

Strategies for SEE Hardness Assurance—From Buy-It-And-Fly-It to Bullet Proof

Ray Ladbury
NASA Goddard Space Flight Center

Table of Contents

1	Introduction.....	1
1.1	Preliminary Concepts	2
2	The SEE Radiation Environment.....	6
2.1	The Terrestrial SEE Environment.....	6
2.2	Space Radiation Environments	7
2.2.1	Galactic Cosmic Rays	8
2.2.2	Solar Particle Events	10
2.2.3	Trapped Radiation.....	11
2.2.4	Particle Transport and Secondary Particle Environment	13
3	SEE Mechanisms and Consequences.....	15
3.1	Destructive SEE Modes	16
3.1.1	Single-Event Latchup.....	16
3.1.2	Single-Event Gate Rupture (SEGR)	17
3.1.3	Single-Event Burnout.....	20
3.1.4	Other Destructive SEE.....	21
3.2	Nondestructive SEE	22
4	SEE Threat Identification	26
4.1	Mission Requirements.....	26
4.2	Environments Revisited	26
4.3	Technology and SEE Vulnerabilities	28
4.4	Threat Identification: Mitigation vs. Robust Design.....	29
5	SEE Threat Evaluation.....	30
5.1	Historical Data.....	30
5.1.1	Sources of Historical Data	30
5.1.2	Using Historical Data.....	31
5.2	Using Proxy Data to Bound SEE Susceptibility	32
5.2.1	Similarity Data	34
5.2.2	Protons as a Proxy for Heavy-Ion SEE Susceptibility	36
5.2.3	More on Proxies.....	39
5.3	SEE Testing.....	39
5.3.1	SEE Testing Goals	40
5.3.2	SEE Rate Estimation Methods and Testing Goals.....	41
5.3.3	Proton SEE Testing.....	44

5.3.4	SEE Risk Evaluation: From Data to Rates to Risk	45
6	Mitigation of SEE Risk.....	46
6.1	Reducing SEE Probabilities	46
6.2	Reducing SEE Consequences.....	47
6.2.1	Specific Strategies.....	47
6.2.2	Redundancy and Supporting Strategies	48
6.3	Putting It All Together	53
7	SEE Hardness Assurance for SmallSats	55
7.1	Why Harden SmallSats	55
7.2	Challenges for SmallSat SEE Hardness Assurance	56
7.2.1	Implications of Using COTS Parts and Systems	56
7.2.2	Limitations of Tolerant Design for SmallSats	58
7.2.3	SEE Test Challenges for SmallSat Missions	59
7.3	Summary	61
8	SmallSat SEE Hardness Assurance	62
8.1	Buy-It-And-Fly-It.....	62
8.2	The Swarm—Safety in Numbers?.....	62
8.3	Board- and Box-Level Testing.....	63
8.4	Partnering—Relying on the Kindness of Strangers	64
8.5	Risk-Informed Testing	65
8.5.1	Evaluating Failure Consequences for SmallSats	67
8.6	Examples	68
9	Conclusions.....	70
9.1	Risk Tolerance and SEE Risk Management	70
9.2	Another Decade of SEE Surprises	71
9.3	Improving SEE HA for SmallSats... and For All Sats	72
9.4	Final Thoughts.....	74
10	References.....	74

1 Introduction

The discovery of single-event effects (SEE) 45 years ago presented radiation engineers with a new kind of threat [1]. SEE could occur with no warning at any time during the mission. They could have consequences ranging from trivial to catastrophic. Moreover, when they were discovered, the closest terrestrial analog to SEE occurred during nuclear tests—for all practical purposes, they were unique to the space environment. Over the next several years, the radiation effects community developed effective methods for evaluating SEE risk. Even the first paper on the discovery [1] discusses basic mechanisms and contains first attempts at rate estimation. SEE test methods began to be developed soon after [2]. Protons were reported to cause single-event upset (SEU) by 1979 [3]. The Nuclear and Space Radiation Effects Conference (NSREC) in 1980 saw the first session devoted specifically to single-event phenomena, and soon, the alphabet soup of SEE modes included destructive effects [4-6]. Throughout the first two decades after the discovery of SEE, rate estimation methods underwent continual improvement [7], and by 1996, the basic infrastructure of SEE hardness assurance (SEE HA) was in place.

The process by which SEE HA developed followed the general approach of risk management. The threat(s) were identified—either by looking at on-orbit data or developing appropriate ground-based test methods. The threat was evaluated by carrying out further testing, modeling the space environment and developing models of the basic SEE mechanisms to foster understanding. Rate estimation models and improved test techniques were developed where possible to bound failure probabilities. Ultimately, mitigation techniques were developed, exploiting understanding of the SEE mechanisms to increase hardness at the device, part, circuit or system level. And for the most part, it has worked amazingly well.

At present, SEE HA faces new challenges, arising from new “risk tolerant” missions that still want a reasonable chance of success; from increased pressure to exploit commercial-off-the-shelf (COTS) technologies for their improved performance; from more complex parts and avionics architectures that complicate efforts to evaluate SEE threats; and from ever more stringent demands for lower costs and increased performance.

Whether they are called, SmallSats, Explorer class, Class D or any of a range of other terms, many new missions seek to avoid the high qualification and verification costs of satellite development by allowing higher tolerance of risk. They can do this by flying in more benign radiation environments over shorter durations and decreasing the costs of the mission by launching as a secondary payload, by shortening design cycles, by minimizing testing and/or by use of less costly, higher-performance but riskier COTS parts. Ideally, allowing designers to take more risks would lead to more missions, more innovation and more rapid evolution of satellite design. Unfortunately, SEE stand as a significant obstacle to mission success and to trading risk tolerance for lower costs. In benign radiation environments over short timescales, SEE may dominate radiation risk and in some cases the risk to mission success.

Use of COTS poses challenges in part because of the broad range of technologies, part types, manufacturers and quality encompassed by the term. It may include technologies so soft that they upset due to direct ionization from protons [7], and it may include parts with exceptional radiation performance [8]. Moreover, the broad range of COTS choices available to designers means there is unlikely to be radiation data for any given part or even for similar parts, and the rapid cycle time for many COTS parts limits the time during which the data are valid. Issues of die and lot traceability and lot uniformity raise questions as to whether test part performance will

be representative of flight parts. And many COTS parts—often precisely the ones an engineer wants to use—are extremely complicated and difficult to test.

It is partly this increase in complexity in COTS parts that has placed more pressure on SEEHA to limit costs. Not only will a part with many modes of operation take longer to test, the high speed and high pin count of the part will require more complicated and expensive test hardware to characterize its response to radiation. Moreover, as SmallSats attempt to economize and trim schedules through increased use of COTS, increased SEE testing costs stand out all the more. The very actions that economize other aspects of the design process can increase either SEE HA costs or—more likely—the risks that SEE may cause the mission to fail.

These significant new challenges present an opportunity to step back and view SEE HA in terms of its roots in risk management. Although SEE pose unique risks, the management of those risks follows the same basic steps as risk management for any other threat/hazard, namely:

- 1) Identify the threat—mainly by examining the mission radiation environment, the technologies used in the mission and their application conditions (e.g., temperature, voltage, etc.)
- 2) Evaluate the threat—in terms of the threat risk, $R(T)=P(T)\times C(T)$, where $P(T)$ the probability of the threat occurring and $C(T)$ is the cost or consequences to the mission if the threat is realized. Both $P(T)$ and $C(T)$ can be evaluated using a variety of data types—from SEE test data on flight-like parts to data on similar parts, etc.
- 3) Mitigate the threat—either reducing the probability that a threat occurs or the consequences if it does occur.

These steps hold true for SEE mitigation for any space mission—be it a nanosat or a National Asset. What varies are the resources available for evaluation and mitigation. In this short course, we examine SEE HA as a risk mitigation activity, capitalizing on the structure this introduces to highlight how SEE HA changes for different mission classes and why it can be difficult to translate additional risk tolerance into savings on SEE HA activities.

We begin by reviewing some fundamental aspects of SEE, including the space radiation environment and various types of SEE and their mechanisms. Next, we examine the conventional approach to SEE HA as it is applied to primary payload missions—from threat identification through threat evaluation and mitigation. Following this overview, we look at how SEE HA differs for SmallSats. In this discussion, the secondary payload status of most SmallSats imposes important constraints, encouraging use of COTS technologies, which in turn affects how SEE HA can be applied to such missions. Finally, we mention some of the approaches taken by various groups to dealing with these constraints.

1.1 Preliminary Concepts

SEE occur when an ionizing particle traverses a sensitive region in an active semiconductor device and generates sufficient charge (called the critical charge or Q_c) to alter the normal functionality of the device. The ionizing particle may either be primary—originating in the mission radiation environment—or secondary—resulting from a scattering event between a primary environment particle and materials within the device in which the SEE occurs. (See Figure 1-1.)




Figure 1-1 Single-event effects occur when an energetic ionizing particle traverses a volume in the device sensitive to the SEE (the sensitive volume, or SV). The ionizing particle may either be primary from the mission radiation environment (red arrow) or secondary—generated by collisions between lightly ionizing primary particles (dashed blue line—usually protons, p+) and nuclei within the semiconductor.

In general, the greater the charge Q generated in the sensitive volume (SV) of the device, the more likely an SEE is to occur. Q depends on both the density of the charge track and on the path length of the ion through the sensitive volume. Charge track density is proportional to the ion's LET (linear energy transfer is defined as the rate of energy loss along the ion track path, normalized to target material density), so

$$Q \propto \int \text{LET}(x) dx \quad (1)$$

where x is the linear distance along the charge track. For Si, an electron-hole pair is generated for every 3.6 eV of energy deposited by the track. In general LET depends on the charge of the ion, its mass and its energy. As Figure 1-2 illustrates for ^{56}Fe ions, LET at high energies is low, rising as ion energy decreases to a maximum value called the Bragg peak, and then falling rapidly to 0 as ion energy goes to 0. (Note, the minimum LET vs. energy (called minimum ionizing energy) occurs when the kinetic energy of the ion is a bit over twice its rest energy. Above this energy, LET rises gradually with increasing ion energy, but ion fluxes in the space environment are negligible in this energy range.) In practical terms, LET for ions with energies significantly higher than the Bragg peak change relatively slowly, making it possible to approximate LET of the ion as a constant value. Then the charge deposited in the SV is proportional to LET multiplied by the chord length through the SV. This simplifies the rate estimation process by separating the model into a part that depends only on the distribution of ion LETs—a property of the space environment—and a part that depends only on the device response versus LET—a property of the device to be determined by testing and modeling.

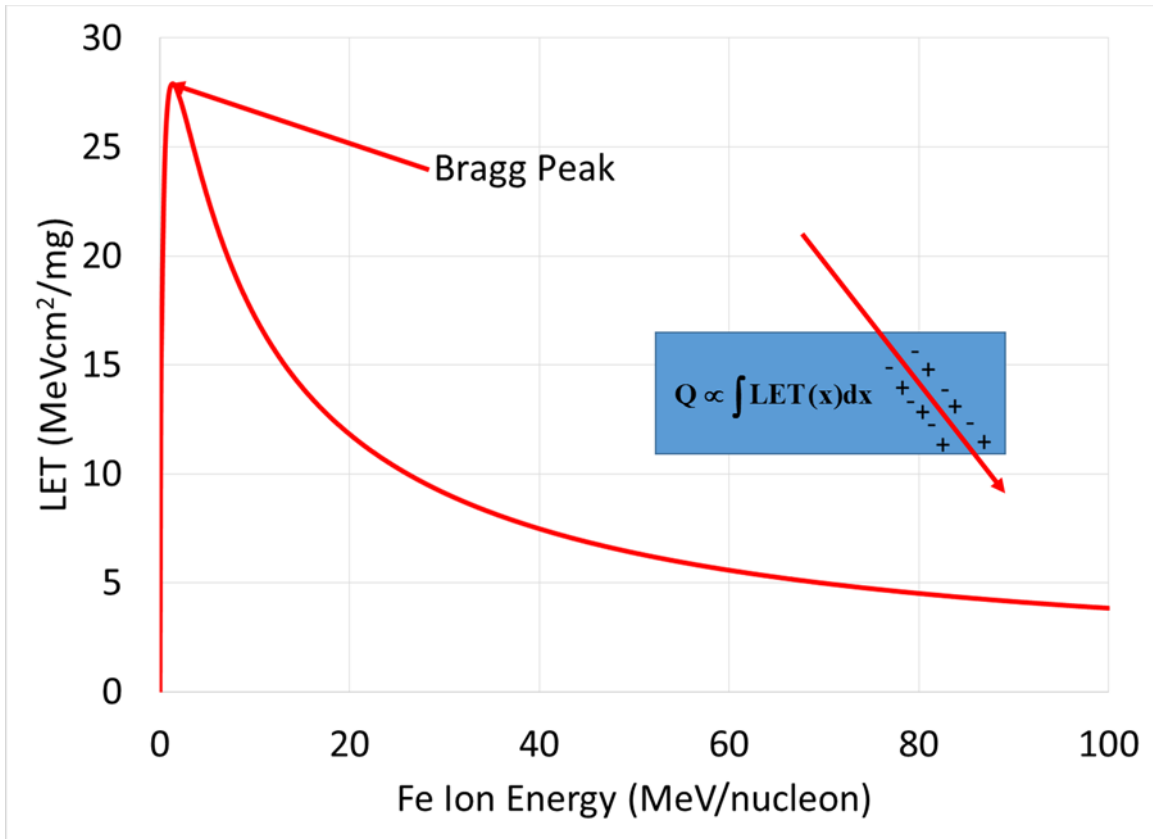


Figure 1-2: The rate of energy deposition/loss of a charged particle can be characterized by its LET—the energy lost per unit length along the ion track normalized to the density of the target material. Since almost all the energy lost goes into ionizing the medium, the total charge generated by the track is proportional to the integral of LET along the track length x . The maximum LET occurs for the Bragg Peak, with LET decreasing gradually at higher energy and rapidly at lower energy as the ion particle ranges out.

The most common measure of the likelihood of SEE occurrence as a cross section, the ratio of the SEE count to the ion fluence (particles per unit area) that cause the errors. As such, SEE cross section has units of area, and can be thought of—very loosely—as the area on the chip susceptible to the SEE. One of the most important products of SEE testing is the curve that measures SEE cross section σ as a function of LET, σ vs. LET. This curve characterizes device response and is used to estimate SEE rates. Figure 1-3 shows the cross section for single-event upsets in a 256K static random access memory (SRAM), the Elpida HM65656. For very low LETs, the path length through the SV is too short for the ion to deposit the critical charge Q_c . Above the onset LET, LET_0 , where one first begins to register SEE, σ rises rapidly. Eventually, the cross section saturates at the limiting cross section (σ_{lim}) or increases only slightly with higher LET. Often, the σ vs. LET curve is fit to a cumulative Weibull form:

$$\sigma(LET) = \sigma_{lim} \times \text{Weibull}(LET - LET_0, w, s), \quad (2)$$

where w and s are the Weibull width and shape parameters. The Weibull fit parameters can then serve as input to SEE rate estimation routines, such as CREME96.

