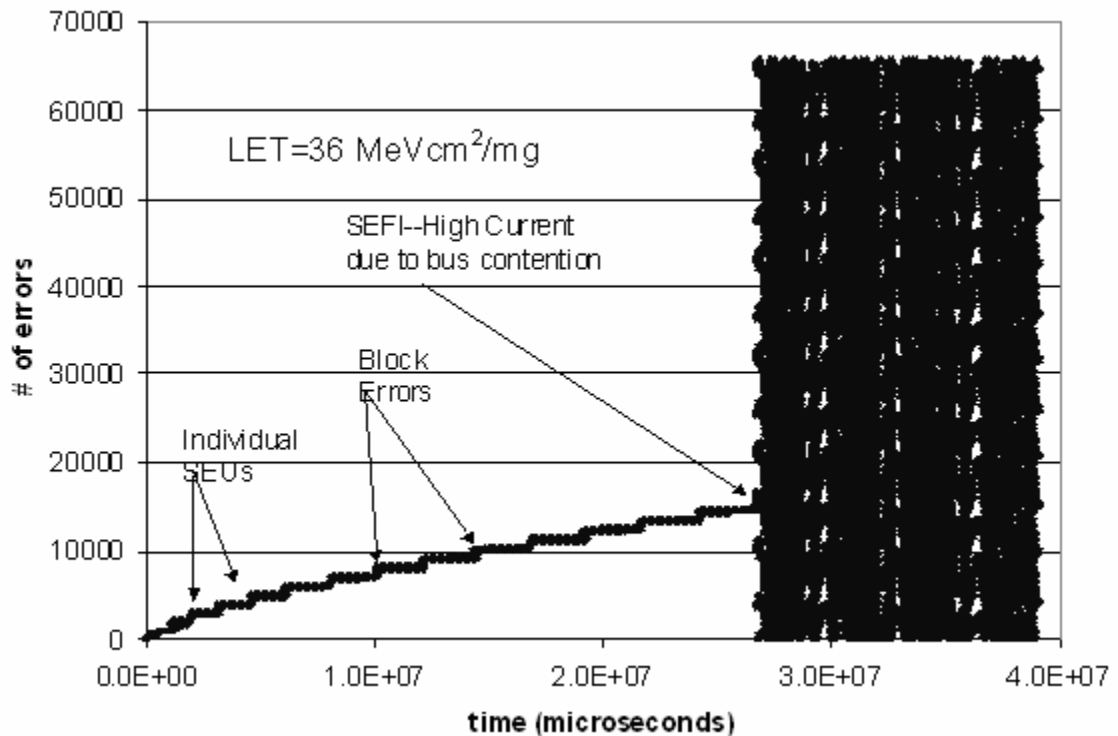


Figure 3-8. A strip chart of errors vs. time for a 1 Gbit DDR SDRAM under irradiation with Xe ions with LET~36 MeVcm²/mg shows errors increasing steadily with occasional small block error. When a SEFI occurs(~28 s), the error count increases dramatically, overflowing the buffer repeatedly. SEFI are often easy to spot.

3.4 Degradation Mechanisms

Unlike SEE, TID and displacement damage are cumulative effects, and so are *more likely* to affect a system near end of life than early in the mission. Moreover, *usually*

failures due to these mechanisms result from gradual parametric degradation rather than abrupt functional failure. Nonetheless, degradation mechanisms have a way of producing nasty surprises if we become too sanguine about them. TID failure have been observed at doses below 5 krad(Si), and the combined effects of TID and displacement damage have resulted in early failures in proton dominated environments. [66].

3.4.1 Total-Ionizing Dose (TID)

TID degradation in individual transistors is well understood. Ionizing radiation generates trapped charge in device dielectrics—both in the bulk dielectric and at its interface with the semiconductor. In MOS transistors, these trapped charges decrease (increase) threshold voltage for n channel (p channel) devices, increase leakage currents and degrade switching performance [67]. In bipolar transistors the main effects of trapped charge are reduced transistor gain, increased leakage currents and degraded AC/DC performance [68]. In electro-optical materials, in addition to the degradation in active portions of the device, one must also consider the ionization-induced degradation of passive materials by color center formation.

In an integrated circuit, a transistor's degradation will likely be masked until it exhausts the margin in its application within the IC. How the IC behaves at this point determines whether we must limit the TID the part sees (e.g. by shielding) or whether we have other mitigation options. For instance, when the Micrel MIC29372 low-dropout voltage regulators were found to degrade much more rapidly when unbiased, one proposed remedy was to maintain bias on the parts even when they were not in use [69]. Alternatively, the degradation modes observed in testing may suggest circuit or operational modifications to compensate for the degradation. For example, current compensation circuitry may be necessary to provide increased drive current to a linear bipolar component as its gain drops. The cost of these strategies is of course power efficiency.

An important aspect of TID degradation is that unlike SEE, devices are vulnerable whether or not they are in use. Indeed, as shown by the example of the MIC29372 above, a cold spare can degrade even more rapidly than the primary device, and this significantly limits the value of redundancy as mitigation for TID. A second important aspect of TID degradation is that it can vary significantly from wafer diffusion lot to wafer diffusion lot,

and even from part to part within a wafer diffusion lot. This elevates the importance of having radiation lot acceptance testing on which to base mitigation decisions—a serious challenge for using commercial parts, since even defining a lot can be difficult for some devices.

3.4.2 Displacement Damage (DD)

Displacement damage is potentially an issue for any technology that may be adversely affected by:

- 1) Minority carrier lifetime reduction
- 2) Reduced carrier mobility
- 3) Carrier removal
- 4) Increased leakage current
- 5) Thermal charge generation

In particular, if minority carriers play a significant role in the operation of the device (as in bipolar technologies), then displacement damage may affect its operation. The complicated physics of defect formation, evolution and annealing make it very difficult to understand or model displacement damage beyond these qualitative susceptibilities. Even such basic questions as damage energy dependence must be resolved empirically by testing at several fluence steps and several energies—usually with either protons or neutrons. The quality of such determinations can become important as we try to mitigate degradation from displacement damage. Fortunately, part-to-part variability and application dependence seem to be less significant for displacement damage than for TID, and most parts degrade fairly gracefully, rather than failing catastrophically.

However, in many cases, devices will be sensitive to both TID and displacement damage, and the predicted damage from both sources must be combined. Reference [70] documents instances where displacement damage effects have caused bipolar and optoelectronic devices to degrade more rapidly in proton- and neutron-rich environments. References [71] and [72] discuss some of the synergistic effects that occur when displacement damage and TID degradation are combined. From the system point of view, the important consideration is that the method of combination bound the degradation. Usually, a linear combination of displacement damage and TID degradation is sufficient to do this.

Mitigation strategies must also consider the part's susceptibility to both TID and DD. Shielding reduces both TID and DD. Likewise, there can be synergy between other mitigations for TID and DD. In this context, optocouplers provide a useful illustration. Because optocouplers are hybrid devices consisting of several discrete parts, they exhibit a broad range of degradation mechanisms in a radiation environment. Displacement damage can reduce LED efficiency due to nonradiative carrier recombination at defects in the semiconductor (see Figure 3-9) and degrade the sensitivity of the photo-receiver increased leakage current and charge trapping. Finally, passive components such as fiberoptics and light guides may have their transparency degraded by formation of color centers by TID. One way of compensating for the lost efficiency of both the light source and the photo-receiver is to overdrive the light source at a higher forward current. As a bonus, it turns out that this will also ameliorate TID damage in passive components by photo-bleaching the color centers in the passive components.

As with TID, displacement damage accumulates whether or not the component is in use, limiting the effectiveness of cold sparing for functional failure due to displacement damage. However, part-to-part and lot-to-lot variation are not typically as significant for displacement damage as they are for TID.

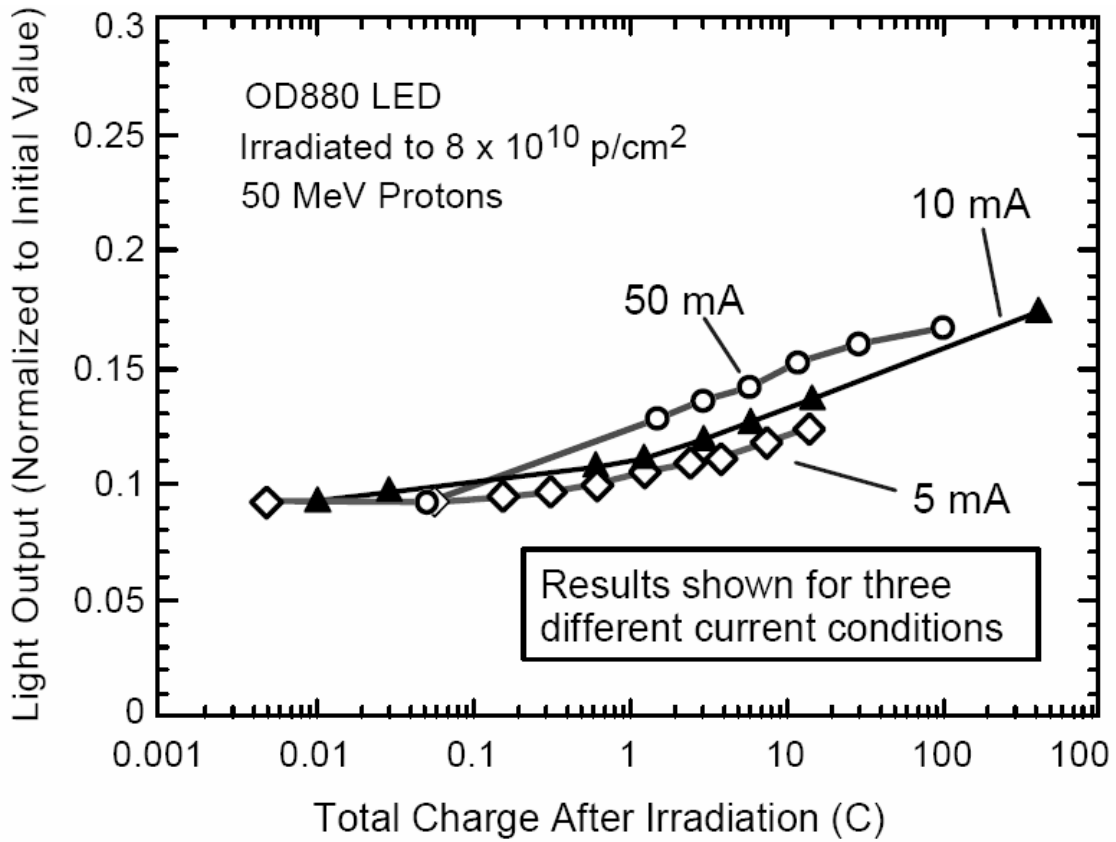


Figure 3-9. Higher forward currents can lead to greater photo-induced annealing in amphotericly doped LEDs. The figure shows the partial recovery of an OD880 LED as charge flows through it at different currents after irradiation with 8×10^{10} 50 MeV protons per square cm. (After reference [73])

4 MITIGATION STRATEGIES

In the discussion of radiation effects in the last section, we pointed out characteristics that are important in developing mitigation strategies. Here we treat mitigation more systematically. In mitigating radiation effects, we are trying to reduce risk—that is, we are trying to either eliminate or reduce the probability of an radiation-induced anomaly or to limit the consequences to the system if the anomaly occurs.

At the system level, we are limited in the strategies we can use to decrease the probability of an anomaly. Because we are often stuck with the radiation performance of the parts the system needs to meet its performance requirements, most system level mitigations seek to decrease the cost of a failure.

Mainly for mnemonic purposes, I have divided system-level mitigation strategies into 5 different categories (see Figure 4-1).

Threat reduction strategies seek to reduce the probability of an anomaly by reducing the stresses related to the underlying radiation mechanism. Shielding can be effective against TID and displacement damage. Derating of operating voltages for power transistors can decrease or eliminate susceptibility to SEB and SEGR. Limiting operating temperature can decrease SEL rates. We may even operate at a lower frequency to decrease the probability that a downstream device captures an SET. Selection of radiation hardened parts in the first place is probably the most effective threat-reduction strategy—albeit, not a system level one.

Performance Matching seeks to avoid overperforming conditions under which a component may be more susceptible to errors. Capacitive filtering or reducing speed to reduce SET capture probability are examples of such strategies.

Redundancy is probably the most versatile of the mitigation strategies—being effective against both destructive (use of cold spares) and nondestructive (use of EDAC, voting, resampling, etc.) SEE. Unfortunately, it is not very effective against degradation mechanisms, since the redundant elements are likely to be as susceptible to degradation as the elements they are trying to protect.

Opportunistic strategies are another versatile category and capitalize on characteristics of the radiation effect to mitigate it. For instance, Figure 3-6 shows that LM124 SET width depends on the resistance values we pick to realize our inverting 10×

amplifier. This suggests that if we have some flexibility in our application, we can get away with less capacitive filtering if we choose our design carefully. Likewise, we can use bias dependence of TID degradation or facilitate annealing of displacement damage in amphoterically doped LEDs (as in Figure 3-9) by running at higher current. These are all opportunistic strategies because they take advantage of some characteristic revealed about the radiation threat during testing to mitigate it.

Relying on Infrastructure is the final category of strategies. These strategies accept the inevitability of some anomalies occurring and rely on software, ground operations or the end user to mitigate the consequences. Examples of such strategies range from implementation of watchdog timer or error counters to detect SEFIs to expecting customers to redial if a satellite link drops their phone call.

Not all of these strategies will be effective for any radiation threat, and in some cases, different strategies will need to be combined to achieve an effective mitigation. Which strategy will be most effective depends on the nature of the threat and on its consequences. If more than one strategy is possible, the way to decide between them is to look at which one achieves system requirements most cost effectively.