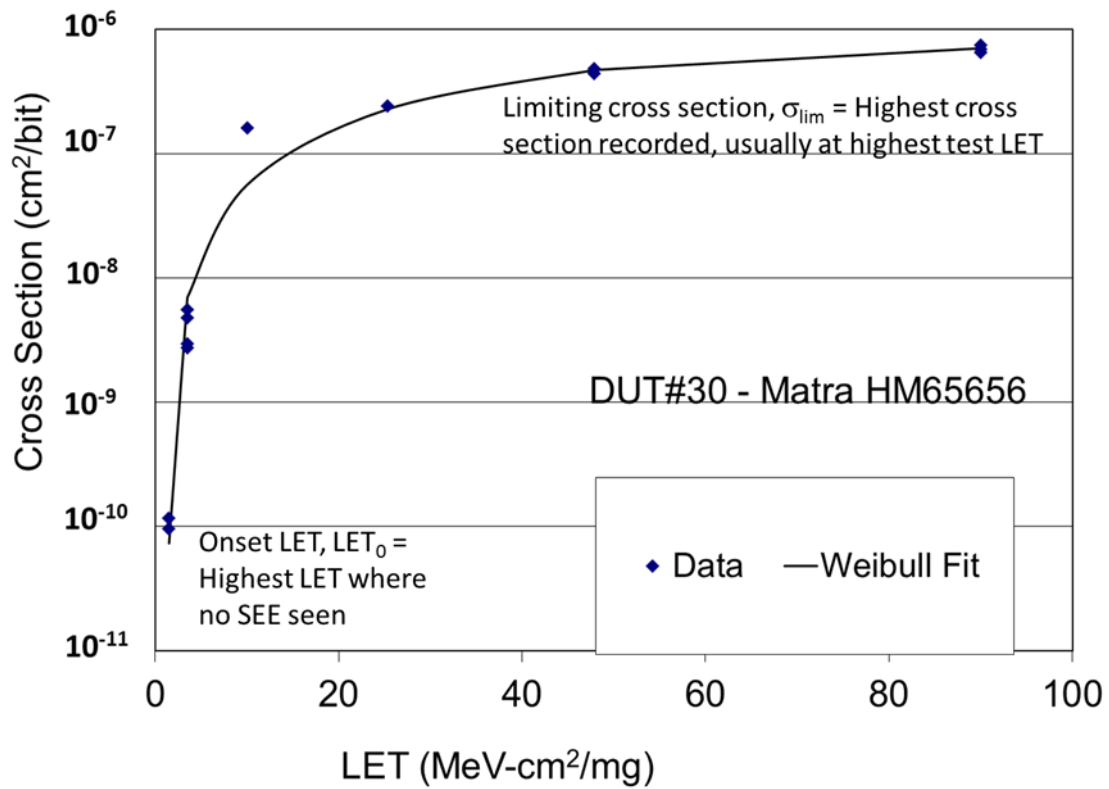


Figure 1-3: SEE susceptibility is measured by the SEE cross section σ —the ratio of the number of events seen to the ion fluence. Usually, σ is 0 up to some minimum LET, LET_0 , and then rises rapidly until plateauing at high LET at the limiting cross section σ_{lim} . The σ vs. LET is often fit to a cumulative Weibull form, with the fit parameters serving as input to SEE rate estimation routines.

2 The SEE Radiation Environment

Although in principle, an SEE could be caused by any ionizing particle, for technologies flown to date, the particles responsible have been neutrons (via indirect ionization), protons (via both direct and indirect ionization), and heavier ions (mainly via direct ionization). Recent SEE experiments have revealed that deep submicron complementary metal oxide semiconductor (CMOS) devices may be susceptible to SEE from muons and even electrons [9-10]. However, since these devices are not at present seen as realistic candidates for flight projects, we confine ourselves for the most part to consideration of neutrons, protons and heavier ions.

2.1 The Terrestrial SEE Environment

In the terrestrial space radiation environment, SEE rates are dominated by neutrons that are the last surviving daughter products of showers of particles from interactions between high-energy galactic cosmic rays (GCR) and the atmosphere. Neutron fluxes at sea level are low—for example only about 14 neutrons/cm² per hour, but fluxes rise roughly exponentially with altitude. Fluxes also rise as one heads poleward, and there are also slight dependencies on longitude (due to geomagnetic field variations). Peak neutron fluxes occur at about 60,000 feet of altitude near the poles and are about 4/cm² per second [11]. Terrestrial SEE are of importance mainly for aircraft avionics (due to the altitude dependence of neutron fluxes and the large number of flight hours), large data storage facilities [12] and highly critical applications such as biomedical devices [13]. SEE are also expected to be of critical importance for drones and autonomous vehicles.

The most numerous energetic (>100 MeV) particles at sea level are muons, which are also produced in cosmic ray particle showers. Recent work has shown that muons can cause SEE in some deep submicron SRAMs, especially for reduced supply voltages [9, 14-15]. However, for the foreseeable future, it is likely that neutrons will continue to be the dominant source of terrestrial and atmospheric SEE. For altitudes higher than ~20,000 m, pions, protons and even some residual GCR heavy ions can contribute to SEE [16]. (See figure 2-1.)




Figure 2-1: In the terrestrial SEE environment, SEE are mainly caused by neutrons—the products of collisions between galactic cosmic rays and nuclei in the atmosphere. For very sensitive devices, muons may also contribute to the rate at sea level. At higher altitudes, one begins to see contributions from mesons and eventually from the most energetic GCR ions, which penetrate as far as the upper stratosphere.

In addition to the natural particle environments, interactions between parts or packaging and the terrestrial radiation environment can also contribute to SEE. Radioactive elements associated with Pb in solders can produce α particles, and low-energy neutrons can cause ^{10}B in borosilicate passivation to fission into an α particle and a ^7Li ion. These mechanisms and the subject of terrestrial SEE in general were covered in an Institute for Electrical and Electronics Engineers (IEEE) NSREC Short Course by Robert Baumann in 2013 [17]. As such, in the remainder of this Short Course, we concentrate on SEE in space environments.

2.2 Space Radiation Environments

In the space radiation environment, the particles responsible for SEE are mainly high-energy protons and heavier ions. Neutrons may be generated during the interaction between the primary environment and spacecraft materials, but their fluxes are usually negligible. Low-energy protons may also contribute to SEE for very sensitive microcircuits, but to date have not been a major source of SEE in space hardware [18].

The space radiation environment encompasses a broad range of possible threat levels. Low-altitude, low-inclination Earth orbits are well shielded by the geomagnetic field, which reduces fluxes of GCR and solar particles and up to ~ 200 km above the planet's surface, includes negligible fluxes of particles trapped by those same fields. In contrast, interplanetary space provides little protection from GCR and full exposure solar particle events, and the trapped radiation belts of Jupiter can result in exposures up to $10^7 > 10\text{-MeV}$ protons/cm² per second [19].

Figure 2-2 illustrates the main components of the space radiation environment, GCR, Solar Particle Events (SPE), and trapped radiation belts. These are described in more detail below. Although GCR energies can reach over 10^{21} eV, most of the flux has energies from a few hundred MeV to a few GeV/nucleon. They consist of all elements ($1 \leq Z \leq 92$) [20]. The Sun is a source of electrons, protons and heavier ions in the solar wind and occasional more energetic bursts of protons and ions during solar particle events, which can include high fluxes of protons and heavier ions over periods of a few hours to a few days [21]. Protons and electrons can also become trapped by planetary magnetic fields, forming toroidal radiation belts. For Earth, this includes an inner proton belt with energies up to a few hundred MeV and an outer electron belt with energies of a few MeV.

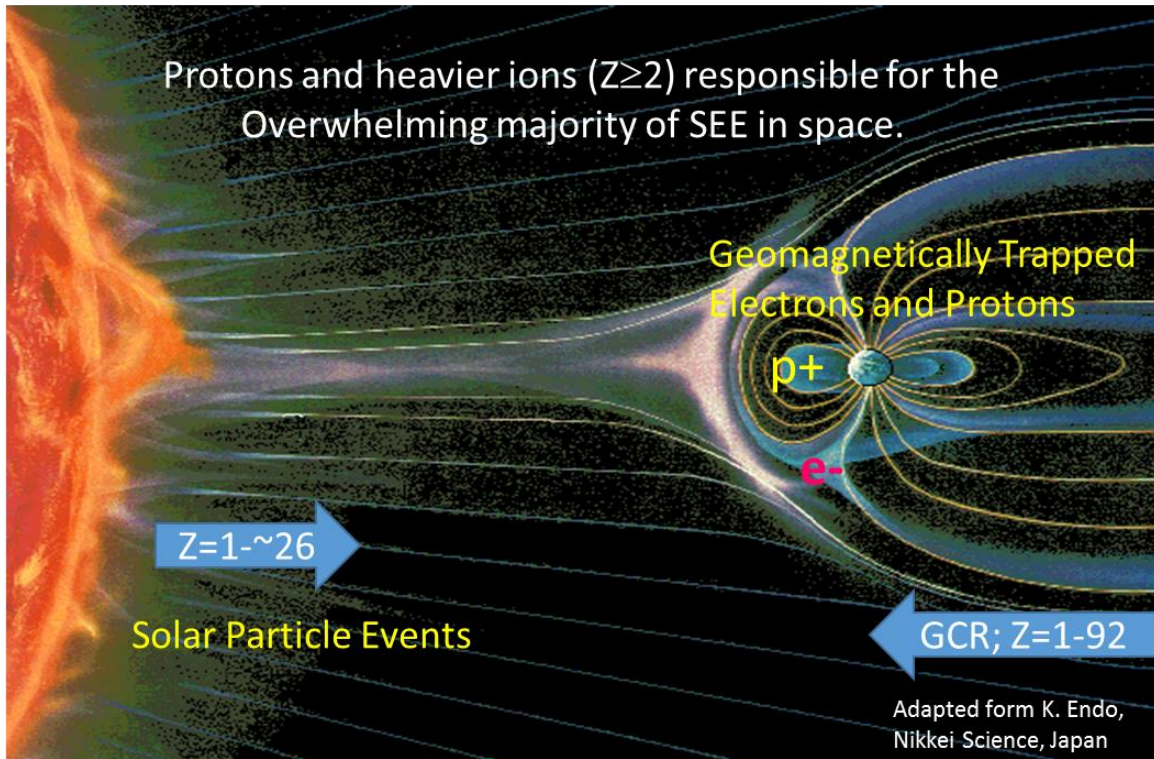


Figure 2-2 Space radiation environments includes several sources of energetic ions. Galactic cosmic rays impinge isotropically on the solar system with a low but fairly consistent flux. Solar Particle events produce episodic high fluxes of protons and heavier ions for times ranging from a few hours to a few days. They can occur any time, but are most likely during solar maximum. Particles trapped by planetary magnetic fields form toroidal belts surrounding the planet and can result in fluxes exceeding 10^5 protons/cm²/s in the heart of the proton belts.

2.2.1 Galactic Cosmic Rays

Galactic cosmic rays are very high energy particles that originate outside the Solar System. Although protons account for about 90% of GCR, the main particles of interest are GCR heavy ions ($2 \leq Z \leq 92$). Because GCR originate in the explosions of supernovae, they are accelerated to high energies (peaking around several hundred MeV per nucleon), making shielding an impractical mitigation strategy against GCR induced SEE [20]. Figure 2-3 shows the flux vs. energy dependence for GCR ions ($1 \leq Z \leq 92$) from the CRÈME96 model for conditions of Solar Minimum and after traversing 100 mil of aluminum equivalent shielding (sufficient to attenuate the low-energy tail). Because Fe ions play an important role in the explosions of supernovae, they are much more abundant than ions of comparable atomic number, and so are important in determining SEE rates.

For purposes of SEE rate estimation, it is convenient to bin GCR ions according to their LET. Figure 2-4 illustrates the ion flux for both a geostationary orbit (GEO) and the orbit of the International Space Station (ISS). The differences between GEO and ISS fluxes illustrate the effect of attenuation by the geomagnetic field, while the differences between Solar Maximum and Solar Minimum illustrate the effect of the solar activity (e.g., the solar wind and heliomagnetic field). The rapid rise in flux seen just above $LET=1$ MeV-cm²/mg is largely due to Fe ions near their minimum ionizing LET, while the drop off just below $LET=30$ MeV-cm²/mg represents Fe ions near their Bragg Peak. GCR fluxes are sufficiently consistent that the

CREME96 model presents just two conditions—Solar Minimum (maximum GCR flux) and Solar Maximum (minimum GCR flux). SEE rates during Solar Minimum are often about 3x higher than those during solar maximum. Attenuation of the GCR flux due to the heliomagnetic field and solar wind also weakens with distance from the Sun, with fluxes increasing by about 4% per Astronomical Unit (AU) for high-energy ions near the peak flux and about 10% per AU for ions with energies less than 100 MeV/u [21]. Thus, GCR SEE rates at Jupiter could be 50% higher than those near Earth, those at Saturn near double those near Earth.




Figure 2-3: Although the vast majority of GCR are protons, there are substantial fluxes of heavier nuclei, especially Fe, which, due to its high nuclear binding energy is the most abundant heavy element. Magenta stars indicate the Energy (in MeV/nucleon) where the ion's Bragg Peak occurs, and the vertical blue bar indicates where the LET is minimum.

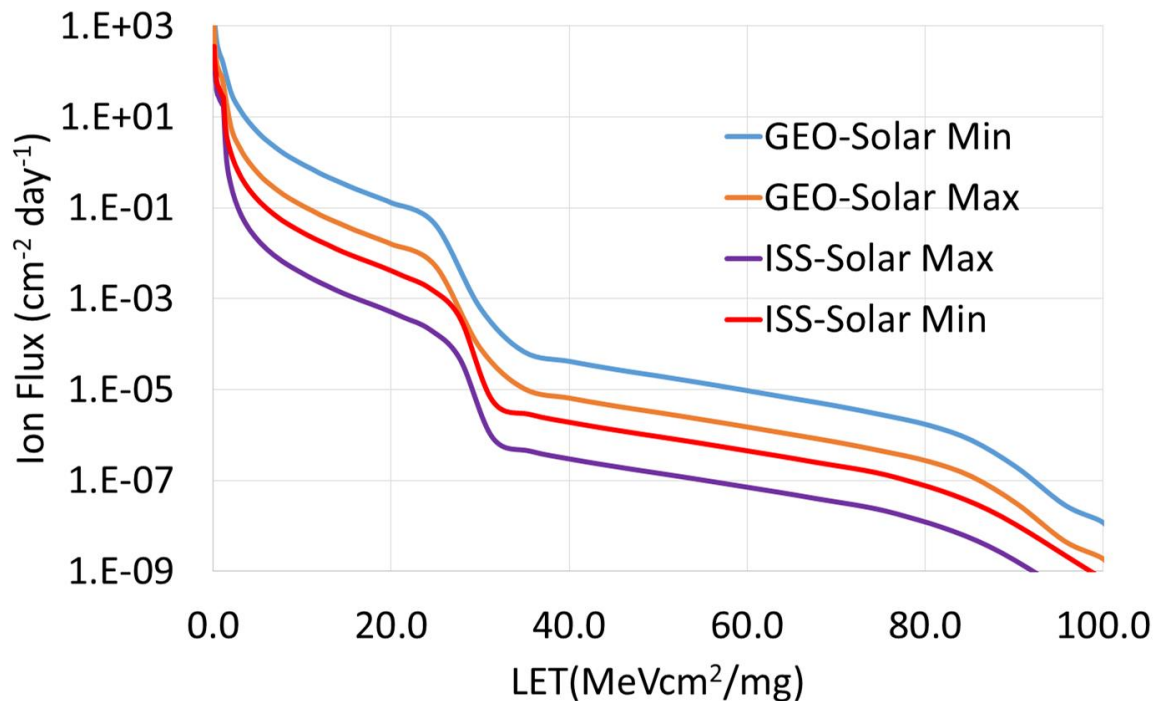


Figure 2-4: Galactic cosmic ray flux vs. LET at geostationary orbit and for the orbit of the International Space Station for Solar Maximum and Solar Minimum conditions in the CREME96 model.

2.2.2 Solar Particle Events

Solar Particle Events occur when a coronal mass ejection (CME) throws off a jet of plasma. The particles follow magnetic field lines through interplanetary space, being accelerated in some cases to 80% of the speed of light. Such events are directional and vary considerably in terms of duration, particle types, and particle energies. However, several generalizations are possible.

- 1) Although SPE can occur at any time, they are much more likely during Solar Maximum, and the more magnetically active the Sun, the more likely they are. (Cumulative solar proton fluences are about 2 orders of magnitude higher in Solar Maximum years than in Solar Minimum years on average.) [20]
- 2) Reference 20 found that SPE fluences roughly follow a truncated power law distribution, and that worst-case events ought to be bounded by $9.6 \times 10^8 \text{ cm}^{-2}$ (70% Confidence), $4.6 \times 10^9 \text{ cm}^{-2}$ (90% Confidence), and $1.2 \times 10^{10} \text{ cm}^{-2}$ (99% Confidence) for 30-MeV protons. For other energy thresholds, the bounds scaled roughly exponentially with the threshold energy.
 - a. SPE proton and ion energy spectra are considerably softer than those of GCR. This means that moderately heavy shielding (>500 mil Al equivalent) can be effective against these events.
 - b. Generally, the ion species in a SPE are representative of solar abundances, although some (including the October 1989 event in Figure 1-8) are enhanced in heavier ions—especially Fe.

Over the worst day, very large SPE like the October 1989 event can increase SEE rates by over 1000× for very soft devices (Onset LET <~1 MeV-cm²/mg) and by >100× for harder devices, compared to GCR background rates.




Figure 2-5: Heavy-ion flux vs. energy for the October 1989 Solar Particle Event as modeled in CREME96.

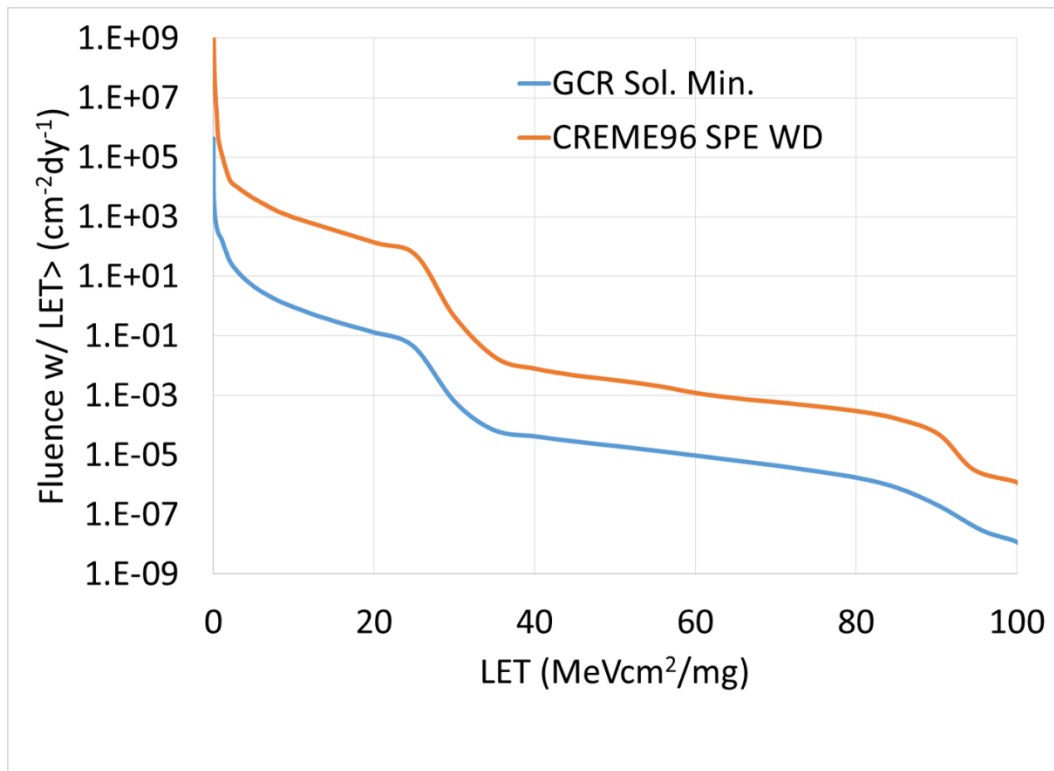


Figure 2-6: Flux vs. LET for the October 1989 SPE in CREME96 compared the Solar Minimum CREME96 Flux.

2.2.3 Trapped Radiation

Permanent radiation belts form as a result of energetic charged particles in the space plasma with planetary magnetic fields. In addition to Earth's Van Allen Belts, Mercury, Jupiter, Saturn, Uranus and Neptune have all been observed to have radiation belts. Although radiation belt

models exist for Saturn, and preliminary models based on Voyager data have been constructed for Uranus and Neptune, Jupiter. In this course, we will concentrate on Earth's trapped radiation. However, interested readers can refer to [19] for discussion of the Jovian magnetosphere, to [23] for discussion of modelling of Saturn's radiation environment and to [24] for a general review of space weather including trapped radiation for planets and moons throughout the Solar System.

The main trapped particles of concern are protons. (Trapped heavy ions have been observed in Earth's magnetosphere, but their energies are too low to be of concern if parts have even very light shielding ($>100 \mu\text{m Al}$ equivalent) [25].) The trapped proton flux of concern varies significantly in terms of the spacecraft's orbit (both radius and inclination), the proton energy required to cause an SEE and when the mission is flying. In Low-Earth Orbits (LEO) with low inclination, fluxes may be $<100 \text{ protons/cm}^2/\text{s}$ for energies $>30 \text{ MeV}$, while in the South Atlantic Anomaly at comparable altitudes, peak fluxes can be over $2000 \text{ protons/cm}^2/\text{s}$. The highest proton fluences occur are found in Medium Earth Orbit (MEO)—about 3500-8000 km above the equator and can exceed $10^5 \text{ protons/cm}^2/\text{s}$.

For trapped protons, there are currently two widely used models AP8, and its successor AP9. AP8 has the advantages of simplicity (the only choice is between AP8MIN and AP8MAX) and familiarity. AP9 allows confidence levels to be selected for the environment that reflect uncertainties due to both measurement errors and limitations and environmental fluctuations [26]. Although most of the differences between the models are minor, for higher energies, which are important for determining SEE rates, differences can be significant for some orbits. The ability in AP9 to select a confidence level is potentially important for SEEHA to ensure that SEE mitigations are not overcome by environmental fluctuations.

Finally, although, we have considered the sources of ionizing particles in the space radiation environment separately, the total particle flux is additive. While some GCR and SPE particles are shielded by the geomagnetic field in a LEO environment, many penetrate and contribute to SEE along with the trapped protons. The total threat from the environment depends on the fluxes of each ion species and their energy spectra.

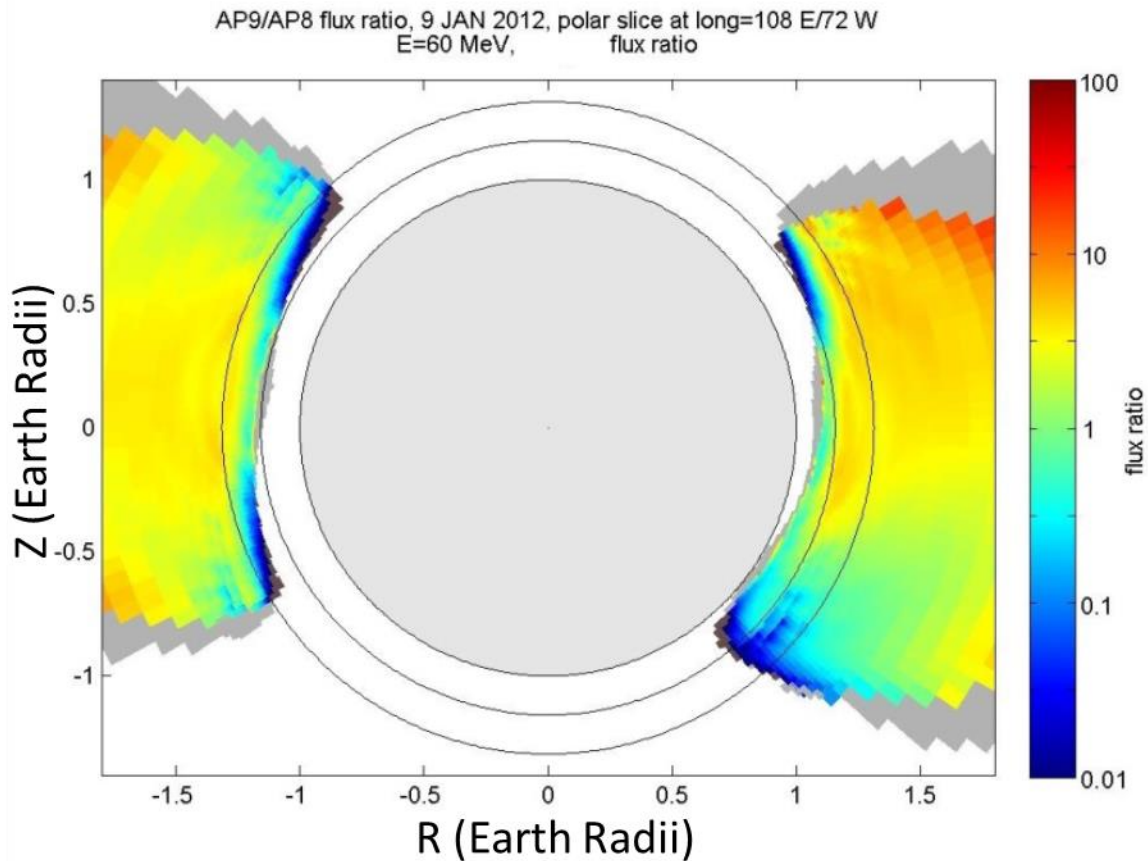


Figure 2-7: AP9(50%)/AP8 Flux ratios for >60 MeV protons on a polar cut 72° W to 108° E [27].

2.2.4 Particle Transport and Secondary Particle Environment

In order to cause SEE in a device, the particles from the space environment must reach the device SV. This means that they must traverse the materials that make up the spacecraft, packaging and other inert materials above the SV. In doing so, it may suffer one of 3 fates:

- 1) It may not have sufficient energy to reach the SV, stopping in the overburden above.
- 2) It may lose energy gradually, changing its LET, angle of incidence, etc.
- 3) It may collide with nuclei, transforming ion species, generating secondary particles.

There exist several options for transporting the external spacecraft environment through the spacecraft, and the degree of fidelity required—both for the spacecraft model and the transport code—depends on the sensitivity of the device being assessed and the space environment being considered. Shielding has relatively little influence on GCR fluxes, other than to filter out the low-energy tail of GCR ions. Usually for GCR, it is sufficient to transport the ions through a 100 mil Al equivalent spherical shell, which can be done in CREME96. GCR ions are sufficiently energetic, and the flux vs. energy distribution is sufficiently flat over a broad range of energies that changing the shielding will not dramatically affect the SEE rate. Better understanding of shielding is more important for dealing with solar heavy ions. The energies of these particles are lower, so LET changes more and more ions range out as the shielding changes. Perhaps the most detailed transport calculations are required if the device in question is susceptible to SEE caused by direct ionization of low-energy protons. In this case, a transport calculation through a

detailed model of the spacecraft may be required using NOVICE or another Monte Carlo transport package [28].

Secondary particle environments are also important when considering nuclear interactions. In such interactions, ions are produced in a range of species (Z), energies, angles relative to the incident primary particle and LETs. These distributions are determined using Monte Carlo codes that simulate the processes of nuclear cascade and evaporation. Examples include recoil ions generated by interactions between Si nuclei in the device SV and high-energy protons from the space environment [29-32], recoil ions generated by light ions colliding with high- Z materials (e.g., W vias) in hardened SRAMs [33] and the discovery in 2015 that protons can generate heavy ions by causing high- Z materials in packaging to fission [34,35]. Although such nuclear interactions are usually rare, there are lots of protons in many space environments, and this can result in SEE rates being dominated by such proton-nuclear processes. As an example, Figure 2-8 compares the relative flux vs. LET distributions for GCR ions with those generated by interactions of proton fluences in the GEO and ISS environments interacting with a 1- μm layer of Au ions. The relative fluxes suggest the threat from fission products could dominate that due to GCR for devices with onset LET $\text{LET}_0 > 15 \text{ MeV}\cdot\text{cm}^2/\text{mg}$ for GEO and above $10 \text{ MeV}\cdot\text{cm}^2/\text{mg}$ for ISS. However, fission daughter products—as well as recoil ions from proton and light-ion scattering—tend to be low energy and have short range, so overburden and SV dimensions significantly affect this conclusion. The SEE threat depends on device susceptibilities as well as the mission radiation environment.

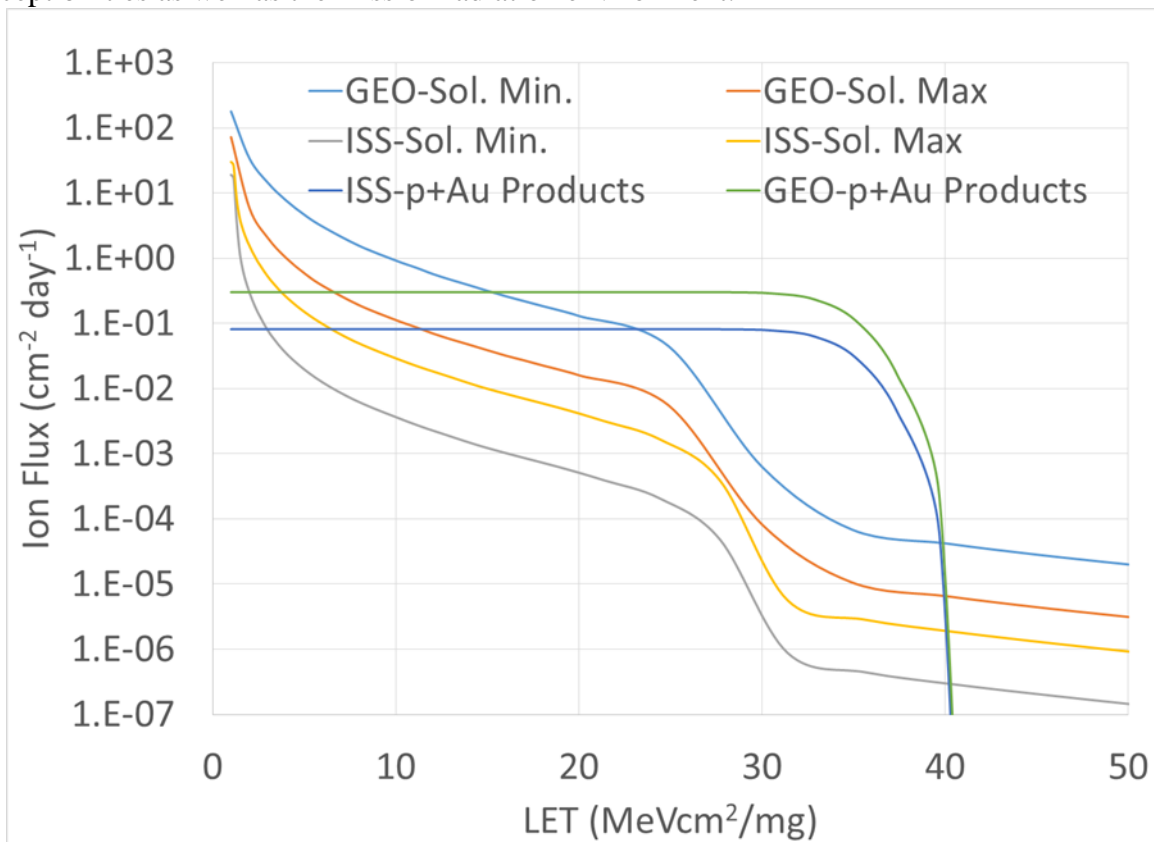


Figure 2-8: Comparison of GCR ion LET flux with that from proton on Au fission fragments for mission proton environments (averaged over mission) incident on 1- μm Au foil for the environment at the ISS and in geostationary orbit.