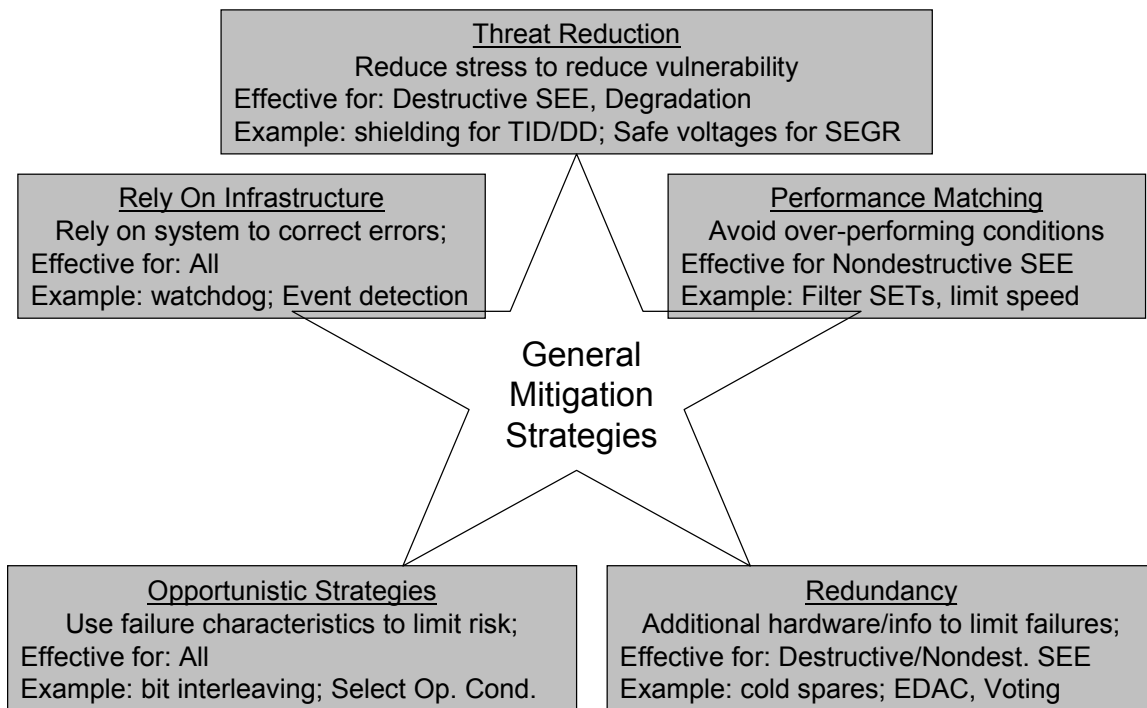


Figure 4-1 General categories of mitigation strategies for radiation effects.

4.1 Mitigation for Destructive SEE

As discussed above, destructive SEE pose a variety of challenges for radiation hardness assurance—for testing and rate calculation, and most especially for mitigation. For this reason, the preferred mitigation strategies for these threats are those that allow us to avoid them, or at least reduce their probability to negligible levels. Either application conditions should be regulated to avoid the failure mode, or if possible, another part that is immune to the failure mode should be substituted. However, if the part is essential to meeting system performance requirements, other mitigations can increase system reliability even with unreliable parts. One strategy is to detect and quench the event before it can damage the part. However, such detection and protection strategies are not always possible—or if they are possible, they may pose unacceptable costs to system performance. If we cannot avoid the potential of destructive, the only effective strategy is to replace the failed part with a cold spare—a process that is both costly and difficult. (See Figure 4-2.)

4.1.1 Threat Reduction

The restriction of power MOSFETs and power BJTs to their safe operating voltages is the preferred method for avoiding SEGR and SEB. Because current levels in circuits using these parts are generally high, and because voltage regulation in these circuits will often necessitate the use of large capacitances, overcurrent sensing and protection are usually not practical or effective. Likewise, given the voltage and current regulation performed by these parts, cold spare substitution would be cumbersome and could reduce system efficiency.

Similarly, maintaining operating temperatures at levels where SEL rates are low can be effective for some parts. In the case of the SDRAMs tested for SDO, no SEL was seen for temperatures below about 50 degrees Centigrade. While maintaining operating temperatures this low is challenging, it has the added benefit of increasing reliability of microelectronic parts.

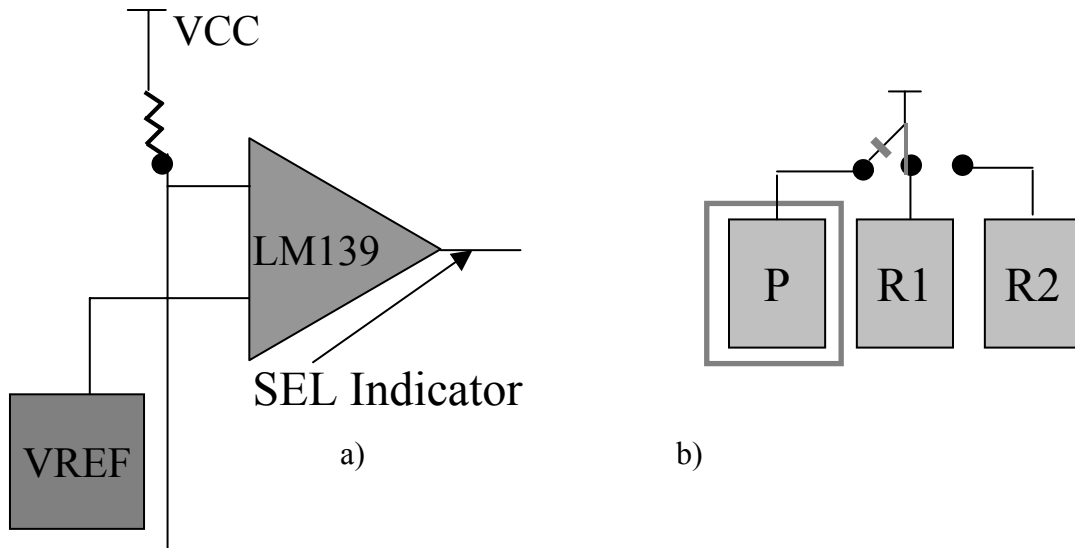


Figure 4-2. a) SEL detection circuitry typically senses an overcurrent due to SEL as a voltage drop across a resistance. b) In a cold sparring mitigation, a redundant part is brought on line to replace a failed part.

4.1.2 Event Detection and Protection

To date, the only destructive SEE mode where detection and protection have been used successfully has been SEL. Generally, an SEL detection circuit operates along the principles illustrated in Figure 4-2. The supply current to the SEL susceptible device flows through a resistor, and the resulting voltage drop is fed into a comparator and compared to a reference voltage. If the current surges due to an SEL, the output of the comparator changes, and usually, VCC is clamped off for a short period ($\ll 1$ s) to quench the SEL.

The challenges associated with such event detection circuitry are that in order to protect against the threat, it must be rapid, while this very rapidity increases its susceptibility to SET that give spurious indications of SEL. In some cases, the SEL detection circuitry has been implemented as a radiation-hardened ASIC to limit such transients.

Another challenge for this strategy is that it must not only be effective against functional failure but also against latent damage. Demonstrating such effectiveness is both challenging and expensive, since the damage may be hidden in a part with several layers of metallization. For this reason, reference [39] recommends a 1000 hour burn-in/lifetest post SEL in addition to microscopic and SEM inspection.

Although to date such solutions have been implemented only for parts susceptible to SEL, they could be also be implemented for parts susceptible to other high-current failure modes, provided these parts were crucial enough to the system to warrant both the effort of developing and validating such circuitry and could sustain the cost of spurious event indications. References [40] and [41] detail efforts to preclude latent damage in a part protected by an SEL detection and protection circuit.

4.1.3 Cold Sparing

If a destructive SEE cannot be avoided and cannot be detected and stopped before it damages the part, the only choice for mitigating the failure is to replace the failed part with an identical part that can fulfill its requirements. Usually, the redundant part is unbiased until it is needed, since it, too is susceptible to the same failure mode(s)—hence the name, cold sparing. This is a very expensive option. Not only must the system bear the dead weight of a part (or parts) that may never be needed, it must electrically isolate the failed part and bring the redundant part on line in the role of the failed part.

Electrically isolating the failed part is challenging in itself, although architectures have been proposed for doing this with radiation-hardened parts [74]. Cold sparing is most straightforward in a 2:1 redundant system, where there is only one failed part whose role must be assumed by the redundant. It is less trivial for a 3:2 redundant system, where the spare must fulfill the role of either part if it fails. As the system and redundancy become more complicated, the benefit of the scheme begins to depend critically on implementation. In an $n:m$ fully redundant system, any redundant part can replace any other part, and if the expected lifetime of a single part is T , then the expected lifetime of the system will be $(n-m+1)T$. A similar situation often arises when a spacecraft includes redundancy for critical functions such as control and data handling (C&DH). The reliability of the system is greatly improved by cross-strapping the various elements of the system—so that even if data storage fails on Side A, and the power system fails on Side B, the system continues to operate with data stored on Side B and power from Side A.

Caution is advised in implementing cold spares, since the need to be able to switch elements in and out introduces the need for additional switching hardware, which can

have reliability problems of its own. Radiation hardening of a system does no good if it reduces the system's overall reliability.

4.2 Hardening techniques for Nondestructive SEE

If mitigation of destructive SEE suffers from a paucity of good strategies, the opposite is true for mitigation of nondestructive SEE. In part, this is because the consequences of many nondestructive SEE are similar to those of noise in signal processing, and techniques for dealing with these errors are well developed. In addition to the voting schemes discussed in a previous session of this short course by Fernanda Kastensmidt, there are also error detection and correction (EDAC) codes and techniques such as interleaving of data that maximize the EDAC effectiveness. To deal with SEFI there are also techniques borrowed from fault-tolerant computing, such as error counting, watchdog timers and so on. Here we first consider techniques for dealing with the loss of functionality caused by SEFI. We then discuss techniques for hardening against SEU, MBU and the data loss resulting from SEFI. Finally, we discuss strategies for SETs and the reasons why these transient errors represent a real threat to mitigation schemes.

4.2.1 Hardening for SEFIs

Although a SEFI generally results from an error in the control logic, they can occur in simpler devices such as the Analog Devices AD6640 12-bit A-to-D converter, which have no control logic. The only thing all SEFI have in common is that they stop or unpredictably alter the normal functioning of the affected device. In addition, the data contents of the affected device are likely to be lost. Recovery of normal operations after a SEFI may be as simple as refreshing the affected control bits that caused it, or it may require cycling power to the device. Here we consider mitigations that facilitate recovery of normal operations, leaving recovery of data to the next section. Restoring functionality after a SEFI depends on having the right infrastructure in place to detect an anomaly quickly when it occurs.

Once a SEFI occurs, the first step in recovery is discovery. Such errors may be obvious in memory devices, where they can corrupt all the data on the device. In this case, an error counter that keeps track of the numbers of errors corrected (see below) can

alert the system to the anomaly. Data format checks can also provide indications of anomalies.

SEFI can be less immediately obvious in processors, FPGAs and other devices which receive data at their inputs, perform a series of possibly time-consuming processes on the data and then output the results. In such a device, a SEFI may result in garbage at the outputs, but more likely the algorithm will crash or will not converge and will continue until it is stopped. An error counter on the output may not register such a problem for a long time. On the other hand, a watchdog timer, heartbeat or other fault detection strategy institutes periodic checks on the functionality of such devices and facilitates discovery of any problems. These checks can be instantiated internal to or external to the monitored device. In either case, it is critical that the monitoring be as independent as possible of the monitored process or hardware.

The cost of such monitoring is efficiency—time spent monitoring the health of data processing is time that cannot be spent processing data. In addition, the monitoring circuitry may itself be susceptible to errors, leading to spurious resets and loss of data. However, usually the overhead for such processes is not large, and they can be instantiated (at least externally) in SEE hardened devices.

Once the error has been discovered, the recovery process can begin. This may involve a series of escalating recovery measures—a partial refresh of the device, followed by a full refresh, followed if necessary by power cycling of the device—or it may involve going for the big guns (a power cycle) first. Rapid recovery requires understanding the various SEFI modes as they affect the system and having recovery strategies in place that can be carried out automatically. Once normal operations are restored, the recovery process continues, perhaps including restoration of recovered software and for reprogrammable FPGAs configuration information. Data too may be recovered—and we discuss the techniques for this below.

4.2.2 Mitigation for SEUs, MBUs and other data loss mechanisms

In trying to mitigate SEUs, MBUs, stuck bits and SEFIs, we are fortunate that data loss due to these error modes resembles that due to noise on a communications signal, and many of the same methods used in correcting data loss in communications can be adapted to mitigating SEE. Here, we will concentrate on voting schemes and on

implementation of EDAC and other mitigations that increase EDAC effectiveness. Both of these schemes use redundant information to identify and correct data errors.

4.2.2.1 Voting

In her section of this short course, Fernanda Kastensmidt covered voting schemes in some detail, particularly as they apply to hardening of FPGA designs. Here we look at how voting schemes can be instantiated at the system level. Voting schemes are appealing because they are intuitive. In a conversation, if we think we have misheard what the other person says, we may ask them to repeat their message, and if what we hear conforms to the initial message we thought we heard, we are satisfied, or if not, we may ask for a third repetition and so on. Voting to mitigate nondestructive SEE operates on a similar principle—minus the social embarrassment.

A triplicate voting system (Figure 4-3 a) compares the outputs of three identical devices bit by bit, relying on the fact that while each bit is equally vulnerable to upset, the probability of the *same* bit upsetting in two independent devices is very low. Even a worst-case SEFI can only corrupt all the bits on a single chip, and so is correctable by triplicate voting and the probability of two SEFIs is negligibly small. The downside of triplicate voting is that it involves over 200% overhead, and for large data words, voting each bit can lead to very complicated voting circuitry.