3 SEE Mechanisms and Consequences

Since their discovery 45 years ago, SEE mechanisms have proliferated into a veritable alphabet soup of acronyms. There are destructive SEE (DSEE), including single-event latchup (SEL), single-event burnout (SEB), single-event gate rupture (SEGR), single-event dielectric rupture (SEDR), single-event stuck bits (SESB) and so on. There are nondestructive SEE (NDSEE), including SEU, multibit upset (MBU—bits in same logical word) and multi-cell upset (MCU—bits in different words), block errors, column errors, row errors, single-event functional interrupt (SEFI), and even multiple types of single-event transients (SET)—digital and analog (DSET and ASET). Further complicating the picture, different papers in the literature may use different definitions for SEFI or MBU. To strip away the confusion, it is useful to consider SEE in terms of the characteristics that are important for risk identification, assessment and mitigation:

- 1) Direct consequences of an SEE are confined to the die in which it occurs. This does not mean that consequences of an SEE (e.g., errors, overcurrent, etc.) cannot propagate and affect other components in the system.
 - a. Exception: Independent SEE can be caused by the same ion in parts where die are stacked on top of each other.
- 2) SEE are Poisson processes. This means that:
 - a. The rate is constant as long as the environment is constant (on average), so the expected number of SEE in a given interval depends only on interval length.
 - b. SEE rates are additive, so that if a part in a system (or satellite in a swarm) has rate λ , a system with N such parts has rate $N\lambda$.
 - c. Times between SEE are distributed exponentially.
- 3) SEE can have the following consequences at the level of the part in which they occur:
 - a. They can produce a transient disturbance in normal function of the part that recovers automatically after a short time (SET on ps to ms time scales)
 - b. They can cause errors in data that can be corrected (rewritten) in the part (e.g., and SEU or MBU)
 - c. They can corrupt large amounts of data, although the data can be corrected (rewritten) at least in principle (e.g., block errors, column errors, etc.).
 - d. They can cause permanent errors that cannot be corrected within the part (e.g., stuck bits).
 - e. They can interrupt normal functionality of the part, requiring either intervention (e.g., power cycling and re-initialization) or long recovery times to restore full functionality (e.g., SEFI). Large amounts of data may also be corrupted.
 - f. They can result in partial or complete failure of the part (e.g., DSEE, including SEL, SEB, SEGR, etc.)
- 4) SEE consequences at the system level depend on the function for which the part in which it occurs is responsible. A destructive SEE in a part responsible for a trivial function may be less consequential than a SET while a part is executing a critical function. Reference [36] discusses the implementation of a SEE Criticality Analysis—analogue to a failure modes, effects and criticality analysis (FMECA).

We next consider some of the characteristics important for risk identification, assessment and mitigation for particular destructive and nondestructive SEE modes.

3.1 Destructive SEE Modes

Destructive SEE pose special challenges for SEE hardness assurance, not just because of their potentially severe consequences, but because understanding of a phenomenon that is inherently destructive—or at the least, very disruptive—will usually be based on poor statistics and because uncertainties may remain about the mechanisms of the effect. However, not all destructive SEE modes pose equal challenges. In some cases, destructive failure can be circumvented during testing by current limiting or power cycling when a high-current is detected. This allows decent statistics to be accumulated. And the mechanisms of some SEE modes are better understood than others. We start with one of the most common and best understood modes, SEL.

3.1.1 Single-Event Latchup

Single-event latchup is a complex, regenerative, parasitic failure mode that has been observed in CMOS microcircuits that have a thyristor or Silicon Controlled Register (SCR, *pnpn*) structure (see Figure 3-1). We can view this structure as two parasitic bipolar junction transistors connected collector to base of the other, resulting in a regenerative gain of the current flowing when charge from an ion is injected into the system. The issues relevant to SEE risk management include:

- 1) The consequences of the high-currents associated with SEL can result in
 - a. Catastrophic thermal failure
 - b. Latent damage, in which metallization is melted, dielectrics cracked [37].
 - c. Recoverable loss of functionality and high current state, but no failure or latent damage—a so-called nondestructive SEL. Demonstrating that no damage has occurred is challenging, requiring extensive SEL testing, and a post-SEL investigation involving microscopic inspection of the die and a burn-in type life test to realize any latent damage that may be present [38].
- 2) Worst-case conditions for SEL occurrence include
 - a. Higher voltage (testing usually done at worst-case operating voltages) [39]
 - b. Elevated temperature (should bound WC application temperature) [40]
 - c. Ions with sufficient energy/range to penetrate to the substrate of the device, usually several 10s of microns. This means that ions with short range—e.g. proton-Si recoil ions, ²⁵²Cf fission products or the products of proton-induced fissions of high Z materials—are likely to underestimate SEL rates for susceptible devices, and may not cause SEL at all. [41, 42]
- 3) A SEL mechanism can be initiated at cryogenic temperature, despite the loss of charge carriers due to freeze-out. The mechanism results when charge carrier mobility increases sufficiently that impact ionization can cause avalanches of charge carriers [42]. Data on this effect are limited, but susceptibility appears to occur below about 30 K.
- 4) Because SEL is a parasitic effect and circuit layout is important, data for similar parts has limited value for assessing SEL risk. Reference [43] found that SEL rates could vary several orders magnitude for similar parts fabricated in the same process, and the even onset LETs for latchup did not appear to form a well behaved distribution.
- 5) Mitigation strategies that may be effective against SEL include:
 - a. Avoidance or SEL susceptible parts. In some cases, rates may also be reduced by avoiding conditions (e.g., high temperature and voltage) where SEL rates are high.

- b. Cold-spare redundant parts can be included to replace parts that have failed due to SEL. Note that may require the failed parts to be isolated from the system, which can complicate system architecture as well as increasing system weight and size.
- c. Current limiting and power cycling when a rise in current is detected followed by a reset/re-initialization of the part. This can introduce the possibility of spurious resets, and in some cases, the current rise is too rapid to avoid damage to the part.

SEL mechanisms are sufficiently well understood that SEL rates can be estimated using standard rate estimation routines, such as CREME96.




Figure 3-1 Single-event latchup occurs when an ionizing particle injects sufficient charge to turn on a parasitic thyristor in a CMOS structure. Once turned on, the SEL is self-sustaining, and resulting high-currents can cause thermal failure in the device.

3.1.2 Single-Event Gate Rupture (SEGR)

SEGR is an inherently destructive failure mode that occurs when a power metal oxide semiconductor field effect transistor (MOSFET) is in its OFF (nonconducting) state. In SEGR, an ion passing through the gate of the MOSFET generates sufficiently high electric fields to break down the gate oxide. SEGR actually involves a two-step process (see Figure 3-2): First the ion passes through the gate oxide, weakening it and reducing the voltage needed to break down the oxide. Next, the ion passes through the channel of the MOSFET, generating charge, which is separated by the vertical field so that the holes pile up under the weakened gate, augmenting the electric field across the gate oxide due to the gate-to-source voltage V_{GS} until the gate dielectric breaks down. The resulting rush of current causes the MOSFET to fail thermally. The following characteristics are important when managing risk for SEGR in power MOSFETs [44]:

- 1) SEGR is inherently destructive. As such every data point represents the destruction of a MOSFET. Nondestructive events may be seen during heavy-ion irradiation (especially when $V_{GS}=0$) that cause current to jump discontinuously, but leave the device still

functional. While these are thought to involve gate-oxide damage, they do not constitute a “nondestructive SEGR” and so cannot be used to facilitate SEGR rate estimation.

- 2) Worst case conditions for realizing SEGR include:
 - a. MOSFET OFF ($V_{GS} \leq 0$ for N-Channel or $V_{GS} \geq 0$ for P-Channel)
 - b. Ion incident normally to the die (for a vertical device the geometry is most likely to pile holes under the weakened portion of the gate oxide)
 - c. High $|V_{DS}|$ and $|V_{GS}|$
 - d. High Z ions [46]
 - e. Ion energy/range sufficient to penetrate to the device substrate (may be 10s of μm). Even more so than for SEL, short range ions (proton-Si recoils, high-Z fission events) are unlikely to reveal SEGR susceptibilities.
- 3) Rate estimation is not possible due to the complicated dependence of SEGR susceptibility on ion Z, energy and angle, coupled with part-to-part variation for a failure mode where each new data point requires a new part. Instead, the product of SEGR testing is a “safe-operating area” of VDS and VGS values where SEGR susceptibility is negligible. (See Figure 3-3.)
- 4) Mitigation strategies for SEGR are limited:
 - a. Avoiding parts susceptible to SEGR for the application conditions for the mission.
 - b. Ensuring VDS and VGS are well within the safe-operating area
 - c. Redundancy may work in some applications, but will likely result in lower efficiency and more complicated circuitry.
- 5) Commercial MOSFETs tend to fail due to SEB rather than SEGR, while most radiation hardened MOSFETs fail due to SEGR. Exceptions exist.
- 6) “Similar” parts may introduce new mechanisms. Si TrenchFETs, SiC MOSFETs, GaN MOSFETs degrade under exposure to heavy ions. MOSFET-like structures in Flash memory charge pumps can fail destructively when erasing or writing memory cells.




Figure 3-2: Single-event gate rupture occurs when an ion traverses the gate oxide, weakening the dielectric, and then generates sufficient charge in the channel that when the charges pile up under the gate, the resulting electric field is sufficient to cause breakdown of the weakened dielectric [45].

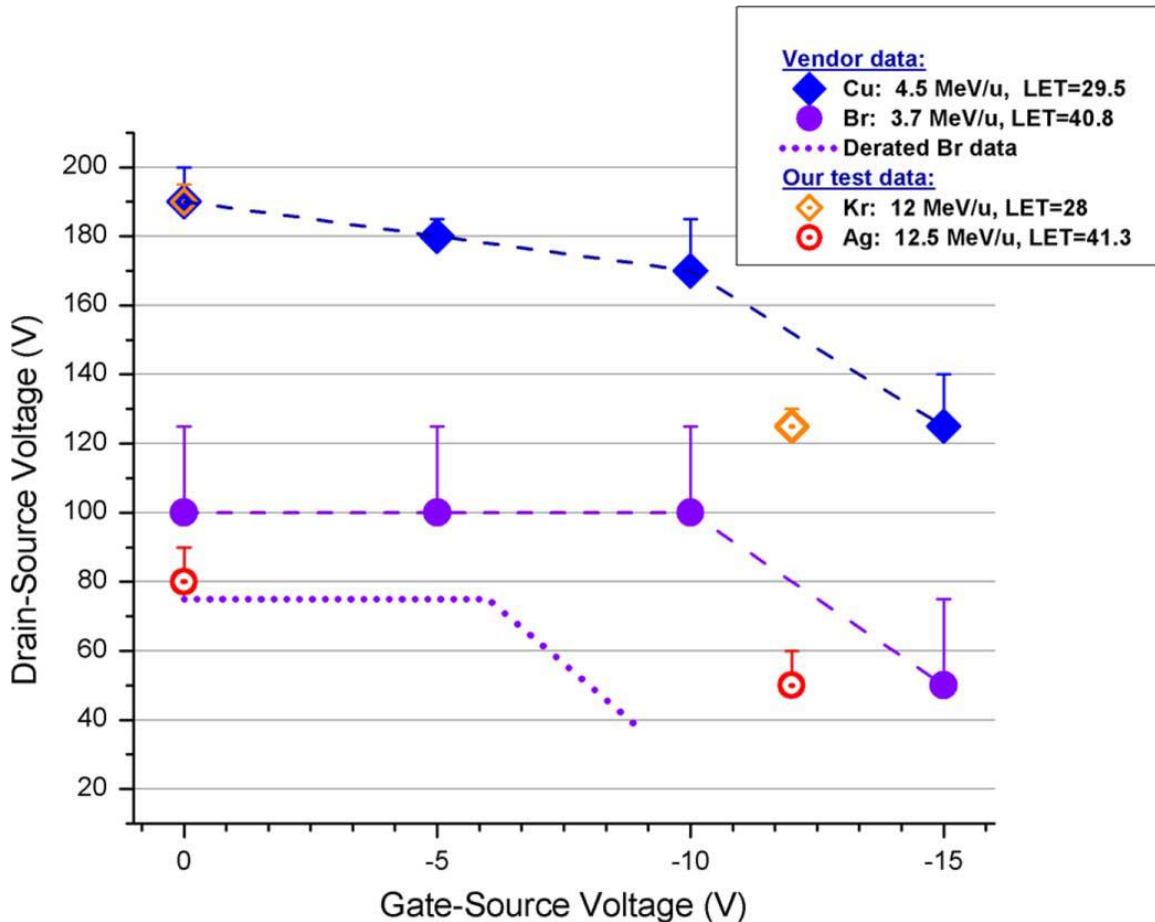


Figure 3-3: The usual product of SEGR testing is a graph of “safe operating” conditions or, more conservatively, a device response curve—for the MOSFET—VDS and VGS conditions where SEGR risk is minimal for the ion used. This device response curve for a radiation-hardened N-channel shows that both ion species and ion energy/range affect device responsibility [46].

3.1.3 Single-Event Burnout

Like SEL, SEB is a parasitic bipolar effect. It can occur in power MOSFETs, bipolar junction transistors (BJT) or field effect transistors (FETs). In SEB, a parasitic BJT in the transistor structure gets turned on by charge injected by the ion. (See Figure 3-4.) If the voltage across the transistor is sufficiently high, charge carriers can become sufficiently energetic to generate additional carriers via impact ionization—a process called second breakdown—and the resulting high-current state can destroy the transistor. Unlike SEGR, SEB is not necessarily destructive. The process can be stopped by current limiting. Characteristics of SEB important for SEE risk management include [44]:

- 1) Worst case conditions for SEB are similar to those for SEGR:
 - a. MOSFET OFF ($V_{GS} \leq 0$ for N-Channel or $V_{GS} \geq 0$ for P-Channel)
 - b. Ion incident normally to the die
 - c. High $|V_{DS}|$ and $|V_{GS}|$
 - d. High Z ions (effect not fully understood) [47]
 - e. Ion energy/range sufficient to penetrate $\sim 30\mu\text{m}$. As with SEL and SEGR, short range ions (proton-Si recoils, high- Z fission events) are unlikely to reveal SEB susceptibilities fully

- 2) Thermal latent damage is possible.
- 3) The fact that SEB can be stopped before it destroys the part makes it possible to accumulate statistics for the same part. This means that σ vs. LET can be measured, and, in principle, rates could be determined. In practice, the complex dependence on ion species, energy and angle makes rate estimation impractical. The normal product of SEB testing is similar to that for SEGR—a “safe-operating area” of VDS and VGS values where SEB susceptibility is negligible.
- 4) Mitigation strategies for SEB are mostly similar to SEGR
 - a. Avoiding parts susceptible to SEB for the application conditions for the mission.
 - b. Ensuring VDS and VGS are well within the safe-operating area
 - c. Current limiting can effectively limit SEB risk, but can compromise efficiency
 - d. Redundancy can be effective in some applications
- 5) SEB is the dominant SEE failure mode for commercial MOSFETs and related devices
 - a. Some parts may exhibit very low voltages for onset of SEB (e.g., 22% of the rated VDS for the International Rectifier IRF640) [48]
 - b. Part-to-part variability can be significant for some commercial devices [49].