Other voting schemes can be practical in some applications. Temporal voting (Figure 4-3 b), in which the same device or data path is polled 3 times and the results stored and voted is effective against SETs. The penalty here is speed. Resampling can occur no faster than a system clock cycle, so speed is reduced by a factor of 3x—or more if transients last longer than a clock cycle [75]. If error rates are low, this speed penalty can be reduced by dispensing with the third sample if the first two agree. In principle, a temporal voting scheme can also work for complicated devices such as processors—for example, by having the calculation performed 3 times starting with the same input. (and configuration for programmable FPGAs). However, this scheme would only be suitable for results that were not time critical, and if a result is not time critical, it is usually better to perform the operation on the ground and send it to the spacecraft. One possible exception to this rule could be for exploration rovers where power is at a greater premium than time and autonomous operation may be needed over extended periods.

A strategy that represents a compromise between spatial and temporal voting is duplicate with retry, in which two hardware instantiations are compared and then resampled or recalculated if there is a mismatch. With this scheme, there is no time penalty incurred unless there is a mismatch. Clearly, the effectiveness of these latter two voting schemes relies on the system's ability to restore the original data after the SEE. In the case of persistent upsets (e.g. those in a memory device) this may not be possible, and only spatial voting is effective.




Figure 4-3. Three instantiations of voting. a) Spatial voting takes the majority of three separate instantiations of hardware. b) Temporal compares the output of the same device sampled three times. c) Duplicate with retry compares two instances of the same hardware, resampling only if there is a miscompare.

Voting schemes need not be majoritarian. A moving average is in some sense a voting scheme, in which influence of outliers is suppressed until they become persistent. The appropriate voting system depends on the error characteristics (both radiation induced and otherwise) one is trying to mitigate.

Voting is not a panacea. In addition to the high overhead mentioned above, all voting schemes have an Achilles heel in a common voter for all three paths. In addition, coordination of three paths to ensure voting doesn't take place too early or too late make it difficult to instantiate voting without at least some other common combinatorial logic. Despite these weaknesses and disadvantages, voting is often the best hardening option against data loss due to SEE. In other cases, there may be better options.

4.2.2.2 Error Detection and Correction and Supporting Strategies

Those who advocate use of EDAC find themselves in good historical company. Such methods were first advocated by Claude Shannon in 1948 [76] although the details had to

await other researchers. In EDAC, redundant bits are added to a data word to enable the system to detect and correct errors in the data (be they due to SEU or transmission). To understand how this works, consider the parity bit that is often appended to a command code. This bit is 0 if the number of bits that are 1 is even and 1 otherwise. The parity bit allows the system to detect (but not correct) any error that corrupts an odd number of bits. A more sophisticated EDAC code is that developed by Richard Hamming[77], a product of his frustration working with the temperamental computers of the early 1950s. If we consider a 4 bit word that can assume values from 0-15 and has data bits D0, D1, D2 and D3, we can correct any single-bit error and detect the occurrence of any double bit error (a so-called single-error-correct-double-error-detect, or SECDED code) if we add 3 error correction bits, E0,E1, and E2. The result is a Hamming (7,4) code, and the error bits are calculated as follows (where \oplus represents an exclusive OR operation).

$$E0=D0\oplus D1\oplus D2 \quad (4-1)$$

$$E1=D0\oplus D1\oplus D3 \quad (4-2)$$

$$E2=D0\oplus D2\oplus D3 \quad (4-3)$$

The resulting 7-bit word, D0D1D2D3E0E1E2 has 128 possible values, only 16 of which correspond to correct or uncorrupted data. (See Figure 4-4.)

Hamming also introduced the concept of the Hamming distance, the number of differences in the bits between two bit strings of equal length. The minimum distance between any two valid symbols in a Hamming (7,4) code is 3. Thus, with one bit flip, one can correct the error by returning the bit string to its nearest valid value. With 2 bit flips, one is equidistant between two valid values—one knows a double bit flip has occurred, but cannot correct it.

In general, if we have a block of n bits with k data bits and n-k check bits, we will be able to correct at most $(n-k)/2$ bits in error (for n-k even, or $(n-k-1)/2$ for n-k odd). While the principle of redundant information is similar for all EDAC codes, some are more sophisticated in the way they implement that redundancy.

Allowed Values for Hamming (7,4)						
Data bits				Check bits		
D0	D1	D2	D3	E0	E1	E2
0	0	0	0	0	0	0
0	0	0	1	0	1	1
0	0	1	0	1	0	1
0	0	1	1	1	1	0
0	1	0	0	1	1	0
0	1	0	1	1	0	1
0	1	1	0	0	1	1
0	1	1	1	0	0	0
1	0	0	0	1	1	1
1	0	0	1	1	0	0
1	0	1	0	0	1	0
1	0	1	1	0	0	1
1	1	0	0	0	0	1
1	1	0	1	0	1	0
1	1	1	0	1	0	0
1	1	1	1	1	1	1

Bit in Error for Hamming (7,4)			
Bit in Error	Eq. 4-1	Eq. 4-2	Eq. 4-3
D0	FALSE	FALSE	FALSE
D1	FALSE	FALSE	TRUE
D2	FALSE	TRUE	FALSE
D3	TRUE	FALSE	FALSE
E0	FALSE	TRUE	TRUE
E1	TRUE	FALSE	TRUE
E2	TRUE	TRUE	FALSE

Figure 4-4. Allowed values and Error indicators for Hamming (7,4) EDAC.

Reed-Solomon (or R-S) codes [78] are frequently encountered in spacecraft data handling system, and are also commonly used to preserve data on compact disk recordings against scratches and other blemishes on the disk. R-S codes partition the data as blocks or sets of m-bit symbols rather than bits. If our data have k symbols, we can view them as points defining a polynomial of degree k-1 or less. Our check bits are determined by oversampling this polynomial. Because they act on symbols rather than bits, R-S codes treat multi-bit errors within the symbol the same way they treat single-bit errors (that is, a multi-bit error misses the polynomial just as does a single-bit error). This makes R-S codes efficient for correcting MBUs, just as they are efficient at correcting burst errors in communications. However, there is a limit. An error that spans more symbols than the code can correct—e.g. a SEFI—results in system-level errors.

4.2.2.3 Enhancing EDAC capability

EDAC codes can be an efficient mitigation against complicated data error, but they are limited in the number of bits or symbols they can correct. Fortunately, we have already encountered a strategy that can increase EDAC effectiveness: interleaving. Here

we must interleave the bits/symbols in such a way that even the worst-case SEE—one that corrupts every bit on a single chip—does not overwhelm EDAC capabilities. This means that we can store no more bits of any data word on a single chip than could be corrected by EDAC. While this complicates system architecture, it does provide immunity to SEFI and nondestructive SEL without resorting to a voting scheme. Indeed, if we can restrict the number of bits/symbols from any word to half the number that can be corrected by EDAC, the system becomes immune to any two SEE affecting—a very low probability occurrence.

In our discussion of EDAC codes we have drawn considerably on the similarities between correcting SEE in data storage systems and correcting corrupted bits in a communications data stream. However, there is one additional threat that a string of bits sitting in bulk memory faces—the accumulation of corrupted bits due to SEE over time. The threat here is that if errors accumulate and a SEFI corrupts a large number of data words, the words affected by both a SEFI and a SEU may not be correctable. In order to counter this threat, it is necessary to scrub the memory for errors periodically. However, usually, the required scrub rate need not be a burden on system efficiency to keep system error rates acceptably low.

4.2.2.4 Which Strategy: Voting or EDAC *et al.*

In the previous three sections we have examined the effectiveness of voting and EDAC in correcting data errors. In this discussion, we saw that triplicate voting is a powerful technique with broad applicability, but that it has a high cost in terms of system overhead, power, weight, etc. On the other hand, EDAC, especially when supported by interleaving and scrubbing, was found to be a very efficient strategy for mitigating SEE induced errors in memory arrays. Unfortunately, EDAC is generally not very effective for other types of ICs such as processors and FPGAs. In such devices (see Figure 4-5), data are input to the device, and a series of operations are performed on the data, and eventually the data arrive at the output. There is usually no external access or monitoring of the data during this process, and the operations involve the entire data word, so interleaving is not possible. Moreover, since SEFI can alter not just the data, but the operations on the data, it is unlikely that any strategy implemented internal to the processor other than triplicate voting can be effective.

In the case of reprogrammable FPGAs, the situation is even more difficult, since hardware as well as software and data can change as a result of upsets to configuration memory. Here, however, scrubbing of configuration memory can be effective.

Thus, for memory arrays, the best strategy if the inherent hardness of the part is not adequate is usually EDAC, supported by interleaving and scrubbing. If the EDAC is sufficiently robust to support more than one worst-case SEFI, the array will be as hard or harder than a triplicate voting system.

For processors, FPGAs and other computationally intensive devices, if the rate of data loss is not a concern, it may be sufficient simply to monitor of device functionality using watchdog timers, error counters, etc. Scrubbing of configuration memory in reprogrammable FPGAs is also essential. If data loss is a concern, triplicate voting is likely to be the most effective strategy. Finally, if the errors requiring mitigation are transient, temporal voting can be effective. In the next section we consider other strategies for mitigating transients.

4.2.3 Mitigating Single-Event Transients

For errors that correct themselves, SET generate a lot of problems. This is because increasing operating speeds of current microcircuits make it increasingly likely that an SET will be captured as an error during its brief sojourn in our electronics. Almost all of the mitigation techniques for SET carry a significant speed penalty. Temporal voting requires sampling outputs of SET susceptible devices three times on a timescale long compared to transient durations. Capacitive filtering slows the output of the device to the point where the transient is insignificant, just as does running at diminished speed to reduce susceptibility to transients.

Such a speed penalty is not acceptable in applications such as SEL detection circuits, where rapid response can make the difference whether a circuit is protected from thermal failure or not. In such situations, it may be possible to capitalize on dependence of SET susceptibility on application conditions—for example, the dependence of the National LM139's SET rate on the input voltage difference for the comparator [79], [80], [81]. (See figure 4-6.) Increasing the value of ΔV_{in} from 0.13 V to 1.17 V (with supply voltage of 13 V) reduces the SET rate by a factor of more than 50, to about one transient in 50 years per device. However, not all devices exhibit such strong application

dependence. If neither capacitive filtering nor opportunistic hardening is feasible, a triplicate voting circuit may be a solution—albeit a challenging one, given the speed and SET hardness required.




Figure 4-6. The SET threshold of the National Semiconductor LM139 comparator exhibits a strong dependence on the input voltage difference. (Supply voltage was 13 volts, adapted from [79])

Recently, several trends have increased concern over SETs. The millisecond-long SET discussed by Ben Blalock are one example [60] seen in the LM6144 and other op amps. Such transients make it very difficult to harden a circuit with capacitive filtering, and the slowness of a temporal voting scheme for such transients would preclude its use in many applications. Moreover, the mere existence of devices exhibiting such long transients highlights the risks involved in relying on “rules of thumb” rather than device specific test data.

While some devices are exhibiting very long transients, the increasing speeds of SOTA CMOS, SiGe and other technologies are making them sensitive to ever-shorter transients [58], [82]. Not only does sensitivity to such short transients raise questions about devices formerly thought to be SEE immune (e.g. RF amplifiers), SET on clock lines in such fast devices can give rise to burst errors several bits long.

Still another area of concern regarding SET is the possibility that they could damage increasingly sensitive devices. The Actel RTAX-S data sheet cautions that the supply voltage should not exceed 1.8 V at any time, and should not exceed 1.575 V for more than 10 μ s [59]. Thus, SETs from voltage regulators have for the first time become a reliability concern.

4.3 Mitigation for Degradation Mechanisms

Other than the usual strategy of adding spot shielding, mitigation of degradation mechanisms presents rather limited options. In some cases, the application dependence of TID degradation may be helpful in this regard—e.g. using higher I_F to facilitate photobleaching of color centers in passive optical components in optocouplers. In other cases, it may be possible to compensate for damage as it occurs.

However, when there is no strategy for reducing the damage a part sustains due to TID or displacement damage, the radiation analyst can still work with designers to ensure that the system can sustain the degradation. This can be achieved by inferring the degraded parametric function the design must accommodate or by estimating a safe RDM that achieves the desired success probability.

As an example of the former, SDO had a need for an op amp in its battery-charging unit. We consider the Linear Technologies RH1014. Because the application has high impedance, an increased leakage current ΔI_{bias} of 15 nA would give rise to an unacceptably large voltage error. A test sample of 8 parts yielded the data in Table 4-1:

Table 4-3 Changes in I_{bias} vs. Dose for two Lots of RH1014s

Dose	60 krad(Si)	100 krad(Si)	200 krad(Si)	Failure Level krad(Si)
ΔI_{bias} -P1 (nA)	10.03	10.631	22.052	128.3
ΔI_{bias} -P2 (nA)	11.86	13.579	26.184	99.3
ΔI_{bias} -P3 (nA)	10.504	13.553	25.934	105.3
ΔI_{bias} -P4 (nA)	9.901	12.51	24.062	115.3
ΔI_{bias} -P5 (nA)	9.768	11.723	22.654	123
ΔI_{bias} -P6 (nA)	8.592	15.198	28.844	102.2
ΔI_{bias} -P7 (nA)	9.548	16.396	31.16	93.6
ΔI_{bias} -P8 (nA)	10.217	15.203	28.757	97

Table 4-2 Summary Statistics

	Mean failure level	Standard Deviation	Lognormal mean= μ_{ln}	Lognormal Standard Deviation= σ_{ln}
Aggregate	108	12.7	4.67	.115

The project uses 10 parts in this application and wishes to achieve 99% success probability with 99% confidence. Figure 4-7 shows the distribution of failures inferred from these data. Applying the one-sided tolerance limits discussed in section 2.2.1 above shows that we achieve this at a dose of:

$$D = \exp(\mu_{ln} - K_{TL}(8, 99.9\%, 99\%) * \sigma_{ln}) = \exp(4.67 - 7.55 * 0.115) = 45 \text{ krad(Si)}$$

For SDO, this corresponds to a shielding level of about 180 mils Al equivalent shielding.