Figure 3-4: Single-event burnout occurs when a parasitic bipolar in a MOSFET is turned on by an ion strike. If the voltage is sufficiently high and a source of current on the collector sufficiently large, the charge carriers can avalanche via impact ionization (second breakdown) resulting in device failure.

3.1.4 Other Destructive SEE

Although SEL, SEGR and SEB are the main destructive SEE threats for most technologies, there are other potentially destructive SEE that may affect some specific part types. Single-event

dielectric rupture (SEDR) is similar to SEGR in that parts are most susceptible when the ions are incident normal to the die. It has been seen in some one-time programmable field programmable gate arrays (FPGAs) [SEDR] and in MOS capacitors in some linear bipolar devices [50]. Like SEGR, the main mitigation technique for SEDR is avoidance—either by using parts hardened to the effect or avoiding conditions where the parts are susceptible.

Single-event snapback is yet another potentially damaging parasitic bipolar effect, mainly in SOI devices. Although snapback can generate currents sufficiently high to cause damage, it is usually not a major threat to reliability [44].

Stuck bits have been seen in a variety of devices—mainly memories. They occur when an ion deposits sufficiently large amounts of charge in a gate oxide of a transistor that it can no longer be programmed. The bit then behaves like a permanently upset bit. Stuck bits have remained sufficiently rare, and mitigation is the same as that for SEU, so to date, they have not been a major issue.

The observation in 2013 that Schottky diodes could fail due to SEE raised significant questions about some basic assumptions of SEE HA [51]. In most cases, diodes were considered to be immune from anything more pernicious than extremely short transients. Although the mechanism for this effect is not well understood, the failure mode is sufficiently common that risk avoidance should be the main mitigation approach. To date, no failures have been seen in parts derated to 50% or below of their rated voltage [52]. If such derating is not possible, testing is recommended.

3.2 Nondestructive SEE

After the previous discussion of destructive SEE, one might be tempted to heave a sigh of relief at the sight of the word nondestructive in the heading above. That relief would be premature. A system can be taken off line as surely by an SEU or SET as by an SEL. The consequences of an SEE depend on the criticality of the function compromised by it—that is, by the function the device was performing when the SEE occurred. The type of SEE that occurs depends on the device in which it occurs. Table I summarizes the main nondestructive SEE modes, the vulnerable technologies and mitigation/recovery strategies.

The risk posed by an SEE depends on the rate at which it is expected to occur and the consequences of the SEE when it occurs. The consequences of the SEE depend on how the SEE affects the struck part and on the part was executing when the SEE occurred. Thus, one must follow the propagation of the error downstream to its ultimate effect. This is especially true for SETs.

Table I: Nondestructive SEE Modes

Mode	Definition	Mode occurs in	Mitigation/Comments
SET	Temporary, self-recovering disturbance in device output or function	Analog or digital devices, usually without bi-stable elements; inputs determine output(s) in absence of transient	Multiple sampling, temporal voting, capacitative filtering
SEU	One or more bit flips in device; Correct value can be rewritten (if not it is a stuck bit)	Digital, data storage cells, including bi-stable devices, DRAM cells, Flash cells	Error correction code, voting multiple outputs or devices, duplicate compare w/ retry...
MCU	Multiple bit flips by the same ion, but not necessarily in the same logical word; simple error code sufficient if bits not in same word	Same as SEU	Same as SEU
MBU	MCU with bits in the same logical word;	Same as SEU	Same as SEU, but requires more complicated ECC or voting bit-by-bit and more complicated architecture.
Block Error	Error in ADDRESS or other control logic corrupting large data blocks; normal part function not affected	Same as SEU, but usually having large data arrays and can affect Column, Row or Block of address	recovery same as for MBU, but done many times over
SEFI	Disturbance or cessation of normal functionality of device requiring intervention for recovery; Note: SEFI can often be viewed as an SEU that upsets control logic	Usually occurs in devices with complex control logic (DRAMs, FLASH, FPGA, processors...), but can occur in fairly simple devices (e.g. ADC)	May recover after part re-initialized or may require power cycle and then re-initialization; data and configuration likely lost; if power cycle required, requires significantly more complicated design or accept higher data loss.

Since the SET by definition lasts only a finite time, the effects at the system level depend on how the transient affects devices downstream of the output of the struck device. Thus, specifying the consequences of an SET can be complicated, depending not just on the SET rate, but also on its duration and amplitude, as well as the functions of devices downstream of the struck device and their state of operation when the SET occurs. An idea of the complexity of the duration-amplitude behavior can be gained from Figure 3-5, which shows an amplitude-duration scatterplot for the Linear Technologies LTC6268 operational amplifier (op amp) irradiated by ions with several different LETs. Although the vast majority of transients last less than 10 μ s, a few transients have both high amplitude (>2 V) and long duration (>1 ms). What is more, these long transients are seen for LET values as low as 4 MeV-cm²/mg. Whether these transients constitute a threat to the mission depends on the application of the device. This plot indicates some of the challenges in risk management for transients. Mitigating the <10 μ s transients with

capacitive filtering or multiple samples of the readout would be straightforward and would likely not reduce performance of the device unacceptably. However, the rare larger, longer transients are more likely to have adverse consequences and would be much harder to mitigate. Would the mitigation be worth it, or would it be better to accept the consequences of a lower probability event? These are the types of questions one has to answer when dealing with transients—or with nondestructive SEE generally.

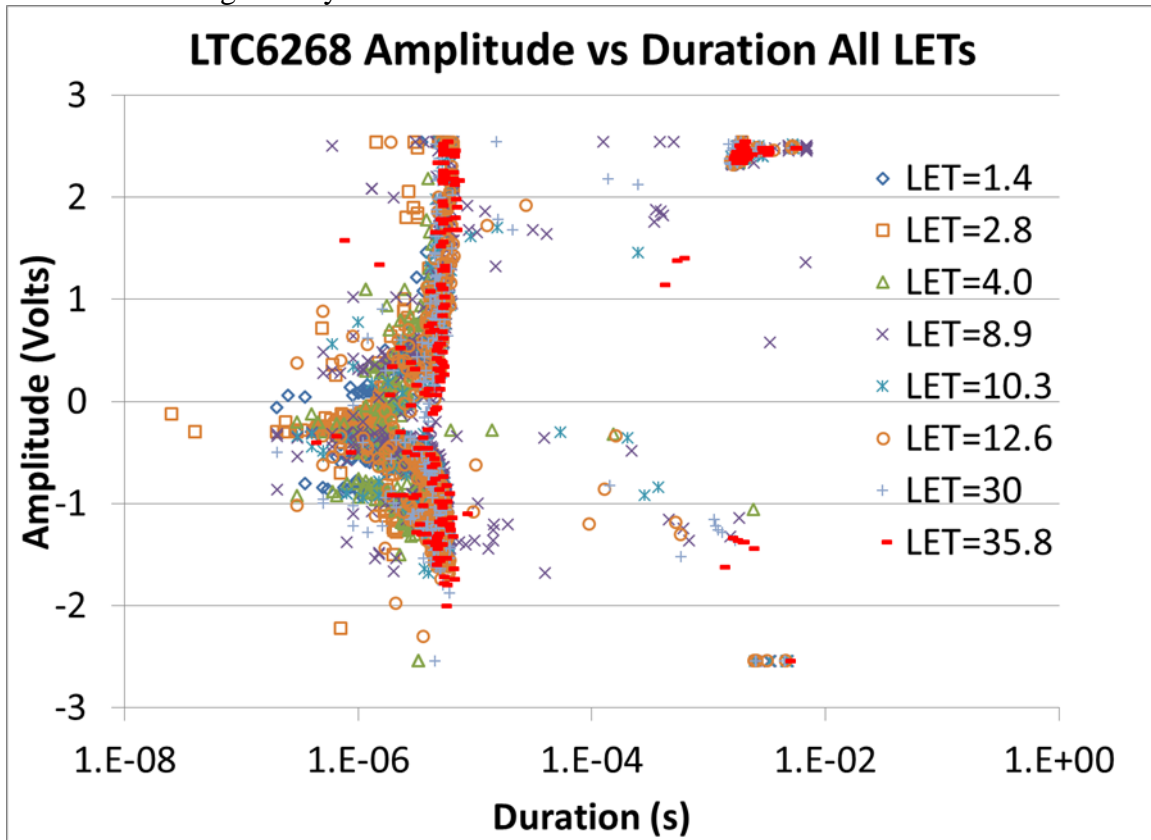


Figure 3-5: Single-event transients (SET) are temporary disturbances to the output of a device caused by an ion strike. Although their waveform may be complicated, they are usually characterized by a duration and an amplitude. Although most SETs are short—ns to 10s of μ s—some devices, such as the LTC6268 exhibit transients that can last several milliseconds.

Similarly, in dealing with a SEU, a block error or a SEFI, accurate risk assessment requires following error propagation or loss of functionality to the system level as well as the SEE rate (determined from the σ vs. LET curve), and the consequences in the struck device. A block error may seem much more severe than an MBU, but the overhead needed for mitigation is the same (it will just be much more active with a block error). Further complicating the situation, a single part may be susceptible to several or all of these SEE modes. (See Figure 3-6.)

There are two competing factors that determine the consequences of an SEE susceptibility:

- 1) The direct consequences of the error/failure mode (#bits upset, loss of availability, etc.)
- 2) The penalties incurred through use of the mitigation in terms of cost, performance, etc.

The goal is to find mitigation strategies that limit consequences of the SEE to acceptable levels without inflicting unacceptable penalties on the performance, schedule or cost of the system—and what is unacceptable to one mission may not be to another. This will be discussed later in the section on mitigation. At present, the important considerations for the next section on threat

identification are the relative costs of mitigating different SEE modes. Mitigation of short transients (e.g., via capacitive filtering) and single-bit SEU (e.g., via voting or a Hamming or other simple error correction code (ECC)) is relatively straightforward. If recovery from a SEFI requires only re-initialization of the part (e.g., by rewriting mode registers of a synchronous dynamic random access memory (SDRAM) or reloading configuration memory for a programmable FPGA), that may also be relatively simple as long as the loss of data and loss of availability during the recovery process are not an issue. If a part is susceptible to multibit upsets or block errors corrupting more than one bit per word, mitigation becomes more complicated. Not only must the capability of the ECC increase, one may also have need a more complicated architecture (e.g., interleaving of bits from each word across multiple die) to avoid overwhelming the ECC.

Likewise, a transient $>100 \mu\text{s}$ may also be very difficult to mitigate for some applications. Generally, the most disruptive nondestructive SEE are SEFI requiring a power cycle and re-initialization to recover functionality. Such power-on/resets (POR) result in loss of all data stored in volatile memory and may require several minutes to execute. Moreover, it is often not practical to engineer a system so that only the affected die is power cycled, necessitating power cycling at the board or box level, greater data loss and longer recovery time. Destructive SEE are often the most difficult SEE modes to mitigate—so difficult that most missions require parts to be immune to these effects. Often identifying destructive SEE susceptibility is the first priority in SEE analysis.

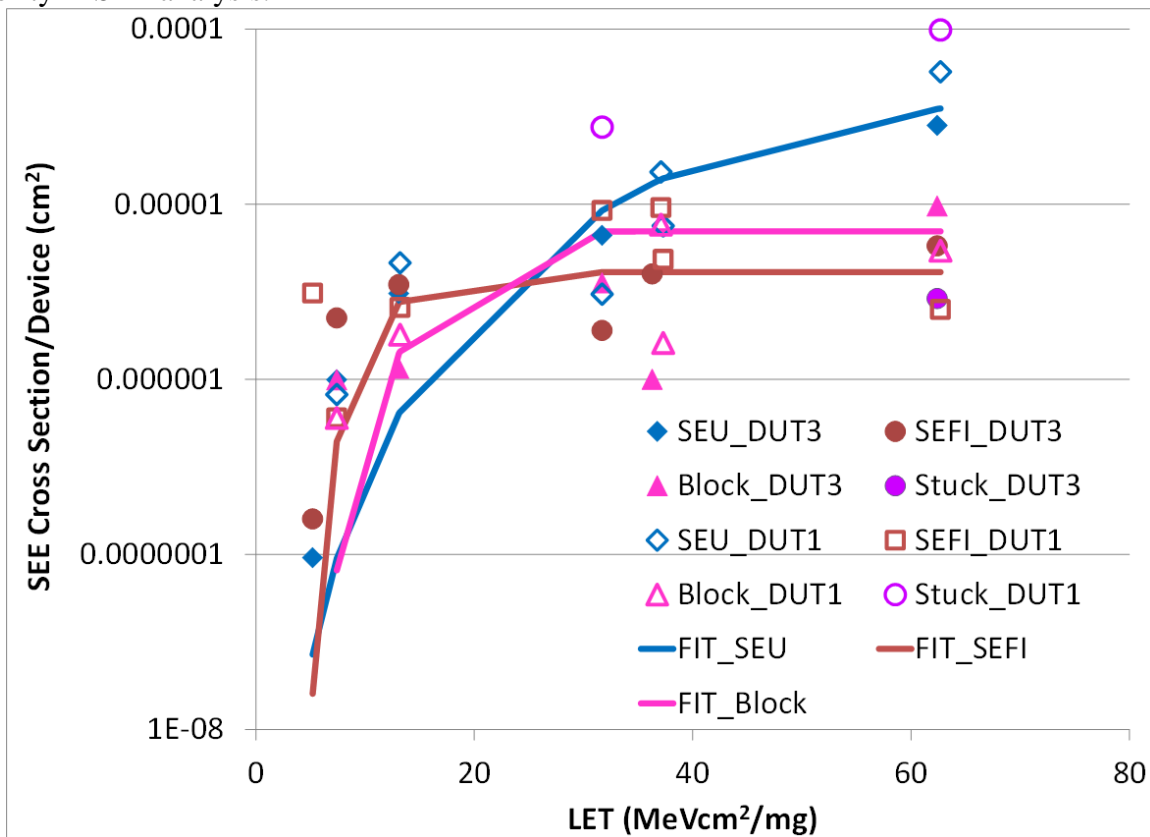


Figure 3-6: Nondestructive SEE can have a range of consequences—from single or multiple bit flips to corruption of large blocks of data to stuck bits to complete loss of functionality requiring a power cycle and re-initialization for recovery. SDRAMs, including the DDR3 device for which these data were taken, exhibit all of the above with comparable cross sections, making testing as well as operation challenging.

4 SEE Threat Identification

An SEE threat is credible when a potentially susceptible technology is deployed in a mission environment where the threat could be realized and the consequences of the threat could compromise the required functionality of the part. As such, the initial threat identification is made based on assessment of the mission requirements and environment and the technologies in the parts slated for the mission. To some extent, such an analysis is subjective, in that it depends on the body of knowledge available to the analyst when the determination is made. For instance, prior to 2012, Schottky diodes would likely not have been thought to be susceptible to destructive SEE. However, after they were observed to fail during heavy-ion testing of a DC/DC converter, it is now recommended practice to derate them to 50% of their rated voltage or to test them if they are to be used at higher voltages [51]. Similarly, we will not know whether low-energy protons are a concern until we know if the parts for the mission include deep-submicron commercial CMOS.

4.1 Mission Requirements

Top-level requirements for a mission lay out mission duration, performance and other conditions that must be met to achieve the mission objectives. Usually, the top-level requirements will be at a sufficiently high level that they will not even mention radiation effects. Rather, the relevant top-level requirements that will generate secondary and tertiary requirements pertinent to particular SEE related threats. For instance, mission life and reliability requirements will generate requirements pertaining to destructive SEE and catastrophic system-level failures due to nondestructive SEE. Availability requirements will generate requirements for recoverable but disruptive SEE, such as SEFI. Data accuracy requirements will ultimately cover data corrupted by SEE (e.g., SETs, SEU, block errors, etc.). In some cases, there may be no requirements that directly mention SEE, in which case, the SEE analysis relies on the discretion of the SEE analyst, responsible engineers and input from reliability experts.

Ultimately, requirements must be pertinent to mission objectives and to their parent requirements, attainable and verifiable. They should also not be so prescriptive that they constrain designers unnecessarily. This usually means that it is best to keep the requirement to general goals and allow designers to decide on how to meet the requirement. SEE requirements are usually verified through analysis, since SEE testing is destructive, and even if SEE test data are available, they must be interpreted using a model of the SEE mechanism to infer likely on-orbit performance. Attainability is also important. If a particular requirement is driving design and/or cost, it may be prudent to discuss with mission planners whether there is any flexibility in the requirement that could provide relief, lower costs or facilitate achieving the rest of the mission requirements.

4.2 Environments Revisited

Once the trajectory or orbit for the mission is known, the analyst can determine the likely environments the spacecraft will face and assemble the models needed to specify those environments. For the threat identification portion of the SEE analysis, the emphasis is on determining fluxes of particles likely to cause SEE for the duration and conditions of the mission and the desired levels of confidence. In terrestrial environments, the particles are neutrons—and possibly muons for very sensitive devices. For the most part in the space environment, the

particles of interest are high-energy protons and heavy ions. We have previously discussed the sources of these particles. For the risk identification portion of SEE HA, we are mainly interested in the adjustable parameters or features in the models of these different sources that can change particle fluxes, energy spectra and other factors that affect SEE rates. Some sets of requirements may be sufficiently detailed that these factors are specified in detail. If not, it is prudent to explore how much variability is possible given the specified requirements.

Galactic cosmic rays will be present in all space environments, and the current de facto standard model for GCR is the CREME96 model. (Note: Although other models exist, they do not significantly affect SEE rates [53].) In reality, the GCR environment is dynamic, depending on modulation of heliomagnetic and solar wind conditions as well as the flux of particles impinging from throughout the galaxy. In practice, however, the flux of GCR ions is sufficiently low that statistics are insufficient to model short-term fluctuations. As such, the main factors that affect the output of the CREME96 GCR model the choice of Solar Maximum or Solar Minimum conditions and, for orbits inside Earth's magnetosphere, whether geomagnetic conditions are quiet or stormy. Since space weather, launch dates and the durations of solar cycles can all vary, it is prudent to explore the effects such variability could have on expected rates. Solar Minimum conditions are worst case, and correspond the conditions observed during 1986-87 [20]. Generally, SEE rates for Solar Maximum and Solar Minimum conditions in geostationary orbit vary by a factor of 3-10, depending on the shape of the σ vs. LET curve (especially onset LET). For the ISS (as an example of an environment within the geomagnetic field), solar minimum rates exceed those at solar maximum by roughly a factor of 2-10. Specifying stormy space weather can result in rates 20% to >2x higher than quiet conditions depending on the σ vs. LET curve for the device of interest. One factor that is not in the CREME96 model is the variability of GCR fluxes from one solar cycle to another. The most recent solar minimum was the deepest seen during the space exploration era, with GCR fluxes from 14% (for protons) to 22% (ions w/ $Z > 2$) higher than during the prior solar minimum [54].

The analyst has more choices when it comes to solar particle events and trapped protons. If the model specifies AP8, the only choice is whether to choose the model for solar maximum or that for solar minimum. Confidence levels can also be selected by simulating several missions to reflect environmental variability and ordering the results. The results can be used to assess the severity of the proton environment (both in terms of flux and energy spectrum) for causing SEE.

For solar particles, there are two concerns. The first is the cumulative fluence of solar protons and heavy ions accumulated during the mission. The higher these fluences, the greater the expected SEE total during the mission and the average SEE per day. The second concern is the worst-case solar particle event. During solar particle events, SEE rates can spike by more than 1000x for soft devices and over 100x for SEE hardened devices. Such rates can overwhelm many mitigation schemes—especially those based on redundancy.

Unlike previous cumulative fluence models for solar protons [55-58], the ESP [59] and PSYCHIC [60] models allow cumulative fluences of solar protons and heavy ions, respectively, to be bounded at a desired confidence level. This allows designers to budget for average error rates in the design over the duration of the mission. Unfortunately, the solar particle flux is rarely average, a single solar particle event can account for a significant fraction of the particle fluence in a given year or even a given solar cycle. The concern for severe solar weather is that device-level SEE rates scale with particle flux, and if, for example, a redundant system fails after two independent upsets or failures, the system level failure rate will scale as the square of the device rate. There are two approaches to bounding solar particle fluxes. The CREME96 SEE

rate estimation package includes a solar proton and heavy-ion model for the October 1989 SPE. This event was particularly large and had a relatively high heavy-ion content, it is likely to bound worst-case SEE rates. However, it is not possible to say how conservative the bound will be. Several approaches have been developed to bound solar proton fluxes and fluences. Reference [61] modeled solar proton fluences for SPE as a truncated power law distribution, which has the important implication that fluences are bounded, and there is not a low-probability tail extending to ∞ . Reference [62] found that solar proton fluxes could exceed 1000 protons $\text{cm}^{-2}\cdot\text{sr}^{-1}\cdot\text{s}^{-1}$ for about 24 hours during an average Solar Maximum year, and that for about 2 hours of that year could be more than 10x higher. Solar heavy-ion fluxes are so much lower than solar proton fluxes that statistics on energy distribution tend to be poor. Usually, they are assumed to scale roughly with proton fluxes, leading to a conservative bound on the actual solar ion environment. However, Reference [63] compared heavy-ion fluxes for the large November 2001 event to the October 1989 event model in CREME96 and found good agreement out to LET~1-3 MeV- cm^2/mg , but the CREME96 model was more severe at higher LET values. This could suggest that the SPE model in CREME96 is conservative, or it may indicate that the October 1989 event was more enriched in heavy ions than that of November 2001.

4.3 Technology and SEE Vulnerabilities

As mentioned above, a SEE threat becomes credible when a vulnerable technology is used in an environment where the threat can be realized. This raises questions not just of how to determine which SEE modes occur in which technologies, but prior to that how to determine the technology of a particular part. Table II lists some electronics technologies and the SEE vulnerabilities associated with them.

Table II: SEE Vulnerabilities of Electronics Technologies

SEL	SEGR	SEB	SEDR	Stuck Bit	SEU/MCU	SET	SEFI
CMOS	MOSFET	POWER MOSFET	One-time Prog. FPGA	SRAM	Digital/bistable technologies	bipolar technology	Complex Microcircuits
Bipolar?	FLASH	Power JFET	Bipolar Microcircuits	DRAM	Deep submicron CMOS more MCU susceptible	Analog microcircuit	ADCs
	Schottky Diode	Power BJT		FLASH		Digital microcircuit	PWMs

Part-Level Consequences

- Catastrophic failure possible
- Destructive but limited
- Nondestructive

How Common is Issue?

- Common in technology
- Catastrophic failure possible
- Not seen but possible in principle

The list of SEE modes is not exhaustive...and even if it were, the rate of discovery of new SEE modes means it would not remain exhaustive for long. In many cases, determining the technology of a part is straightforward. Discrete components are usually straightforward, although there are exceptions (e.g., is a diode a Schottky diode or a super barrier junction device?). State of the art parts like processors and memories, etc., will almost certainly have CMOS (although of what feature size is another question). If a part has many operating modes, it is reasonable to expect complicated control logic and therefore SEFI modes. In some cases, a vendor may produce devices only with a single or a few different technologies, and it may be

fairly straightforward to ascertain which technology is used in a particular part, especially if one is familiar with the process and consults the datasheet for the part. Datasheets are often very useful, as simplified schematics often depict MOS transistors for CMOS, BJTs for bipolar and both for a BiCMOS. If the datasheet is not helpful, the vendor may be. This may be as simple as looking at the reliability data for the different processes. For example, Analog Devices Reliability page list the technology for most part numbers in their inventory at: <http://www.analog.com/en/about-adi/quality-reliability/reliability-data.html>. Otherwise, the analyst may be forced to try and find a knowledgeable and helpful source at the vendor. Moreover, the preceding discussion presumes that the SEE modes for the technology are known. This may not be the case, as new materials and device architectures are becoming increasingly common as the microelectronics industry strives to keep pace with Moore's Law despite the failure of conventional CMOS scaling over a decade ago.

If one is confronted with a truly novel technology, the best approach may be to use the closest analogue to the new technology as a rough guide for what to expect. The physics, structure and materials of the device can also give clues of what to expect and how seriously to take the threat. However, new technologies always run the risk of introducing new vulnerabilities and should be a high priority for heavy-ion testing to characterize any SEE modes.

4.4 Threat Identification: Mitigation vs. Robust Design

The previous discussion raises the question of where to draw the line between threat identification and threat evaluation. While it is true that a threat is identified as soon as the potentially vulnerable technology is slated for a use in an environment likely to realize the threat, not all threats are equal. If we know a part is fabricated in a technology that is highly susceptible to SEL, it should be a higher priority for attention than a part from a technology that has never exhibited SEL. Thus, the initial threat determination should use all information readily available, and if this allows threats to be ranked, evaluation and mitigation resources can be applied to maximize risk reduction for the mission.

Similarly, there are several reasons why it may make sense to design in mitigation for various SEE even before we know for certain that the susceptibility is present in the design. First, certain SEE modes are quite common in some part types—for instance if we have an SDRAM, we can expect it to exhibit SEFI block errors and SEU. Second, many mitigation measures ameliorate the effects of an SEE rather than reducing the chances of its occurrence. These can often be straightforward to implement and have little impact on performance or other factors. Finally, some mitigation strategies are much easier to implement early in the design process than later on. There is also the question of the criticality of the application and the requirements for the mission. If a particular function absolutely must work, it may be worth purchasing some insurance against failures whether they are realized or not.

Looked at in this way, the SEE hardness assurance seems less a three-step process than a recursive effort loosely broken into threat identification, evaluation and mitigation.

5 SEE Threat Evaluation

Threat evaluation is usually the process that takes up most of the time in an SEE analysis. Threat identification is usually straightforward for most devices. Threat mitigation approaches are fairly standardized (at least in concept, although they may pose challenges in implementation), but expensive. As such, the threat evaluation stage is needed to ensure that finite resources for mitigation are expended to reduce SEE risk in an efficient manner.

Threat evaluation is a process of gathering information that may change the initial threat assessment. The information may be about the part, the application, the mission requirements, the environment or anything else that affects mission risks. The information may be gathered by part testing, discussions with design and other responsible engineers, discussions with part manufacturers and suppliers. If test data for the part already exist, this historical data can be assessed for applicability, or if there are data for similar parts, they can be assessed to determine whether they place a meaningful bound on part performance. All of these efforts have costs, and the goal is to reduce risk via the most economical path. Often, the first effort is made in gathering as much information as possible about the application from responsible engineers and about the part from vendors and parts engineers. Not only is this effort fairly straightforward, it also provides useful information that will be useful in later assessment and mitigation efforts. SEE testing is often an expensive, so after obtaining a full understanding of the application, the next task is often the hunt for prior test data on the part or on similar parts.

5.1 Historical Data

If you are very lucky and have lived a virtuous life, you may find that someone else has tested the very part you want to fly. If life were fair, you would be done. Unfortunately, as parts have become more complicated, SEE test results have become more application specific. Unless the application is the same, or unless one understands how to extend the test results to the application in the current mission, the data can serve only as a rough guide for expectations of SEE performance. This has made it even more important to understand data in context and to understand the sources of the data.

5.1.1 Sources of Historical Data

Because radiation effects is a small, specialized field, sources for radiation test data have traditionally been limited to journals of record (e.g., IEEE Transactions on Nuclear Science) and conference and workshop proceedings (e.g., IEEE NSREC Radiation Effects Data Workshop, SEE Symposium, RADECS, etc.). The IEEE XPLORE search tool provides access to most of these resources. Also, Dave Hiemstra's summary of the parts with test data in the previous year's REDW appears the following year. Other conventional sources of radiation data include the radiation group websites for space exploration agencies (some require registration to look for data).

As SmallSats have proliferated, some have sought to increase their reliability in the space environment, so they have begun to produce radiation test data. Most of the test data uses only protons, but because SmallSats tend to use more COTS parts than conventional platforms, SmallSat conference proceedings are a good place to find what limited data may exist on such parts. In addition to the above suggestions, the right key words in a web search can sometimes turn up a resource the above suggestions would have missed.

Once the data are in hand, the analyst must evaluate it to determine whether it is representative for the flight parts and if so, what it says about likely flight part performance. Table III summarizes some of these source.

Table III: Sources of SEE Test Data

Data Source	Data Found There	Website Address
IEEE XPLORE	Mainstream Rad Effects Data (NSREC, TNS, REDW)	http://ieeexplore.ieee.org/
GSFC NASA Database	Test reports for GSFC Projects/Programs	https://radhome.gsfc.nasa.gov/
JPL RADATA	Test reports for JPL Projects/Programs	https://radcentral.jpl.nasa.gov/
ESCIES data base	Test reports for ESA missions	https://escies.org/labreport/radiationList
ARC Search Page	Publications of AIAA and other SmallSat related venues	https://arc.aiaa.org/search
Web search	You'd be surprised what you can find!	Your favorite web browser, Google Scholar may help

5.1.2 Using Historical Data

Having found historical data for the part being investigated, it is tempting to declare victory. However, quality checks on the data must be performed to ensure it is suitable for the application of the part (see Figure 5-1). Usually, part-to-part and lot-to-lot variation of SEE susceptibility are treated as negligible, and in most cases this assumption is borne out, especially for nondestructive SEE. Since destructive SEE involve parasitic structures and/or processes that do not occur during normal operation of the part, there is less confidence that variations—especially lot-to-lot—will be negligible. For example, in recent testing of two samples of a DDR3 SDRAM, the data suggested well over two orders of magnitude in the stuck bit rate [64]. Moreover, for destructive SEE, one also has the challenge that statistics may be poor, and—if statistics for each part are poor—it may be difficult to distinguish between part-to-part variation and Poisson fluctuations in fluence to failure. If part-to-part or lot-to-lot variation cannot be neglected, the qualification becomes much more complicated, and using historical data instead of testing is not recommended.