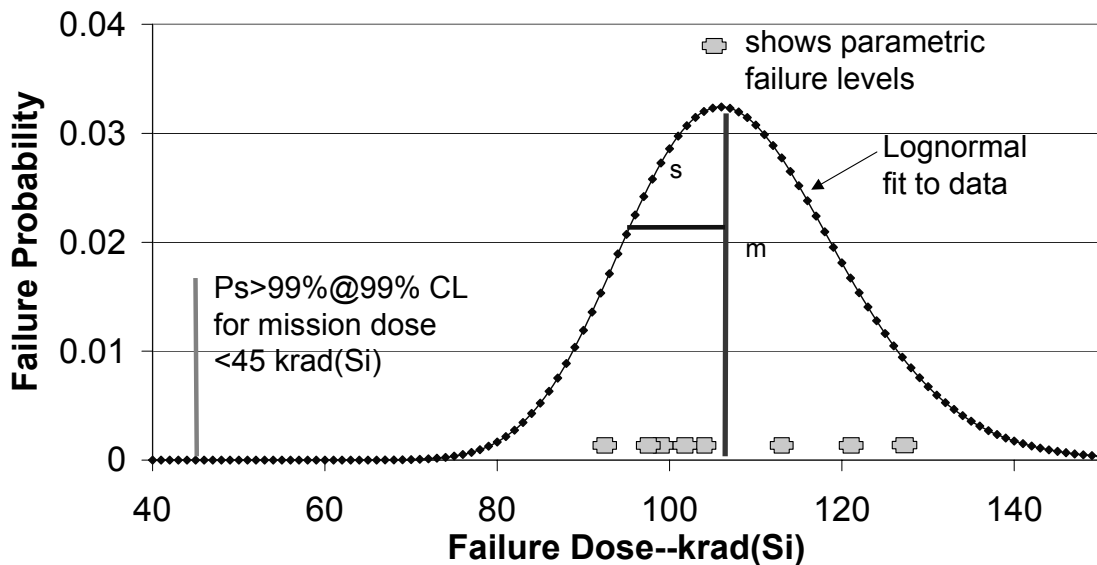


Figure 4-7 Failure levels and inferred failure distribution and dose level for 99% success probability @ 99% confidence for RH1014 op amps in a battery voltage monitor for SDO.

Now, consider the Micropac 53272 MOSFET optocoupler, which is used in several applications on SDO. For a forward current of 10 mA, we wish to know by what factor we can expect the charge transfer ratio (CTR) to degrade after exposure to a damage dose equivalent to 2×10^{11} 50 MeV protons/cm². Figure 4-8 shows the fractional degradation in CTR as a function of 63-MeV proton fluence. The main culprit in this degradation is

the GaAsP LED. Since damage does not follow the total NIEL for III-V semiconductors, we use the damage function from reference [83] and find that a fluence of 2.14×10^{11} 63-MeV protons/cm² is equivalent to our desired damage level. For this fluence, the degradation distribution of the 4 sample parts fits a lognormal with mean -2.0 and standard deviation 0.119 . Again, using one-sided tolerance limits, we find that 99% of the optocouplers will retain at least 7% of their original CTR at the desired fluence. If this level of degradation poses a problem, using $I_F=30$ mA—the highest value measured for the test yields a 99/90 fractional CTR of 26% for the same fluence—albeit at 3 times the power.

Mitigation for degradation strategies usually involves predicting the damage the part is likely to sustain at end of life and working with designers to ensure that the application can continue operating even for that level of damage. If the damage sustained for a given dose is problematic, spot shielding can be added to decrease the expected damage by decreasing the stress (TID or DDD) the part is exposed to. In some cases, application conditions can be manipulated to minimize damage or compensation can be built into the system to mitigate the damage.

Often for commercial devices, however, the most important system-level question that will arise in dealing with TID is ensuring that flight performance can be predicted reliably from test data. This is because lot traceability can be difficult to obtain for commercial devices. We will discuss this in more detail in the next section as we consider the qualification of SDRAMs used in flight data recorders for SDO and other projects.




Figure 4-8 Proportion of CTR remaining as a function of proton exposure.

4.4 Validation

Once mitigation has been implemented, it must be validated. Whether validation should be by test or analysis depends on the consequences the mitigation addresses and whether system-level testing can be implemented effectively. The most likely mitigations to require validation by test are those for detection and protection circuits for destructive SEL, because of the severe consequences if the mitigation should fail and because it is difficult to establish confidence in a model of a worst-case SEL or the failures or latent damage that might result from it. If cold spares are being used to mitigate destructive failures, the strategy for isolating the failed part and swapping in the redundant part must be validated. SET mitigation may also require validation by testing, since SET can assume a broad range of forms, and the validation can often be done using a pulsed laser system rather than heavy-ion irradiation. Testing may also play an important part in validating that compensation circuitry to mitigate TID or DD degradation will perform as intended. However, since there is a nonzero probability that flight parts could perform worse than test parts, analysis is also necessary.

In many cases, since the mitigation has been tailored to the best understanding of the radiation threat based on the available data, performing additional radiation testing would be uneconomical and would add little confidence in the mitigation in any case. In such cases, analysis and modeling will be the validation techniques of choice. Fault-injection has proven an excellent technique for identifying weaknesses in mitigation in a controlled environment. It can be used to test error correction as well as testing watchdog timers, error counters, checksum and other monitoring techniques to detect anomalous behavior. Circuit level simulations and tests can validate whether designs can operate at end of life by simulating likely degradation. In some cases, especially for very complicated systems, a proton, neutron or heavy ion test on vulnerable components with simulated mitigation in place may be needed for full confidence to be established.

5 PUTTING HARDENING INTO PRACTICE: AN EXAMPLE

Early in the planning stages of the James Webb Space Telescope (JWST), it became clear that the spacecraft would need a large amount of memory for data storage—a need that was also shared by other projects, including SDO (see section 7). The large amount of volatile memory necessitated the use of SDRAMs for the solid-state recorder (SSR). The SDO requirements were particularly stringent, in that some of the most interesting science for the project would occur during solar particle events. JWST, on the other hand, is precluded by the sensitivity of its detectors from taking data during inclement solar weather. To minimize part qualification costs, it made sense to seek a common solution across flight projects for the memory needs.

SDO needed 3 Gbit of usable memory, organized in a $\times 8$ configuration. The board space for the memory was very limited, so the project wanted to use parts with six 256-Mbit SDRAM die stacked in a single package. The $\times 8$ configuration makes it difficult to interleave fewer than 8 bits per 32-bit word on a single die. If we require the system to be able to preserve data quality through one worst-case SEFI or nondestructive SEL, then we need to be able to correct 2 nibbles simultaneously. This means we need to use a robust EDAC such as a Reed-Solomon (48,32) code, which uses 16 EDAC bits (4 nibbles) to correct up to 2 nibbles in any word. This brings the SSR memory size up to 4.5 Gbits, or three 6-die memory modules.

At the time the effort started, the latest generation SDRAM was 256 Mbits, and most parts in this generation had a high SEL rate [84], [85] that precluded their use in flight. SEFI were also common. The best candidate in terms of radiation performance was the Hitachi (later Elpida) HM5225805B. While this part did exhibit SEL, it did so only at elevated temperatures (>50 °C), and the SELs observed to that point had been nondestructive, and recovery from most SEFIs could be accomplished by refreshing mode registers. However, there remained several issues. TID performance had been variable. It was not known whether all SELs were truly nondestructive or whether rare destructive modes had been masked by more common nondestructive modes, and the issue of latent damage had never been addressed. Lot-to-lot variability of stuck bits was another open question. The TID issue was particularly difficult because SDRAM

vendors must compete in a high-volume, low-margin market and issues such as wafer-diffusion-lot traceability are not part of their business model.

Moreover, funds for testing to resolve these issues did not permit a full-blown testing campaign. The solution the projects arrived at was to work with a value-added parts supplier, Maxwell Technologies, to obtain satisfactory resolution of the above issues. Maxwell projected sufficient future demand that they were able to resolve the lot-traceability issue by purchasing a full wafer lot and packaging them in their RadPak™—either as single die or stacked multiple die to a package. This not only answered the problem of lot traceability, it supplied enough shielding so that for SDO and JWST, the parts would see lower doses than the lowest failure level previously observed.

The issues that still remained to be resolved were—in order of decreasing risk posed to the application:

- 1) Does the die exhibit destructive SEL modes and/or modes that cause latent damage?
- 2) Now that we have lot traceability, is TID performance acceptable in the application?
- 3) Are stuck bits likely to accumulate sufficiently that error mitigation is compromised by end of life?
- 4) Do nondestructive SELs, SEFIs, MBUs occur at sufficiently high rates that they are likely to overwhelm error mitigation, especially in light of stuck bit accumulation?

5.1 Destructive SEL and Latent Damage

The issue of destructive SEL was critical for the application primarily because there were no good mitigation strategies. Even if a destructive SEL could be mitigated using SEL detection and protection circuitry, it would preclude using stacked die, and space was at a premium on the data storage board. If stacked die could not be used, a major redesign would be necessary. For this reason, the project decided to resolve the issues of destructive SEL and latent damage by testing. Both Maxwell and Goddard Space Flight Center undertook independent heavy-ion irradiations of SDRAM die (note: Figure 2-5 is a product of one of these studies). In addition, GSFC undertook laser testing that pinpointed the SEL susceptible sites in the control logic and stimulated hundreds of SELs. After these tests, die were tested functionally and examined for microscopic evidence of latent damage. In addition, die were subjected to 1000 hour burn-in post SEL with no failures seen.[40] Throughout the testing, all SELs—at temperatures

ranging from 50-125 °C were observed to be nondestructive and recovery could be achieved by cycling power to the DUT.

5.2 TID and RLAT

As mentioned above, the packaging of the SDRAM die reduces the TID the die see (~7 krad(Si)) to well under the lowest failure level seen. During RLAT, one part failed functionally between 42 Krad(Si), while the remainder were fully functional when the test was suspended at 50 krad(Si). Since the lot is the flight lot, an analysis assuming a lognormal form for the failure distribution shows we have $P_s=99\%$ with 90% confidence at 14 krad(Si)—giving us a 2x margin even on our 99/90 TID level. This suggests that TID hardness is adequate for the mission.

5.3 Stuck bits

Stuck bits are another mechanism for which there is no good mitigation. As such, we addressed this issue by testing to determine both susceptibility to stuck bits and how likely they are to anneal. The results showed fewer than 100 stuck bits forming in a typical heavy-ion irradiation, with over 90% annealing within an hour. This means that stuck bits are very unlikely to impact error mitigation during SDO's 5 year mission.

5.4 Information loss: Nondestructive SEL, SEFI, MBU and SEU

Although the requirement to operate through a solar particle event poses the greatest challenge, SDO will spend most of its time in orbit during periods of solar quiet. Therefore it is important to ensure that data can be reliably stored in the SSR for extended periods under these conditions. Since the consequences of a nondestructive SEL and a SEFI are the same in terms of interrupted operations and lost data (up to all data stored on a single die), we consider the effectiveness of four mitigation scenarios in terms of such errors, together in four mitigation scenarios. Scenario one includes no mitigation and produces a combined rate for all such errors of $4 \times 10^{-4} \text{ dy}^{-1}$ per die. This means that about 9 times during the 5-year mission, the SSR would experience an error that would result in the loss of its data.

Scenario 2 involves the addition of a Reed-Solomon (12 nibble, 8 nibble) EDAC, interleaving (8 bits per die) and daily scrubbing. This reduces the rate for such outages to 0.026 per mission.

If this rate were still too high, it would be possible to store only 4 bits per word on any die (Scenario 3). This would considerably complicate operations. For instance, it would necessitate reading and writing words 2 at a time onto the 8-bit die or redesigning the SSR to use a memory organized in a $\times 4$ configuration rather than an $\times 8$ configuration. This would reduce the error rate to 0.0002 per mission. We add error counting to our list of mitigations, so that a SEFI is discovered in less than 3 hours, rates for the EDAC protected system decrease by another factor of 8 (because the additional SEFI(s) would have to occur in 3 rather than 24 hours).

For scenario 4 we look at the efficacy of a voting scheme (again with a once per day scrubbing requirement), which yields a rate of 0.021 large data losses per 5 year mission. This is only slightly better than Scenario 2, and it requires twice as much memory to implement (six 6-die stacks as opposed to 3 for scenario 2).

We can further reduce the error rates for scenarios 2, 3 and 4 by another factor of 8 if we add error counting so that a SEFI is discovered in 3 hours rather than 24 hours. This comes with a very small cost to the system and has the benefit of providing a way of monitoring system health.

For a data storage system, the resulting 0.325% chance of a large data loss is acceptable, so we choose scenario 2 as the most cost-effective strategy.

In addition to these large errors, there will also be small amounts of data lost whenever a SEFI occurs—roughly 6 corrupted bits per SEFI occurrence. Since these errors scale with the number of memory die, this rate will scale with the overhead needed to implement the mitigation. Thus, we expect fewer than 15 such small data losses during the mission.




Figure 5-1. An unprotected 3 Gbit memory array (Scenario 1-black) would experience roughly 9 outages in the 5 year SDO mission. The addition of a Reed-Solomon (48,32) EDAC with a scrub rate of once a day reduces the error rate to 0.026 outages per mission (Scenario-2—dark gray, with 8 bits per word interleaved over different die) or 0.0002 per mission (Scenario 3-light gray, 4-bit interleaving). A triplicate-voting scheme with daily scrubbing (Scenario 4—crosshatched) has an outage rate of 0.021 per mission.