Even if we expect limited variability, there is still the question of whether the test is sufficiently general to cover the application for the current mission. As mentioned above, SEE test data are often application specific. Complicated devices such as DDR SDRAMs have so many different modes of operation that testing in all of them may not be feasible. SEE rates in processors will depend on the algorithm being executed, while SEE susceptibilities in programmable devices (e.g., FPGAs) will depend on device configuration. SET rates, amplitudes and durations can be strongly dependent on applied voltages. If one is operating the device in the same or a very similar manner as the test for which the data were collected, the data may still serve to bound expected SEE rates and effects.

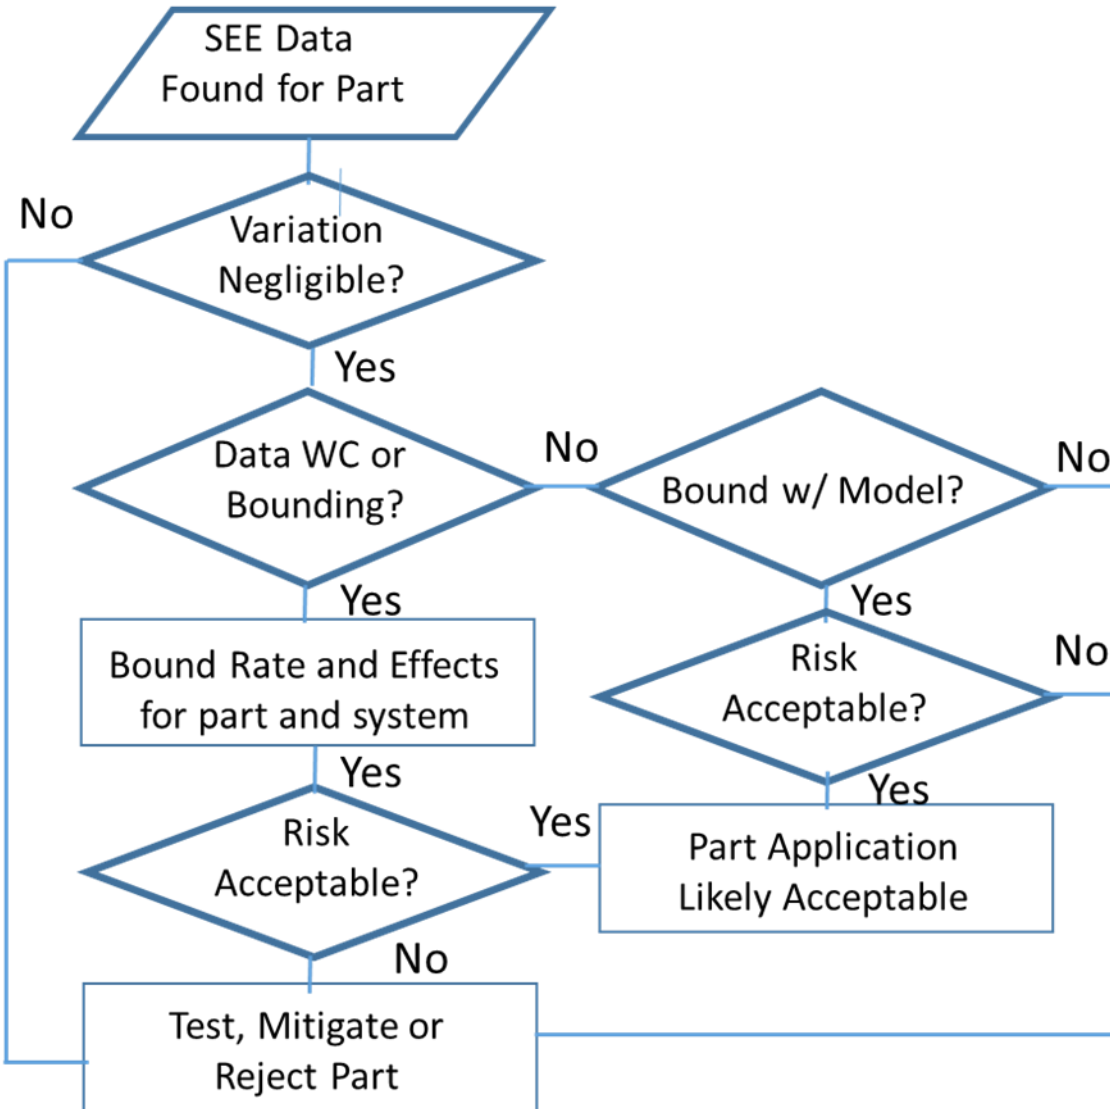


Figure 5-1: Flow chart for using historical data

If there are multiple test data sets enveloping mission application conditions, SEE performance may be bounded by constructing a statistical model of the variability. Alternatively, if the dependence of the SEE mechanism on the relevant application conditions is well understood, it may be possible to use a physics-based model to bound flight-part performance. Statistical models play a much greater role when discussing how to bound SEE performance for a part using data for other parts with similar designs fabricated in the same process.

5.2 Using Proxy Data to Bound SEE Susceptibility

A proxy variable for flight-part SEE performance of a device is a quantity that is not directly relevant to that performance, but which correlates with the SEE performance sufficiently well that we can infer it from the behavior of the proxy. Since, in reality, we can never measure flight part SEE performance until the mission is underway, any data we use to infer such performance is a proxy. The test conditions will always differ from the actual mission environment, and we will need to use a model to move from proxy data to the quantity that interests us. Even heavy-

ion SEE test data are collected in an environment very different from the space environment: accelerator flux rates are much higher, and the accelerator provides a single ion species at a time with a given energy and angle of incidence as opposed to the space environment where $1 \leq Z \leq 92$, energy can range up to several GeV/nucleon and the angle of incidence can vary over 4π steradians. The important question is whether ways the test environment differs from the mission environment are important for the quantity that interests us. For instance, knowing that the SV for SEL can extend $>30 \mu\text{m}$ into the substrate, we can anticipate that testing with low-energy ions from a ^{252}Cf source could underestimate SEL sensitivity. Similarly, because SEE testing is destructive, we cannot test the flight parts themselves for SEE susceptibility, but rather must test a representative test sample from which we can then infer flight performance using a model (yes, assuming that part-to-part and lot-to-lot variation are negligible is a model).

One advantage an analyst has in trying to use proxy data for SEE HA is that the goal is not to predict, but rather to bound flight-part SEE performance. This means that the distribution of proxy data need not be narrow, but must merely envelop the flight-part performance. Of course, if the distribution is too broad, it will not be possible to generate useful bounds, but this will usually be evident from the analysis. Figure 5-2 illustrates the situation. The SEE behavior of the flight parts may have some small variation (as indicated by the small diameter of the yellow circle). Variation of parts from flight lot or other lots may vary slightly more, but still (hopefully) negligibly and so can serve as a stand-in for inference of flight-part performance. Variation will likely be significantly greater for “similar parts” fabricated in the same process as the flight parts. Below, we will discuss the use of a statistical model to bound flight-part performance from similar parts. Whether such “similarity data” will produce useful bounds depends on how tightly the performance of the similar parts is clustered.

As one moves beyond the circle of similarity data in Figure 5-2, it can become more difficult to construct useful models for bounding SEE performance. Technology trends, expert opinion and the physics of the SEE mode or of the part can serve as a guide, especially if the SEE mechanisms and part functionality are well understood. We will next look at two proxies for heavy-ion SEE susceptibility in the space environment: similarity data and use of protons to constrain heavy-ion upset rates.

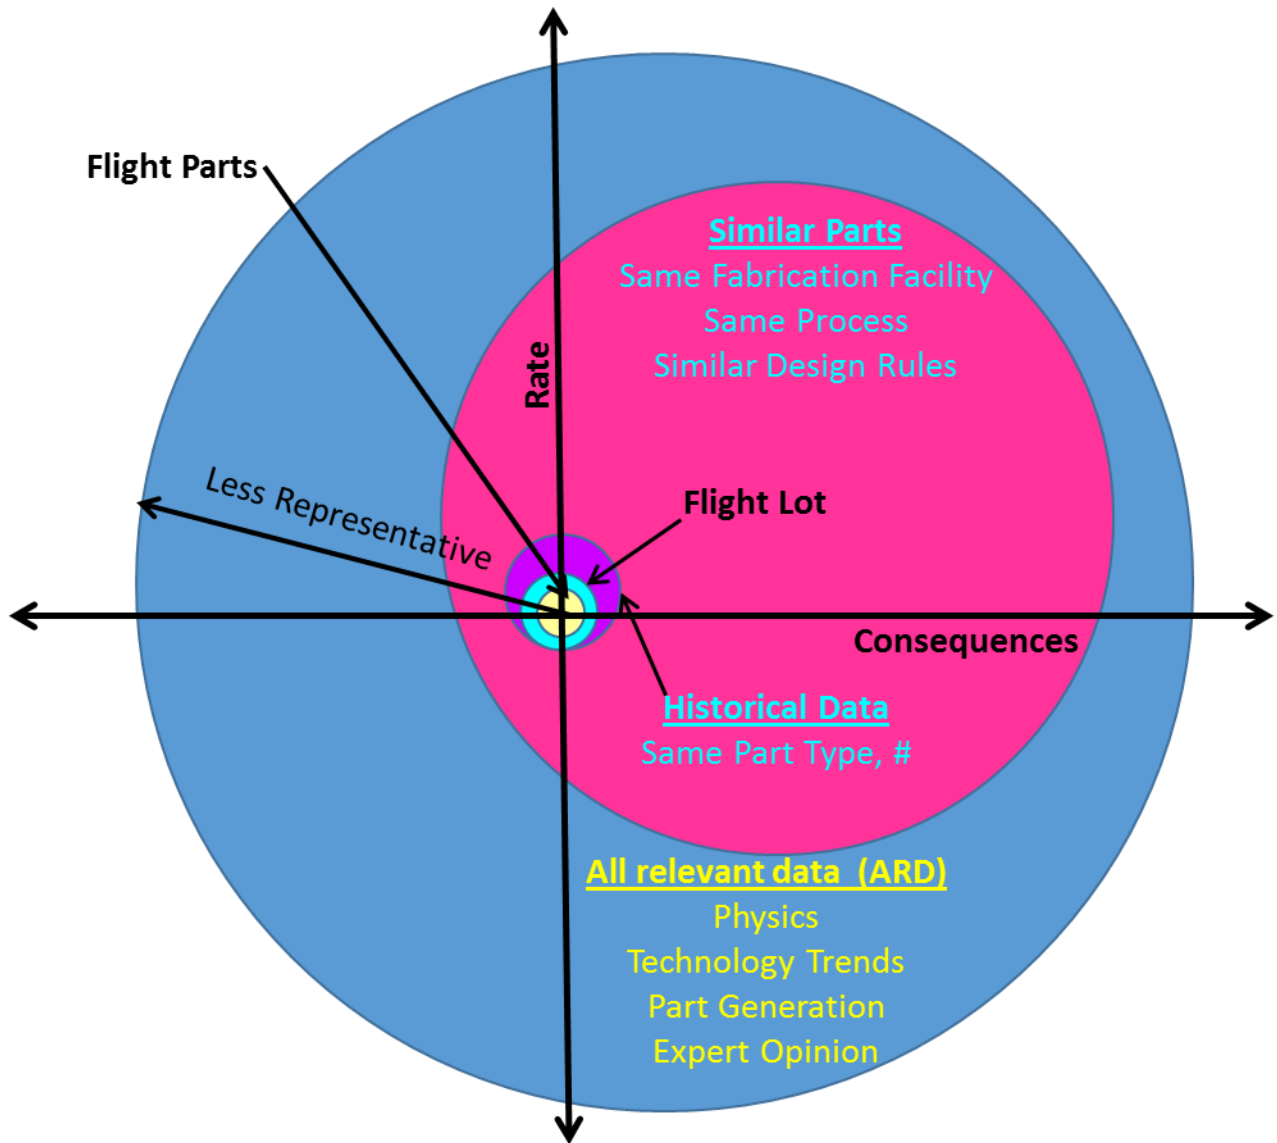


Figure 5-2: Although what we really want to know is how flight parts will behave, the destructive nature of SEE testing precludes knowing this until the mission is finished. As such, flight-part performance is inferred from less representative data. Test data on a sample from the flight lot, historical data for the same part are usually assumed to differ from flight part performance that they can stand in for the latter. Statistical analysis of less representative data for similar parts can also bound flight part performance. In some cases, the physics of the part or failure mechanism, technology trends or expert opinion can yield meaningful bounds as well.

5.2.1 Similarity Data

When data are not available for a part slated for use in a space mission, it is common to look at how similar parts have performed. However, what constitutes similar and how many different parts one should have in the dataset has rarely been defined. Is it sufficient for the parts to be fabricated in the same process (including minimum feature size) and foundry, or should the type of part be similar as well, and if so, how similar? Is a comparator be sufficiently similar to an operational amplifier? The answer to these questions depends on how the data are to be used and on the model to be parameterized by the data.

For instance, if one is simply interested in the question of whether SEL is possible for a given fabrication process, the analysis is minimal, and gathering as much data on as many parts as possible is desirable. Thus, looking at parts from Analog Devices' (ADI) 0.6 μm CMOS process and finding that the AD7664 latches up [43] is sufficient to establish that SEL is a threat for this process. On the other hand, one may look at data for dozens of parts fabricated in ADI's BiCMOS process without finding a single part susceptible to SEL [43], and this does not completely exonerate the process. However, it does increase confidence that the process may be immune, and perhaps the risk of not testing the technology could be worthwhile.

More quantitative bounds on SEE performance are possible if one looks at how data for several similar parts are distributed. Reference [43] used Bayesian probability analyses to look at distributions for SET rates and SET durations for two different families of operational amplifiers and constructed distributions that allowed both quantities to be estimated for each family at any required confidence. The analysis also looked at SET amplitudes, but found that transients going to the voltage rail are sufficiently common that assuming rail-to-rail transients is prudent. The analysis in [43] requires computation of the sample means and standard deviations for the rate and duration distributions, so the minimum number of similar parts that can be used for this technique is 3, and the quality of the results increases rapidly as the count of part types increases. Figure 5-3 illustrates the process.

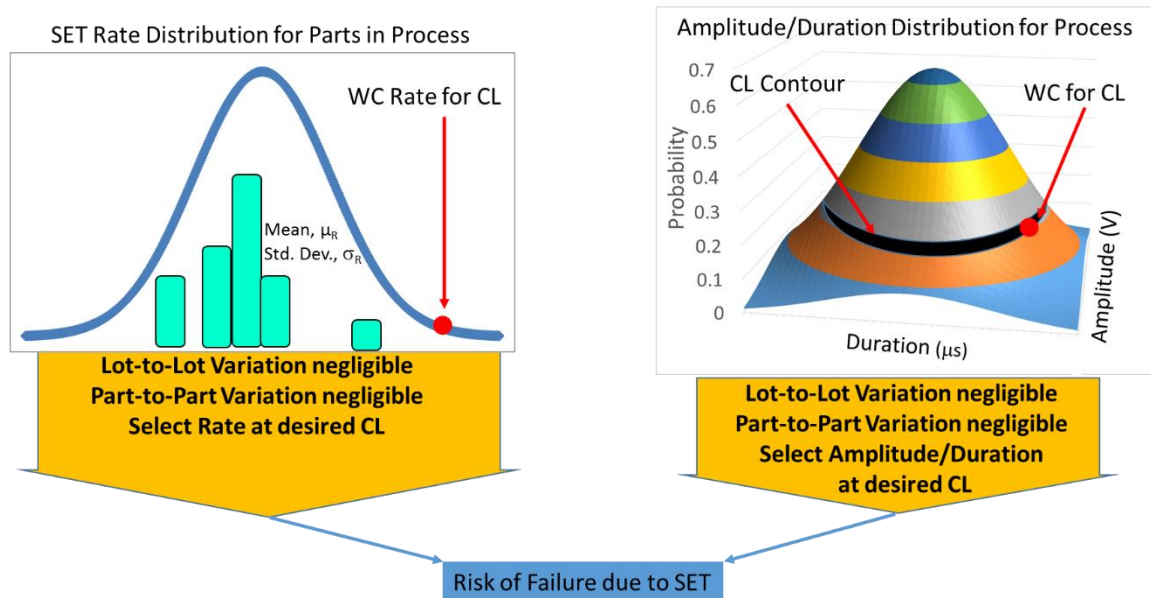


Figure 5-3: Determining SEE risk requires determining both the probability of an SEE and the consequences if it should occur. Fitting SET rates to a distribution and selecting the worst case rate for a given confidence level defines failure probability. Consequences can be determined by fitting SET amplitude and duration to a bivariate distribution and selecting the appropriate bounding SET profile (amplitude + duration).