The question of outages during solar particle events (SPE) is more challenging. Initial determinations of SEE rates during SPEs are roughly a factor of 1000 higher. This means that there is a high probability of 2 SEFIs affecting the same word during a CREME96 worst-case SPE, causing a large data loss. Fortunately, solar heavy ions have a soft spectrum, so rates may be significantly lower for our well-shielded SDRAMs than the initial estimate. Figure 5-2 shows the SEFI rate trend vs. shielding, where for 500 mils Al equivalent our rate is down roughly 2 orders of magnitude and back in the acceptable range (2% chance during the 24 hours of the CREME96 Worst-Day Environment). Data loss would also be minimal, with a 29% chance of a SEFI occurring and causing the loss of about 10 words of data.

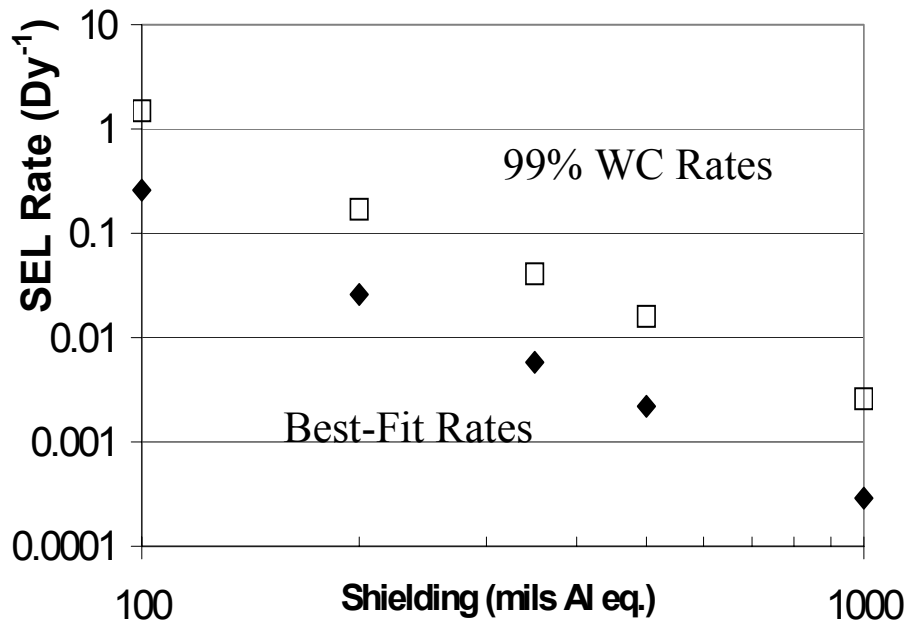


Figure 5-2. Because of the soft solar heavy-ion spectrum, SEE rates for solar particle events decrease rapidly as shielding levels increase.

5.5 Summary

The analysis above illustrates the system-level hardening for a system with a broad range of radiation effects. We worked with designers and system engineers to understand the requirements for the solid-state recorder and to ensure that they understood the radiation threats to which the system was susceptible. We also discussed potential mitigations for these threats, their underlying assumptions, effectiveness and possible costs to the system in weight, power, performance and resources.

Initially, there were many uncertainties about whether the parts would be suitable. In particular, questions about destructive SEL and latent damage, stuck bits and lot traceability and variability for TID were significant obstacles. On the other hand, we were confident that we could effectively mitigate the data loss and outages resulting from SEU, MBU, SEFI and even nondestructive SEL. Moreover, if we could resolve the issue of lot traceability, we were confident we could shield the parts adequately to meet mission TID requirements.

The purchase by Maxwell Technologies of an entire wafer lot of die resolved the lot-traceability issue, as well as providing a packaging option that added significant shielding for the parts (a 2x margin even on the 99/90 TID level).

Since there are no good mitigations against destructive SEL or stuck bits, both Maxwell and GSFC performed additional testing. Heavy-ion and laser testing followed by inspection and burn-in life testing found that all SELs exhibited by the device were nondestructive and did not cause latent damage. Similarly, stuck bits were found not to accumulate to levels of concern, because they anneal nearly as quickly as they form. The elimination of these concerns, the mitigation of which could have forced a redesign, allowed mitigation efforts to focus on mitigating against data loss and ensuring data quality.

The choice of a $\times 8$ configuration for the memory meant that we had to choose robust EDAC such as Reed-Solomon (12 nibble, 8 nibble) in order to correct a worst-case SEFI. This, along with memory scrubbing and interleaving no more than 8 bits per die, is the minimum level of protection that will be effective against all SEFI and nondestructive SEL. The resulting system achieves hardness to data loss nearly equal to a triplicate voting scheme with a memory array half the size. The addition of an error counter that would discover the occurrence of a SEFI more rapidly would further limit the probability of data loss with little additional cost.

The remaining issue of outage rates during a SPE was resolved fortuitously by the fact that the high shielding for the SDRAM die coupled with the soft SPE heavy-ion spectrum gave rise to a much lower SEFI rate than would be predicted using a nominal 100 mil Al equivalent shielding. If this had not been true, our mitigation strategy would have broken down—whether we had chosen Scenario 2 (R-S EDAC + interleaving + scrubbing) or Scenario 4 (triplicate voting + scrubbing)—because the error rate would have simply been too high for mitigation to work effectively. In the next section, we examine the conditions under which various mitigations break down

6 WHEN MITIGATION BREAKS DOWN

The discussion of radiation effects in section 3 noted that different radiation effects have different spatial and temporal characteristics. SET usually affect only a single output and last only a short time, and SEU while they persist indefinitely until corrected still affect only a single bit. In contrast, degradation mechanisms can have permanent global effects—causing degradation in primary and redundant parts alike. Many of the mitigations discussed in section 4 make explicit or implicit use of these spatial and temporal characteristics. When these assumptions are violated—e.g. when a SET lasts a long time or fans out to affect large amounts of surrounding circuitry—then mitigation techniques can break down or at least lose much of their effectiveness. If this is discovered during testing, the impact can perhaps be limited. However, if the mitigation is based on a “rule of thumb,” rather than testing or analysis, the vulnerability may lead to anomalies during the mission.

Similarly, if the characteristic on which an opportunistic mitigation is based changes, the mitigation may also break down. Since opportunistic mitigations are rooted empirically observed trends (that is, test data), such breakdowns only occur if the behavior in flight is different from that observed in testing. This can happen if the test parts are not representative of the flight parts or, if flight application conditions deviate from test conditions or if the behavior of flight parts changes on orbit (e.g. due to TID degradation, aging etc.). It is not usually possible to shore up opportunistic mitigation strategies—they work or they don’t.

Error correction strategies based on redundancy can also break down. Since redundant strategies necessarily increase the number of parts (or at least the cross section) being used, they fail if the error rate climbs to the point where errors are likely during the period between scrubs. This can happen if the test parts on which the mitigation was based are not representative, or if the error rate increases on orbit—as happens when a FLASH memory or SDRAM degrades due to TID at end of life. Although increasing the scrub rate may shore up the mitigation, this reduces the efficiency of the system.

Likewise, mitigations based on infrastructure (software or hardware error monitoring, overcurrent detection systems, etc) will also fail if they are based on faulty testing or if the behavior of parts changes on orbit.

Based on the above discussion, we can divide threats to mitigation into 5 broad categories:

- 1) Violations of assumptions about the radiation response being mitigated
- 2) Issues with fidelity of test conditions to application conditions
- 3) Issues with whether the test sample was representative of mission components
- 4) Synergistic effects causing parts to behave differently in the mission than the test
- 5) Threats to the cost effectiveness of risk reduction

In this section, we examine each of these threats to the integrity of system-level mitigation, suggest safeguards for avoiding them and look at current trends that may exacerbate the threat. We end with a brief assessment of the prospects for system-level hardening in coming years.

6.1 Breakdown of assumptions

Although all system mitigations are based on radiation test results, the high cost of radiation testing often leads designers to base mitigation on a “worst-case” threat rather than tailoring the mitigation to the performance of a specific part. Unfortunately, what has been worst-case for parts used in the past is not necessarily worst case for all time, Transients are an excellent example of a threat the rapid development of the field makes it difficult to generalize about the threat. The concern generated by the discovery of millisecond-duration transients in the LM6144 op amp provides an example.[60]. Such long transients decrease the periodicity for re-sampling to rates less than 1 kHz and a temporal voting scheme would have to operate at less than half that speed. Such rates would be unacceptable for many applications.

Moreover, what constitutes a “long” transient depends on the system. Figure 6-1 shows a scatter plot of error magnitude vs. error duration for the 8-bit 1-GSPS MAX108 A-to-D converter [86]. Although even the longest errors last under 300 ns, at 1 GSPS, this represents nearly 300 corrupted samples. Using a moving average to mitigate transient errors would not work for this device unless the period covered roughly 1000 samples. Although the periodicity of the moving average is usually handled by software and so is reprogrammable, discovery on orbit that such long periodicity is needed could impact mission requirements for some programs. The only way to avoid such unpleasant

surprises is to base the mitigation on test data specific to the part and system under consideration.




Figure 6-1. Maximum Error Magnitude vs. Event Duration under irradiation by Au ions with incident LET=83.9 MeVcm²/mg [86].

The assumption that a transient's effects will be localized to a single device or output can also be violated. This was seen for upsets to clock distribution circuits in a SiGe test chip [82]. Fanout from a device that experiences a transient can spread the effects of the transient throughout a system.

Another area where assumptions can break down and cause trouble has to do with heritage designs and archival test results. The increasing speed of CMOS and other technologies means that such devices can be vulnerable to transients lasting under 1 ns. If a device such as an rf amplifier was tested a decade ago, such short transients might have been missed, and the device could be considered immune. Use of the same device in current hardware would introduce a SET threat for today's high-speed digital logic. Even the assumption that transients are recoverable cannot be taken for granted given the potential for damage to input circuitry from overvoltage [59].

However, vulnerable assumptions are not confined to SETs. Any new technology, or any old technology used in a new way, can challenge conventional wisdom. As shown by the failure of the IRF640 200 V MOSFET at $V_{DS}=44$ Volts, one cannot take for granted that commercial MOSFETs will exhibit the same hardness to SEGR and SEB as their radiation hardened counterparts [43]. The use of MOSFETs at cryogenic temperatures on the James Webb Space Telescope necessitated a challenging cryogenic SEGR test to rule out that threat. The failures of bipolar ICs due to second breakdown discussed in section 3.2.4 shows that SEL in CMOS is not the only destructive failure mechanism in microelectronics with which we need to be concerned. In recent years, even the assumptions that ELDRS is an issue only for bipolar technologies [87] and that CMOS technologies are not prone to bulk damage [88], [89] have been challenged. Finally, as the data in Figure 3-7 show, we cannot be complacent that the amount of interleaving in memories will necessarily prevent the occurrence of MBU—particularly given the much higher energies of ions in space compared to those in the lab. Although these threats do not jeopardize hardening for current systems, the rapid pace of change in electronics means that the status quo can never be taken for granted.

6.2 Test Fidelity

Although current testing methods are adequate for most applications, achieving 100% fidelity to all applications is not feasible. The inability of commonly used heavy-ion accelerators to produce ions at GCR energies is one such shortcoming. Although it is unlikely that this shortcoming has adversely affected system-level mitigation strategies, continued shrinking of feature sizes means that the fidelity of future testing will depend increasingly on the effectiveness of interleaving and other strategies of avoiding MBU at the part level. The concern over the discrepancy between test and space energies is further raised by the increasing use of packages with several die stacked vertically. For such packages, a single ion could cause upsets or SEFI in multiple die within a stack. Figure 6-2 shows the LET of three of the most common heavy ion species (iron, neon and oxygen) vs. the residual range for GCR ions incident on the package with the peak GCR energy (~ 1 GeV per amu). Clearly, for large memory arrays, the risk of such multiple SEE becomes less negligible as the onset LET for SEFI and other serious error modes

decreases. It is possible that modeling could provide the additional understanding required if the vendor supplies logical bit map or if it can be reverse engineered.

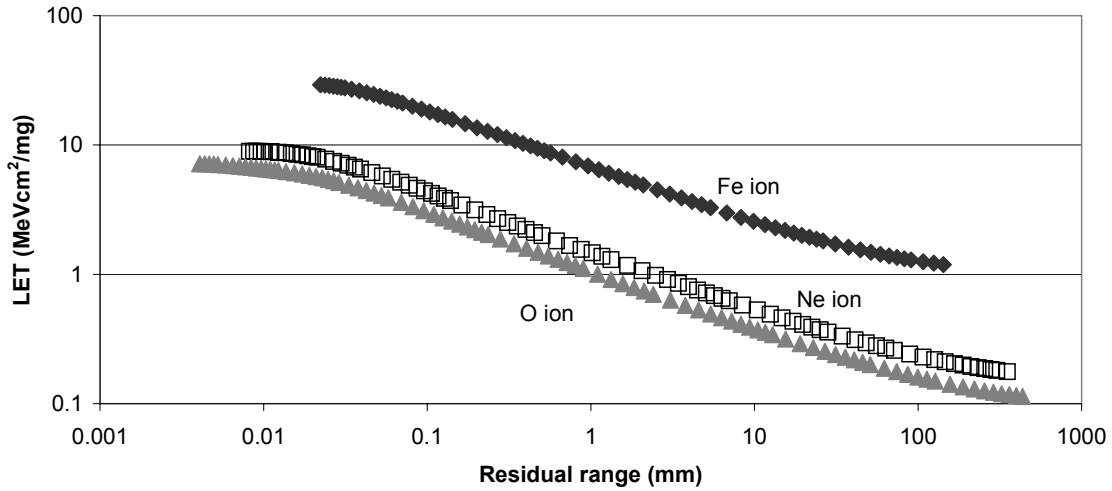


Figure 6-2. LET vs. residual range in Si for Fe, Ne and O ions incident with energies at the peak of the GCR spectrum (~ 1 GeV per amu) (Calculated using SRIM.)

ELDRS testing remains another area where 100% fidelity is impossible to achieve given the very low dose rates in most space environments. To date, while testing has revealed continued worsening of degradation for a few devices as dose rates continue to fall, the differences have been quantitative rather than qualitative [66]. If there is concern about a particularly sensitive parameter, most cases could be handled by assuming an additional $2\text{-}3\times$ RDM on tests conducted according to MIL-STD 883 Method 1019.8.

Another barrier to achieving fidelity to applications is the sheer volume of testing required. The example of the DDR SDRAM's 68 modes of operations provides an indication of the complexity of many microelectronic devices. Even relatively simple analog devices may be difficult to test for all application conditions, and SET susceptibility can be very application dependent. For some applications, laser testing and modeling may increase the level of fidelity if required.

6.3 Ensuring Representative Test Samples

In section 5, we briefly discussed the difficulty of ensuring lot traceability for commercial parts. The reason why this is important is that without lot traceability we have no way of knowing whether our test sample is representative of the flight lot. While

this is most important for TID, some SEE—especially destructive mechanisms—have varied significantly from lot to lot. Without lot-traceability, we may have to measure the full variability of the parts across wafer diffusion lots, and this would require very large sample. Although there are solutions to the lot-traceability issue, as discussed in section 5, there are other threats to ensuring that our test samples are representative of flight parts. Usually the lifecycle of a commercial part is on the order of 2-3 years—especially for high-volume sectors such as memory. By the time a project defines its need for such a technology and gathers enough test data on the part to establish confidence in its probability of success, the parts available will probably be a new die revision—and for the purposes of radiation hardness, a new part. The project then must decide between re-testing and making do with this new revision or generation or trying to obtain the old version. Since the performance of the new part may or may not meet specifications, the temptation will be to try to obtain the old part.

Unfortunately, this opens the project up to another worsening problem—that of counterfeit parts. According to a recent briefing by David Meshel of the Aerospace Corporation [90], the Electronic Resellers Association International, Inc. has identified 2857 independent brokers selling counterfeit parts. These parts may have correct markings, pass initial screenings and even have counterfeit certificates of compliance. Counterfeit parts may be recycled from old computer boards, may be rejects given new markings or they may be fresh from a fabrication facility in a country where regulation is limited.

On the other hand, if the project chooses to use the new version of the die, the radiation performance may not be adequate, and the search will have to start over again. One way to mitigate these problems is to purchase an adequate supply of parts for the program before testing or, if possible, make purchase conditional on radiation test results.

Another question that affects how representative test samples may be of flight parts is whether new materials, device structures and manufacturing techniques affect part-to-part or lot-to-lot performance. As feature sizes continue to shrink, it is virtually certain that there will be greater variability from transistor to transistor [91]. Will the small sample sizes typically used for radiation testing (4-10 devices for TID and 2-3 for SEE) be able to capture adequately this variation? Will the migration of CMOS circuits toward non-

native oxides introduce new sources of variability? These are questions that will be confronted in new generations of commercial parts very soon, and they can only be answered empirically.

6.4 Synergistic Effects

The investigation of synergistic effects in radiation is still fairly new. Several studies have looked at the impact of aging on TID hardness [92], [93], [94], [95]. More recently some researchers have begun to examine whether SEE response could change at end of life as a result of degradation effects. References [88] and [89] are pertinent to the question of whether stuck bits are affected by degradation. Reference [96] looked in detail at whether TID exposure affects SET response of linear bipolar microcircuits. To date, the effects seen would not have a large impact on the effectiveness of most system-level hardening schemes unless margins were extremely tight.

6.5 Cost-Effective Risk Reduction

Having just discussed all the radiation threats a system-level hardening effort must address and all the potential complications that may result from future trends, this is where we get the bad news that we need to do it more cheaply. Density, integration and speed of parts being used on satellites are all increasing, and these raise the cost of testing—not to mention the cost and difficulty of test sample preparation for heavy-ion SEE testing. Moreover, since many parts are too complicated to thoroughly qualify, more and more tests are application specific, and test results for one system do not necessarily generalize to a system using different mitigation. While not surprising, increased cost of qualifying complicated parts may be a concern, since it means that projects must either increase their budget for radiation testing and mitigation, restrict themselves to more conservative technologies for all but the most pressing needs or accept more risk. The latter two options are obviously bad news for radiation specialists, but they are also bad news for satellite designers, since one places designers into a conservative straightjacket while the other increases the probability of failure. In order to ensure continued innovation, close cooperation is needed between designers and radiation experts to ensure that scarce resources achieve maximum reduction of risk. Strategies radiation engineers can follow to facilitate such efficiency include:

- 1) Early involvement—Although projects are sometimes resistant to early expenditures for radiation specialists, early involvement usually pays dividends. Unnecessarily risky technologies are avoided. Technology needs can be assessed and the project can determine which needs can be met with existing radiation-hardened or radiation-tolerant solutions, which ones require an RHBD ASIC or FPGA solution and which needs require special circuit-level or system-level hardening. For the latter, testing efforts can be tailored to validating and selecting among the best candidate components given the available mitigation strategies.
- 2) Cooperation among institutions—Although cooperation has long been a goal, it has been difficult to realize given the different technology needs, schedules and competitive sensitivities between organizations. However, given the nearly universal need for bulk volatile memory and some other needs, there has begun to be increasing cooperation. Different organizations bring different strengths to the table when it comes to organizing a testing campaign—ranging from part preparation to high-speed testing to relationships with vendors and semiconductor manufacturers and access to design information. Sharing information on promising parts does not necessarily diminish a company’s competitive edge, since even if one knows what part to use, realizing a design in which that part succeeds remains a challenge.
- 3) If cooperation can be extended, sharing of data on commonly used parts opens up the possibility of doing in-depth studies on lot-to-lot variability for both SEE and TID. For some parts it may be possible to replace RLAT with periodic testing while still maintaining the same level of risk.[23] This would allow more resources to be devoted to effective risk reduction .
- 4) Laser testing has proven a valuable technique for mapping out the complicated SET response of linear bipolar microcircuits for different application conditions. It may also prove useful in mapping the response of complicated digital microcircuits such as DDR SDRAMs.
- 5) Given the difficulty of fully qualifying a complicated state-of-the-art microcircuit, modeling is going to play an increasingly important role. Such an effort could require establishing a relationship with the vendor or reverse engineering of the part. Alternatively, it may involve detailed analysis of system response using techniques like fault injection.

6.6 Summary

Although we have identified several challenges above, the system hardening techniques we have discussed are robust. As long as the data on which we base our hardening are representative of the parts used in the mission, hardening at the system level can provide an economical solution to meeting system requirements with parts having marginal radiation response. There are several keys to realizing an effective hardening effort. First, the radiation analyst must have a good understanding of system

requirements so that he or she can translate those requirements into validation requirements that have a clear relevance to mission requirements while also being meaningful in terms of radiation test procedures. Once the requirements are in place, and testing reveals whether a component is likely to meet its system requirements, the radiation engineer works with system and design engineers to understand how critical the component is in the design. If a part cannot be replaced without major redesign, the analyst weighs mitigation options with the system-design team and tries to select the option that meets requirements most cost effectively. When such a strategy is feasible, it is then implemented. If it is not feasible, the analyst works with the project to assess the impact to requirements and develop recovery strategies. The final step in the process is validation of the system and mitigation strategy—whether by test or analysis.

Many system-level mitigation strategies predate the space age. They have been extended from communications, from fault tolerant computing and reliability, and applied effectively to radiation threats. Because they concentrate on mitigating system impacts as well as reducing failure probability, it is likely that they will continue to be useful even as new generations of electronic components are applied in future satellite missions.

7 CONCLUSION: CASE STUDIES

In trying to find suitable case studies for this short course, we were limited by several factors. First, commercial satellites consider design and radiation performance to be highly proprietary subjects. Strategic systems tend to be similarly off limits. Moreover, difficulty of access and NASA regulations precluded discussion of projects by foreign governments. Fortunately, NASA's research satellites are sufficiently diverse and rich to serve as examples—both positive and negative. For our part, we have chosen to emphasize the positive—choosing one mission that could serve as a poster child for its systematic approach to radiation hardening, the Solar Dynamics Observatory and one mission that has already been a stellar success, the Mars Rovers,

7.1 Solar Dynamics Observatory (SDO)

The Solar Dynamics Observatory (SDO) is part of NASA's Living With a Star program. Scheduled to launch in 2009, SDO's primary mission is to observe the Sun for at least 2 years with a variety of instruments to learn more about energetic solar phenomena and the Sun-Earth connection. Given NASA's renewed emphasis on manned space exploration, understanding such energetic phenomena will be crucial for accurate Space Weather Forecasting and the protection of crew and other valuable agency assets. Specifically, the science goals of the program are to:

- 1) Understand the Solar Cycle.
- 2) Identify the role the heliomagnetic field plays in delivering energy to the solar atmosphere and its many layers.
- 3) Study how the outer regions of the Sun's atmosphere evolve over time (on scales from seconds to centuries) and space.
- 4) Monitor solar radiation across from visible to extreme UV (EUV)

To fulfill these goals, SDO is equipped with a suite of three main scientific instruments:

The Atmospheric Imaging Assembly takes images spanning 1.3 solar diameters and over multiple wavelengths, with 1 arcsec resolution 10 second cadence.

The EUV Variability Experiment seeks to characterize the Sun's irradiance and its variability in the extreme ultraviolet, using a Multi-EUV-Grating Spectrograph and EUV spectrophotometer.

The **Heli seismic/Magnetic Imager** will study origin of solar variability, characterizing the Sun's interior and magnetic activity by means of a series of filtergrams in a set of polarizations and spectral line positions at a regular cadence throughout the mission.

Because much of the science of interest to the mission is related to periods of intense solar activity, SDO must not only survive periods of intense solar activity, but also continue to collect and record data. Given that during intense solar particle events, SEE rates may rise by a factor of 100 or more, this is indeed a challenging requirement. In addition, the moderately high data rate from the instruments necessitates a relatively large bulk memory for data storage that requires the use of large capacity SDRAMs, which are susceptible to many radiation effects

Neither is the TID environment for the spacecraft particularly benign—especially for lightly shielded applications. (See Figure 7-1.)




Figure 7-1. Dose vs. Depth curve for SDO, along with contributions from trapped electrons and solar protons.

In addition, since the instruments are being built by different contractors—from satellite builders to Universities—and must be integrated with hardware built by NASA, there is also the challenge of bringing consistency to organizations that may have very different RHA approaches. SDO has responded aggressively to radiation threats with both testing and mitigation and has found creative solutions to potential problems.

7.2 Exploration Rovers: Spirit, Opportunity and Beyond

In applying the adjective “plucky” to a spacecraft, a designer might be excused for blushing slightly. Yet, it would be difficult to come up with a better adjective in describing the success the Mars Exploration Rovers Spirit and Opportunity in surmounting obstacles both geological and technical. Few robotic missions have so captivated the public, perhaps in part because anyone who has played video games can envision themselves controlling the joystick. Perhaps somewhat less obvious to the public were the serious scientific goals and technical hurdles faced by such a mission.

The science goals of the rover mission were to:

- 1) Characterize soils and minerals for evidence of hydrologic activity
- 2) Investigate distribution of minerals, rocks and soils near the landing site.
- 3) Determine geology that has shaped terrain and chemistry
- 4) Calibrate and validate Mars Orbiter instruments
- 5) Identify and quantify minerals related to hydrologic processes
- 6) Characterize rocks and determine how they formed
- 7) Search for clues of conditions when Mars had water and whether these conditions were conducive to life.

In terms of instruments, the Rovers are equipped with a suite of instruments well suited to a geology field trip:

Miniature Thermal Emission Spectrometer (Mini-TES)—to assist in identification of samples

Mossbauer Spectrometer (MB) —to assist in identification of samples

Alpha-particle X-ray Spectrometer (APXS) —to assist in identification of samples

Rock Abrasion Tool (RAT)—to determine mineral hardness, expose unweathered surfaces of samples and so on

Magnet array (MA) —to identify magnetic properties of samples

Panoramic Camera (Pancam)—for surveying terrain and selecting targets for investigation

Microscopic Imager (MI)—for examining mineral crystalline structure and other microscopic details

In addition, there are seven cameras for navigation (navcams) and avoiding hazards (hazcams).

Given our lack of understanding of the processes that have shaped the Martian surface and the need for the rovers to operate in an unfamiliar and hostile environment, it is clear that the main key to successfully achieving these objectives is flexibility. For the most part, the rovers have sufficient power to communicate directly with Earth only occasionally and must relay communications via one of the orbiters over the red planet. This means that the vehicles must operate independently for extended periods and recognize potential hazards before it places itself in danger. The long, cold nights near the Martian South Pole mean that the electronics must survive at very low power levels provided by a radioisotope thermal generator (RTG) and low temperatures with heat provided by radioisotope heaters through the winter. The required flexibility all but necessitates reprogrammability for the “brains” of the rovers to supplement the radiation-hardened processor (RAD6000). For Spirit and Opportunity, the reprogrammability is supplied by several Xilinx Virtex I FPGAs, in which the configuration memory is highly susceptible to upsets. In addition, the Rovers require large amounts of volatile and nonvolatile memory a requirement filled by DRAMs and FLASH memory. The nonvolatile memory allowed data to be stored even when power was not available for the memory.

The Mars rovers also faced a variety of radiation threats, beginning with transit to Mars and continuing through the entire mission. Perhaps the most critical electronics were those that control the entry, descent and landing stages of the mission. The thin atmosphere of Mars made it a challenge to survive the still intense heat of re-entry and to shed enough momentum for the rovers to survive landing. During this phase, the spacecraft initially decelerated from roughly 5 km per second via friction between the aeroshell and the atmosphere. When the atmosphere was sufficiently thick, a parachute deployed to further decelerate the probe from 400 m/s to 85 m/s. Because the thin

Martian atmosphere provided insufficient deceleration via the parachute, retro-rockets were fired to slow the vehicle to its impact speed. Just prior to impact, air bags surrounding the vehicle deployed to absorb the shock of landing and the cushioned vehicle rolled to a stop. At this point, the air bags deflated and ramps and other mechanisms deployed to help the rover with egress from the spacecraft. An SEU or SET that interfered with any of these steps could have resulted in the loss of the mission.

Radiation hazards did not end once the orbiter was safely deployed on the Martian surface. The thin Martian atmosphere and lack of planetary magnetic field provided little protection from the space radiation environment, and the neutron flux at the Martian surface is quite high. (See Figure 7-2.) Finally, the RTG and radioisotope heaters also contributed to the radiation environment.

The radioisotope heaters maintain temperatures above the lower operation limit of $-40\text{ }^{\circ}\text{C}$ even as the Martian nighttime temperatures plummet to $-100\text{ }^{\circ}\text{C}$. Even so, ensuring electronics operates properly at these temperatures is challenging—as discussed in the Short Course section by Ben Blalock.

The RHA approach the Rover design team took favored mitigation over testing. The configuration memory of the FPGAs was scrubbed continually. Memory was protected by EDAC. That such an overdesign approach was a general policy for the rovers is evident from the fact that they are still operating nearly 3 years past their design life.

7.3 Summary

The approaches to hardening taken by SDO and the Mars Rovers differ considerably, in part because of SDO's more challenging radiation requirements. SDO chose to use radiation-hardened components for most of its needs, and where it had to use commercial technologies, it undertook thorough testing and made efforts to capitalize on these test results for hardening purposes.

The Mars Excursion Rovers used a fault-tolerant approach that sought to ensure rover safety if errors occurred, even if that meant a lengthy safe hold while ground control figured out the correct mitigation. Such a philosophy works well for an exploration mission, where mission objectives are not time critical and can continue to be met (and exceeded!) as long as the rover survives.

On the other hand, much of the most interesting science for SDO occurs when its radiation environment is at it worst. For such a mission, reliable operation in harsh radiation environments is essential, and the intensive testing effort provides assurance that SDO will be able to meet its requirements.

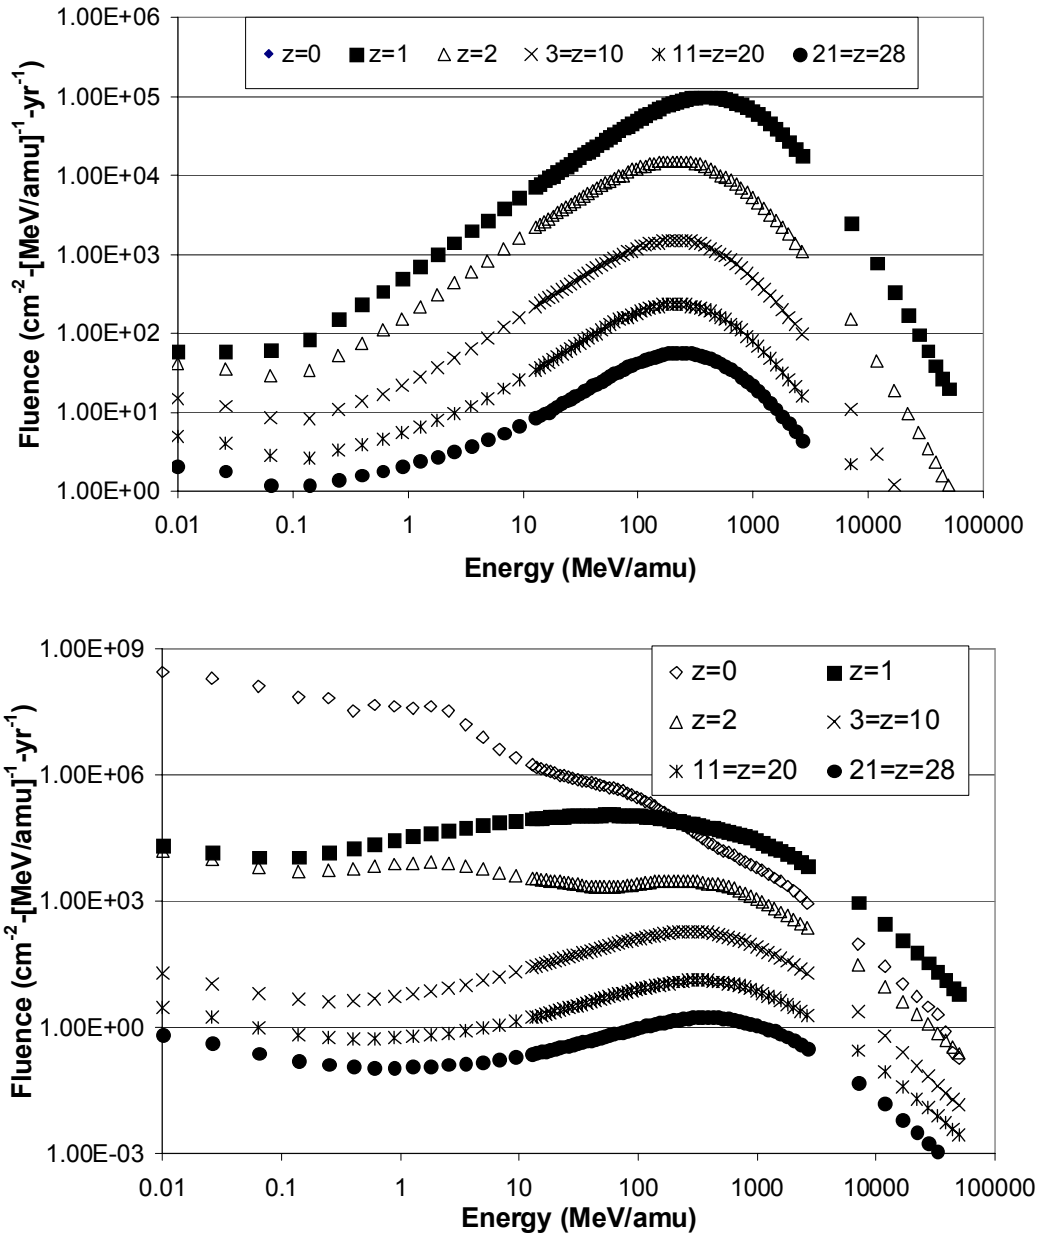


Figure 7-2. Heavy ion and neutron fluxes at the Martian surface are high, as indicated by these plots of the GCR environment before (a) and after (b) transport through the Martian atmosphere using the HZETRN deterministic transport routine (Courtesy of Martha Cloudsley, NASA Langley Research Center [97]).

8 REFERENCES

- [1] Pease R. L., "Microelectronic Piece Part Radiation Hardness Assurance for Space Systems," Part II of the Short Course presented at the 2004 Nuclear and Space Radiation Effects Conference, Atlanta, GA, 19 July 2004
- [2] Aeroflex "RadHard by Design SRAMs," http://ams.aeroflex.com/ProductPages/RH_4Msram.cmf, 10 May, 2007.
- [3] Hillman, R. et al. "Space Processor Radiation Mitigation and Validation Techniques For an 1,800 MIPS Processor Board," Radiation and its Effects on Components and Systems. 2003. RADECS 2003. Proceedings of the 7th European Conference on, 15 - 19 September 2003, pp. 347-352.
- [4] Samson J. R. et al., Technology Validation: NMP ST8 Dependable Multiprocessor Project, Aerospace Conference, 2006 IEEE, 4-11 March 2006, Page:14
- [5] Irom, F. et al., Single-Event Upset and Scaling Trends in New Generation of the Commercial SOI PowerPC Microprocessors, IEEE Trans. Nucl. Sci., Vol. 53, No. 6, Part 1, Dec., pp.3563 – 3568, 2006.
- [6] Datasheet, "HMXADC9225 – Radiation Hardened 12-Bit, 20 MSPS Monolithic A/D Converter," http://www.ssec.honeywell.com/avionics/datasheets/hmxadc9225_atod.pdf
- [7] Layton, P. et al., "TID Performance Degradation of High Precision, 16-bit Analog-to-Digital Converter," Radiation and its Effects on Components and Systems. 2003. RADECS 2003. Proceedings of the 7th European Conference on, 15 - 19 September 2003, pp. 553-557.
- [8] Rennie, R. et al., "Interpretation of SEE Test Results on a Mixed Signal Device," 2007 Single-Event Effects Symposium, Long Beach, CA 10-12 April, 2007.
- [9] M. Loose et al., "Instrument Design and Performance for Optical/Infrared Ground-based Telescopes," Masanori Iye, Alan F. M. Moorwood, Editors, Proceedings of SPIE Vol. 4841 (2003)
- [10] Poivey, C. et al., "Single-Event Effect (SEE) Response of Embedded Power PCs in a Xilinx Virtex-4 for a Space Application," Submitted to RADECS 2007.
- [11] Quinn, H. et al. "Radiation-Induced Multi-Bit Upsets in SRAM-Based FPGAs," IEEE Trans. Nucl. Sci., Vol. 52, No. 6, Part 1, pp. 2455-2461, 2006.
- [12] B. Fodness et al. "Monte Carlo Treatment of Displacement Damage in Bandgap Engineered HgCdTe Detectors," " Radiation and its Effects on Components and Systems. 2003. RADECS 2003. Proceedings of the 7th European Conference on, 15 - 19 September 2003, pp. 479-485.
- [13] Marshall, C. et al., "Comparison of Measured Dark Current Distributions with Calculated Damage Energy Distributions in HgCdTe. " accepted for publication in IEEE Trans. Nucl. Sci. 2007.

- [14] Pickel, J. et al., "Transient radiation effects in ultra-low noise HgCdTe IR detector arrays for space-based astronomy," *IEEE Trans. Nucl. Sci.*, Vol. 52, No. 6, Part 1, Dec., , 2006.
- [15] NASA-1995] NASA Systems Engineering Handbook, SP-610S, National Aeronautics and Space Administration
- [16] LaBel, K. A. et al., "Emerging Radiation Hardness Assurance (RHA) Issues: A NASA Approach for Space Flight Programs," *IEEE Trans. Nucl. Sci.*, vol. 45, No. 6, pp. 2727-2736, 1998.
- [17] MIL-HDBK-814, Military Handbook: Ionizing dose and neutron hardness assurance guidelines for microcircuits and semiconductor devices, Feb. 8, 1994.
- [18] Lum, G. K., "Hardness Assurance for Space Systems," Part I of the Short Course presented at the 2004 Nuclear and Space Radiation Effects Conference, Atlanta, GA, 19 July 2004
- [19] Poivey, C., "Radiation Hardness Assurance for Space Systems," Part IV of the Short Course presented at the 2002 Nuclear and Space Radiation Effects Conference, Phoenix, AZ, 15 July 2002.
- [20] Kinnison, J. "Achieving Reliable, Affordable Systems," Part IV of the Short Course presented at the 2002 Nuclear and Space Radiation Effects Conference, Newport Beach, CA 20 July 2002.
- [21] Heidergott, W., "System Level Mitigation Strategies," Part IV of the Short Course presented at the 1998 Nuclear and Space Radiation Effects Conference, Norfolk, VA, 12 July 1998.
- [22] Sievers, M. W. and Gilley, G. C., "Fault-Tolerant Computing: An Architectural Approach to Tolerating Radiation Effects," Part III of the Short Course presented at the 1985 Nuclear and Space Radiation Effects Conference, Monterrey, CA, 21 July 1985.
- [23] Ladbury, R. and Gorelick, J.L., "Statistical methods for large flight lots and ultra-high reliability applications," *IEEE Trans. Nucl. Sci.*, Vol. 52, No. 6, pp. 2630 – 2637, 2005.
- [24] Jordan, T. M., "An Adjoint Charged Particle Transport Method," *IEEE Trans. Nucl. Sci.* Vol. 23, p. 1857, 1976.
- [25] A. H. Johnston and T. F. Miyahira, "Characterization of Proton Damage in Light-Emitting Diodes," *IEEE Trans. Nucl. Sci.*, **47** No. 6, pp. 2500-2507 (2000).
- [26] Marshall, P. et al., "Single event effects in circuit-hardened SiGe HBT logic at gigabit per second data rates," *IEEE Trans. Nucl. Sci.*, vol. 47, no. 6, pp. 2669–2674, Dec. 2000.
- [27] Petersen, E. L., "Predictions and observations of SEU rates in space," *IEEE Trans. Nucl. Sci.*, vol. 44, no. 6, pp. 2174–2187, Dec. 1997.

- [28] Gates, M. et al., "Single Event Effects Criticality Analysis," NASA Report, See NASA GSFC Radiation Effects & Analysis Home Page, <http://radhome.gsfc.nasa.gov/radhome/papers/seecai.htm> .
- [29] Galloway, K. and Johnson, G. H., "CATASTROPHIC SINGLE-EVENT EFFECTS IN THE NATURAL SPACE ENVIRONMENT," Part IV of the Short Course presented at the 1996 Nuclear and Space Radiation Effects Conference, Indian Wells, CA, 15 July 1996.
- [30] Sexton, F. W., "Destructive Single-Event Effects in Semiconductor Devices and ICs," IEEE Trans. Nucl. Sci, Vol. 50, pp. 603-621, June 2003.
- [31] O'Bryan, M.V et al., "Current single event effects and radiation damage results for candidate spacecraft electronics," Radiation Effects Data Workshop, 2002 IEEE, 15-19 July 2002 Page(s):82 – 105
- [32] Johnston, A. H. et al., "The Effect of Temperature on Single-Particle Latchup," IEEE Trans. Nucl. Sci. Vol. 38, No. 6, pp. 1435-1441, 1991.
- [33] Kolasinski, W. A. et al., "The effect of elevated temperature on latchup and bit errors in CMOS devices," IEEE Trans. Nucl. Sci., vol. 33, pp. 1605–1609, 1986.
- [34] Johnston, A.H. and Hughlock, B.W., "Latchup in CMOS from single particles," IEEE Trans. Nucl. Sci. Vol. 37, No. 6, Part 2, pp.1886-1893, 1990.
- [35] Levinson, J. et al., "Single Event Latchup (SEL) in IDT 7187 SRAMs-dependence on ion penetration depth," Radiation and its Effects on Components and Systems, 1993., RADECS 93., Second European Conference on, 13-16 Sept. 1993 Page(s):438 – 440, 1993.
- [36] Schwank, J. R. et al., "Effects of Particle Energy on Proton-Induced Single-Event Latchup," IEEE Trans. Nucl. Sci, Vol. 52, pp. 2622-2629, 2005.
- [37] Schwank, J. R. et al., "Effects of Angle of Incidence on Proton and Neutron-Induced Single-Event Latchup," IEEE Trans. Nucl. Sci, Vol. 54, pp. 3122-3131, 2005.
- [38] Becker, H. N. et al., "Latent damage in CMOS devices from single-event latchup," IEEE Trans. Nucl. Sci. Vol. 49, No. 6, Part 1, pp. 3009 – 3015, 2002.
- [39] R. Ladbury, "Latent Damage, Best Test and Analysis Practice for Managing Risks," GSFC NASA Advisory, Advisory Number NA-GSFC-2005-05, 29th August, 2005.
- [40] McMorrow, D. et al., "Laser-Induced Latchup Screening and Mitigation in CMOS Devices," IEEE Trans. Nucl. Sci., Vol. 52, No. 6, pp. 1819-1825, 2005.
- [41] Layton, P. et al., "SEL Induced Latent Damage, Testing, and Evaluation," IEEE Trans. Nucl. Sci., Vol. 53, No. 6, pp. 3153-3157, 2006.
- [42] T. F. Wrobel, et al., "Current induced avalanche in epitaxial structures," *IEEE Trans. Nucl. Sci.*, vol. 32, pp. 3991–3995, Dec. 1985
- [43] O'Bryan, M.V et al., "Single event effects results for candidate spacecraft electronics for NASA," Radiation Effects Data Workshop, 2003 IEEE, 21-25 July, 2003, pp65-76.

- [44] Allenspach, M. et al. "Evaluation of SEGR Threshold in Power MOSFETs," IEEE Trans Nucl. Sci., Vol. 41, pp. 2160-2166, 1994.
- [45] Titus, J. L. et al., "A Study of Ion Energy and Its Effects Upon an SEGR-Hardened Stripe-Cell MOSFET Technology," IEEE Trans. Nucl. Sci., Vol. 48, No. 6, pp.1879-1885, 2001.
- [46] Titus, J. L. et al., "Prediction of Early Lethal SEGR Failures of VDMOSFETs for Commercial Space Systems," IEEE Trans. Nucl. Sci., Vol. 46, No. 6, pp. 1640-1651, 1999.
- [47] Swift, G. et al., "An Experimental Survey of Heavy Ion Induced Dielectric Rupture in Actel Field Programmable Gate Arrays (FPGAs)," IEEE Trans. Nucl. Sci., Vol. 43 No. 6, pp.967-972, 1996.
- [48] Koga, R. et al., "Ion-Induced Sustained High Current Condition in a Bipolar Device" IEEE Trans. Nucl. Sci., Vol. 41, No. 6, pp. 2172-2178, 1994.
- [49] Lum, G. K. et al., "The Impact of Single-Event Gate Rupture in Linear Devices," IEEE Trans. Nucl. Sci., vol. 47, No. 6, pp. 2373-2379, 2000.
- [50] O'Bryan, M. V. et al., "Recent Radiation Damage And Single Event Effect Results For Microelectronics," Radiation Effects Data Workshop 1999, 12-16 July 1999, pp.1-14.
- [51] Koga, R et al., "Heavy-Ion Induced Snapback in CMOS Devices," IEEE Trans. Nucl. Sci., Vol. 36, No.6, pp. 2367-2374, 1989.
- [52] Schwank, J. R. et al. "Radiation Effects in SOI Technologies," Vol. 50, No. 3, Part 3, pp.522-538, 2003
- [53] Oldham, T. R. et al. "SEE and TID Characterization of an Advanced Commercial 2Gbit NAND Flash Nonvolatile Memory," IEEE Trans. Nucl. Sci. Vol. 53, No. 6, pp. 3217-3222, 2006.
- [54] S. Duzellier, et al., "Protons and heavy ions induced stuck bits on large capacity RAMs," Radiation and its Effects on Components and Systems, 1993., RADECS 93., Second European Conference on, 13-16 Sept. 1993 Page(s):468-473, 1993.
- [55] Sternberg, A. L. et al., "Effect of Amplifier Parameters on Single-Event Transients in an Inverting Operational Amplifier," IEEE Trans Nucl. Sci., Vol. 49, No. 6, p1496-1501, 2002.
- [56] Gadlage, M. J. et al., "Single Event Transient Pulse Widths in Digital Microcircuits," IEEE Trans Nucl. Sci., Vol. 51, No. 6, pp.3285-3290, 2004
- [57] Buchner, S. et al., "Pulsed-laser testing methodology for single event transients in linear devices," IEEE Trans. Nucl. Sci., Vol. 51, No. 6, pp. 3716-3722, 2004
- [58] Benedetto, J. M. et al., "Digital Single Event Transient Trends With Technology Node Scaling," IEEE Trans Nucl. Sci., Vol. 53, No. 6, pp. 3462-3465, 2006.
- [59] Actel RTAX-S RadTolerant FPGAs Datasheet v. 3.0, September 2006, www.actel.com/documents/RTAXS_DS.pdf, 2006.

- [60] Boulghassoul, Y. et al., "Investigation of Millisecond-Long Analog Single-Event Transients in the LM6144 Op Amp," IEEE Trans. Nucl. Sci., Vol. 51, No. 6, pp3529-3537, 2004.
- [61] Gadlage, M. J. et al., "Digital Device Error Rate Trends in Advanced CMOS Technologies," IEEE Trans. Nucl. Sci., Vol. 53, No. 6, pp. 3466-3471, 2006.
- [62] Buchner, S. et al., "Investigation of single-ion multiple-bit upsets in memories on board a space experiment," IEEE Trans. Nucl. Sci. Vol. 47, No. 3, pp. 705-711, 2000.
- [63] S. Petit, et al., "Memories Response to MBU and Semi-Empirical Approach for SEE Rate Calculation," IEEE Trans. Nucl. Sci., Vol. 53, No. 4, pp. 1787-1793, 2006.
- [64] Radaelli, D. et al. , "Investigation of Multi-Bit Upsets in a 150 nm Technology SRAM Device," IEEE Trans Nucl. Sci., Vol. 52, pp. 2433-2437, 2005.
- [65] LaBel, K. A., "Radiation Test Challenges for Scaled CMOS Electronics," 2007.
- [66] Johnston, A. H. and Rax, B. G., "Testing and Qualifying Linear Integrated Circuits for Radiation Degradation in Space," IEEE Trans. Nucl. Sci., vol. 53, No. 4, pp1779-1786, 2006.
- [67] Schwank, J. R., "Total Dose Effects in MOS Devices," part III the Short Course presented at the 2002 Nuclear and Space Radiation Effects Conference, Phoenix, AZ, 15 July 2002.
- [68] Schrimpf, R, "Physics and Hardness Assurance for Bipolar Technologies,"Part IV of the Short Course presented at the 2002 Nuclear and Space Radiation Effects Conference, Vancouver, BC, 16 July 2001
- [69] Pease, R. L. et al., "Enhanced Low-Dose-Rate Sensitivity of a Low-Dropout Voltage Regulator," IEEE Trans. Nucl. Sci., Vol. 45, No. 6, pp. 2571-2576, 1998.
- [70] Johnston, A. H. et al. "Proton damage effects in linear integrated circuits," IEEE Trans. Nucl. Sci., Vol. 45, No. 6, pp. 2632-2637, 1998.
- [71] Barnaby, H. et al., "Proton Radiation Response Mechanisms in Bipolar Analog Circuits," IEEE Trans. Nucl. Sci. Vol. 48, No. 6, pp.2074-2080, 2001.
- [72] Gorelick, J. L. et al., "The effects of neutron irradiation on gamma sensitivity of linear integrated circuits," IEEE Trans. Nucl. Sci. Vol. 51, No. 6, pp. 3679-3685, 2004.
- [73] Johnston, A. H. , "Optical Sources, Fibers, and Photonic Subsystems" Part III of the Short Course presented at the 2004 Nuclear and Space Radiation Effects Conference, Atlanta, GA, 19 July 2004.
- [74] Benedetto, J. M. and Jordan, A. "A radiation-hardened cold sparing input/output buffer manufactured on a commercial process line," Radiation Effects Data Workshop 1999, 12-16 July 1999, pp. 87-91, 1999.
- [75] Benedetto, J. "Heavy Ion-Induced Digital Single-Event Transients in Deep Submicron Processes," IEEE Trans. Nucl. Sci., Vol. 51, No.6, pp. 3480-3485, 2004.

- [76] Shannon, C. E., "A mathematical Theory of Communication," Bell System Tech. Jour., **27**, pp. 379-423 and 623-656, 1948.
- [77] R. Hamming, "Error Detecting and Error Correcting Codes," Bell System Tech. Jour., **29**, pp.147-160, 1950.
- [78] Reed, I. S. and Solomon G., "Polynomial Codes Over Certain Finite Fields," J. Soc. Ind. Appl. Math., **8**, pp. 300-304, and Math. Rev., **23B**, p. 510, 1960.
- [79] Koga, R. et al. "Single Event Upset (SEU) Sensitivity Dependence of Linear Integrated Circuits (ICs) on Bias Conditions," IEEE Trans. Nucl. Sci., Vol. 44, No.6, pp. 2325-2332, 1997.
- [80] Koga, R. et al. "Single Event Transient (SET) Sensitivity of Radiation Hardened and COTS Voltage Comparators," [Radiation Effects Data Workshop, 2000](#), 24-28 July 2000 pp. 53 - 60, 2000.
- [81] Lalumondiere, S. D. et al. "Laser-induced and heavy ion-induced single-event transient (SET) sensitivity measurements on 139-type comparators," IEEE Trans. Nucl. Sci., Vol. 49, No. 6, pp. 3121-3128, 2002.
- [82] Marshall, P. et al. "Autonomous Bit Error Rate Testing at Multi-Gbit/s Rates Implemented in a 5AM SiGe Circuit for Radiation Effects Self Test (CREST)," IEEE Trans. Nucl. Sci., Vol. 52, No. 6, pp. 2446-2454, 2005.
- [83] Messenger, S. R. et al., "NIEL and Damage Correlations for High-Energy Protons in Gallium Arsenide Devices," IEEE Trans. Nucl. Sci. Vol. 48, No. 6, pp. 2121-2126, 2002.
- [84] Koga, R. et al., "SEE sensitivity determination of high-density DRAMs with limited-range heavy ions," [Radiation Effects Data Workshop, 2001 IEEE](#), 16-20 July 2001 pp. 182-189, 2001.
- [85] Koga, R. et al., "Permanent single event functional interrupts (SEFIs) in 128- and 256-megabit synchronous dynamic random access memories (SDRAMs)," [Radiation Effects Data Workshop, 2001 IEEE](#), 16-20 July 2001 pp. 182-189, 2001.
- [86] Heidergott, W. F. et al., "Complex SEU Signatures in High-Speed Analog-to-Digital Conversion," IEEE Trans. Nucl. Sci., Vol. 48, No. 6, pp. 1828-1822, 2001.
- [87] Witczak, S. C. et al., "Dose-Rate Sensitivity of Modern nMOSFETs," IEEE Trans, Nucl. Sci., Vol. 52, No. 6, pp. 2602-2608, 2005.
- [88] Shindou, H. et al., "Bulk Damage Caused by Single Protons in SDRAMs," IEEE Trans. Nucl. Sci., Vol. 50, No. 6, pp. 1839-1845, 2003.
- [89] David, J.-P. et al. "Light Particle-Induced Single Event Degradation in SDRAMs," IEEE Trans. Nucl. Sci., Vol. 53, No. 6, pp. 3544-3549, 2006.
- [90] Meshel, D., Personal Communication.
- [91] Frank, D. J., "Power-Constrained CMOS Scaling Limits," *IBM J. R&D.* vol. 46, no.2/3, pp. 235-244, 2002
- [92] Shaneyfelt, M. R. et al., "Impact of Aging on Radiation Hardness," IEEE Trans Nucl. Sci., Vol. 44, No. 6, pp. 2040-2047, 1997.

- [93] Pershenkov, V. S. et al., "Effect of aging on radiation response of bipolar transistors," IEEE Trans Nucl. Sci., Vol. 48, No. 6, Part 1, pp. 2164-2169, 2001.
- [94] Rodgers, M. P. et al., "The Effects of Aging on MOS Irradiation and Annealing Response," , IEEE Trans Nucl. Sci., Vol. 52, pp. 2642-2648, 2005.
- [95] Batyrev, I. G. et al., "Effects of Water on the Aging and Radiation Response of MOS Devices," IEEE Trans Nucl. Sci., Vol. 53, pp. 3629-3635, 2006.
- [96] Buchner, S. et al., "Impact of TID Effect on the Analog SET Sensitivity of Linear Bipolar Integrated Circuits," Accepted for 2007 IEEE Nuclear and Space Radiation Effects Conference, Honolulu, 23-27 July 2007.
- [97] Cloudsley, M. et al "Prediction of Martian Surface Neutron Environment," http://hesperia.gsfc.nasa.gov/sspvse/Group_D/Cloudsley_Mars_neut.ppt and Personal Communication.