Reference [43] also applied this method to SEL data in Analog-to-Digital and Digital-to-Analog converters. However, in this case, the distribution of consequences is uninteresting, as assuming destructive SEE to be destructive is usually prudent, and the distribution of SEE rates was too broad (spanning >2 orders of magnitude) to produce meaningful constraints for the process. This is likely because SEL is a parasitic process that depends on circuit layout, so whether a particular

element is susceptible is somewhat hit or miss. Even the distribution of onset LET for SEL was quite broad for both families—and perhaps even bimodal.

The previous discussion illustrates that the quantitative methods can yield useful quantitative bounds on SEE performance, and even when such bounds are not possible, they yield useful information about the parts in the process being analyzed (e.g., that SEL rates may not be amenable to statistical treatment using similarity data or that SET amplitudes should be treated as rail-to-rail for a similarity data analysis).

5.2.2 Protons as a Proxy for Heavy-Ion SEE Susceptibility

Although protons have been known to cause SEE via production of recoil ions from proton-Si collisions for over 35 years, the idea of using proton SEE data as a proxy for constraining heavy-ion SEE vulnerability was first developed in detail in [29]. The method was refined in [30-31, 65]. The basic idea is that the recoil ions produced by a proton collision are “heavy” ions and will leave ionization tracks through SV just like any other ion. However, as mentioned above, whether a proton recoil can be treated like “any other ion” depends on whether the differences between the recoil-ion and GCR/SPE spectra are significant for the SEE mode under investigation. One important difference between the proton-recoil and space heavy-ion environment is the low energies of proton recoils. Most recoil ions from 200-MeV protons are on the low-energy side of the Bragg peak, making use of LET as a metric problematic unless the sensitive volume is shallow. As such, proton recoil ions tend to underestimate destructive SEE susceptibilities due to the deep SV for these SEE modes [32]. One way of compensating for this tendency is to assign an equivalent LET (LET_{EQ}) to the recoil ions that takes into account the range of the ions relative to the SV depth

$$LET_{EQ} = \frac{E_{dep}}{(\rho_{Si} \times \text{depth}(SV))}, \quad (3)$$

where ρ_{Si} is the density of Si, and the depth of the SV is assigned based on what is known about the SEE mode and the technology. Figure 5-4 illustrates that LET_{EQ} has the effect of compressing the LET distributions to the left, where SEE cross sections are low, thus decreasing the effective fluence of particles that can cause the SEE mode. Using LET_{EQ} , [66] showed that a test with 10^{10} 200-MeV protons/cm² would fail to detect SEL in a susceptible device for all but the most favorable SEL σ vs. LET curves (that is, low onset LET and rapidly rising with increasing LET. Moreover, when the SV depth was >10 μ m, detecting SEL was highly unlikely regardless of proton fluence or energy.

Proton SEE testing faces other challenges as well. Because recoil ions are produced with $3 \leq Z \leq 15$, and a range of angles and energies, one cannot know the characteristics of the ion that caused a given SEE observed in the device. Not only does this complicate bounding SEE rates, it also makes it difficult to learn more about SEE mechanisms and impossible to optimize an SEE test to detect a particular SEE mode, since one cannot tune the beam to have worst-case ion characteristics for that mode. Also, because proton recoils have $Z \leq 15$, and produce few ions w/ $Z > 12$ the majority of ions have $LET < 12$ MeV-cm²/mg., leaving one effectively blind to SEE susceptibilities at higher LET or Z.

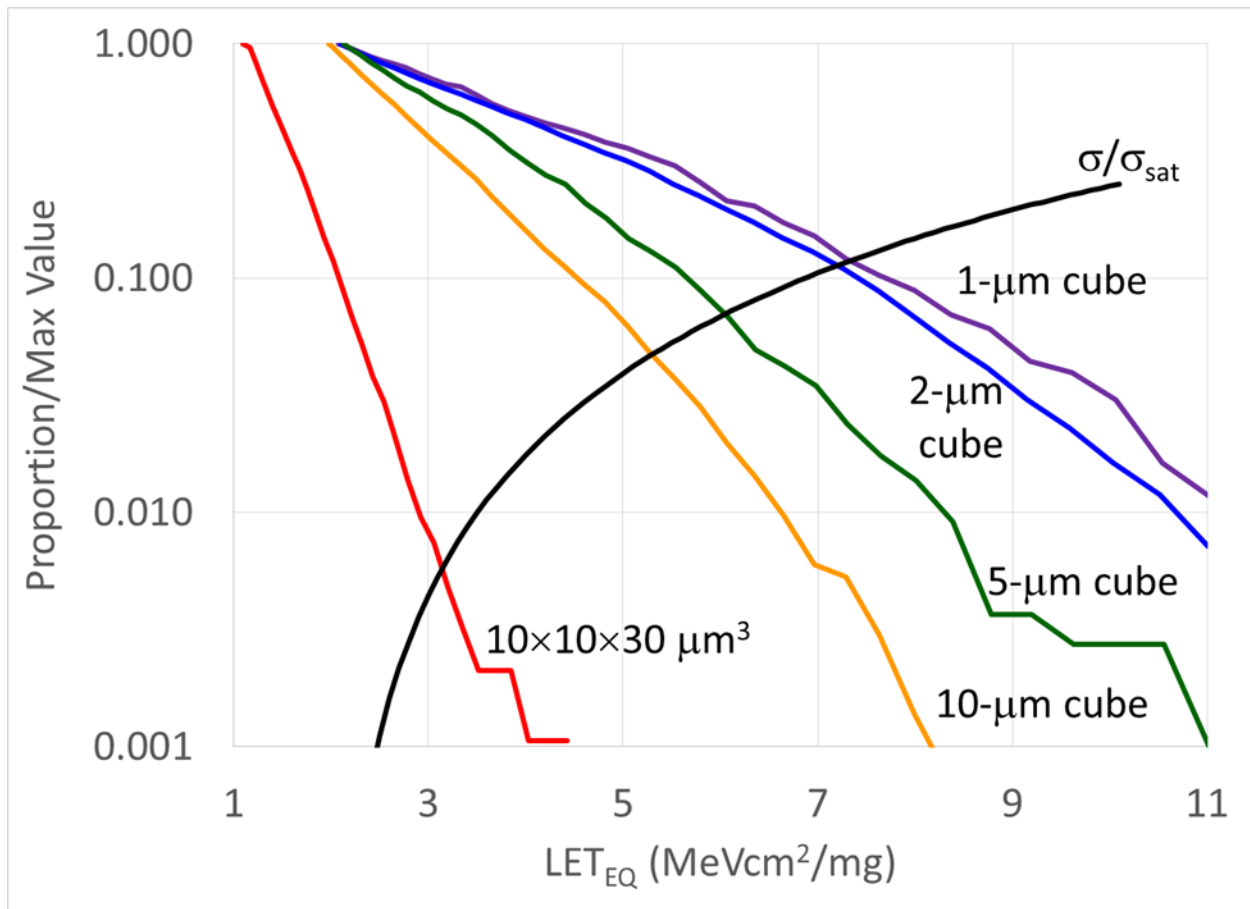




Figure 5-4: Equivalent LET, given by equation 3 facilitates comparison between range-limited proton recoil ions and GCR, SPE and accelerator ions. For deep sensitive volumes, LET_{EQ} pushes the fluence distribution to the left, where SEE cross sections are lower, presenting a more realistic picture of the efficacy of the ions at causing an SEE.

Perhaps the most significant limitation for proton testing is that only 1 out of every 289100 200 MeV protons generates a recoil ion, while every proton contributes to TID. Testing with recoil ions adds almost 200 times as much dose as would an equivalent ion LET spectrum coming from an accelerator. The result is that although the proton fluence used for testing seems high, the actual number of ions probing the susceptibilities of the device is quite low—e.g. about 34600 ions/cm² for a fluence of 10¹⁰ 200-MeV protons/cm², or one ion every 2891 μm². Figure 5-5 illustrates this situation visually by depicting simulated recoil ion strikes superimposed on an infrared micrograph of an Elpida EDS5108 512 Mbit SDRAM. A fluence of 10¹⁰ 200-MeV protons/cm² generates only 3 ions in the area shown (an upward fluctuation from a mean of 1.45) and does not come close to covering all of the discernibly different features in the micrograph, while a typical heavy-ion test run fluence of 10⁷ ions/cm² does a good job of coverage.

The recoil-ion fluence—and therefore the coverage of the test—increases with the proton fluence. However, so does TID—a fluence of 10⁷ ions/cm² requires a fluence of 2.89×10¹² 200 MeV protons/cm², for a dose of ~141 krad(Si). One way to accumulate recoil ion fluence without dosing the parts is to test, say 10 parts to 10% of the desired cumulative fluence a TID of 14.1 krad(Si). This not only reduces TID to each part, but also presents an opportunity to expose part-to-part variation if it is present, at least for common error modes.



10^{10} 200-MeV protons/cm²
 34589 recoil ions/cm²
 0.49 rad(Si)
 1 ion per 2891 μm²
 $3 \leq Z \leq 15$
 All angles, LETs, Energies



10^7 heavy ions/cm²
 1.25 rad(Si) (for Ar ions)
 1 ion per 10 μm²
 Single, Z, angle, LET,
 energy

Figure 5-5: Only one out of 289100 200 MeV protons generates a recoil ion. Thus in the $\sim 60 \times 70 \mu\text{m}^2$ area of the Elpida EDS5108 512-Mbit SDRAM, a) a fluence of 10^{10} 200-MeV protons generates only 3 recoil hits (an upward fluctuation from the mean of 1.45). b) Compare this to the coverage generated by 10^7 ions/cm² typically used in a heavy-ion accelerator test.

Other recommendations for improving results when using proton test data to bound heavy-ion SEE rates include:

- 1) Using LET_{EQ} to interpret results will avoid underestimating SEE susceptibilities, especially for destructive SEE modes.
- 2) Because proton-recoil ions differ significantly from those in the space environment, it is important to understand as much about the parts being tested and the SEE mechanisms as possible, so that one can ascertain whether the differences affect susceptibility
 - a. For novel technologies where SEE mechanisms may not be well understood, use of protons to bound heavy-ion rates is not feasible.
 - b. This is especially important when proton testing is done at the board or box level, where information on part-level response may not be available and each part may

have a different SV depth and so a different fluence vs. LET_{EQ} spectrum. (See Figure 5-6.)

- Because each part tested with protons has a residual upper bound on its failure rate due to modes not detected by the test, the bound on a system built with several parts subjected to proton testing will scale with the number of parts. This will be true whether the parts were tested one at a time or at the system, board or box level.

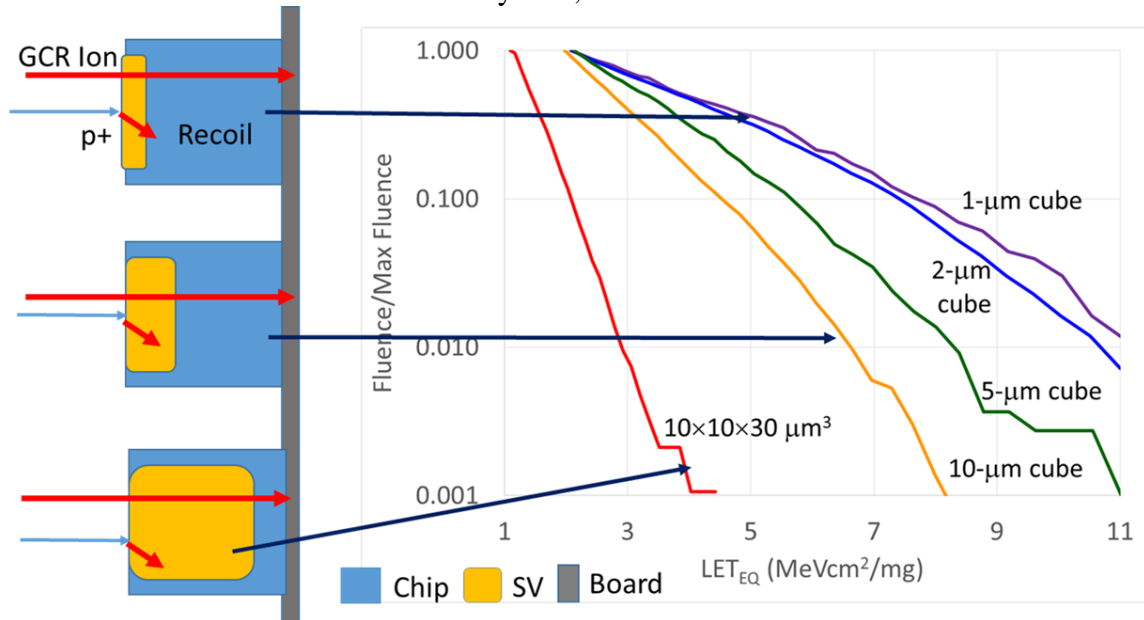


Figure 5-6: When testing with protons at the board level, each device on the board—indeed different SEE modes on the same chip—may have different SV depths. This means the same proton fluence may have very different implications for the SEE susceptibility for each chip on the board.

5.2.3 More on Proxies

Discussion of proxies for SEE Hardness assurance could easily merit a short course of its own, and the previous sections are intended only to introduce the reader to some more quantitative methods for proxies. As mentioned above, it may also be possible to bound SEE response of parts based on the physics of the part and/or SEE mechanism or expert opinion. In addition to protons, it is possible that laser testing [67] and even synchrotron x-rays may be used to bound SEE rates for some parts once sufficiently detailed models are developed [68]. The possibilities depend on how ambitious the goals are for testing and analysis. As with any tool, the utility of the tool depends on how ambitious the goals of the testing or analysis are.

5.3 SEE Testing

If a part performs a critical function, there is no other substitute part known to be SEE hard and information available about a part is insufficient to ensure it will meet its required performance, SEE testing will likely be the only way to verify its performance. Because SEE testing is expensive, it is usually regarded as a last resort. However, performing an SEE provides an opportunity to tailor the test specifically to ensure that it produces the data needed to ascertain the parts suitability. SEE testing has become highly specialized. Each of the general

standards, guidelines and procedures in Table IV details different test methods and conditions optimized to detect susceptibilities to different destructive and nondestructive SEE modes. MIL-STD-750, Method 1080 is a more detailed test method that takes into account the more complicated mechanisms of SEB and SEGR.

Table IV: Governing SEE Test Documents

Type of Testing/ Environment	<u>Available Standards, Guidelines and Methods</u>
Proton SEE Testing	JEDEC JESD234: TEST STANDARD FOR THE MEASUREMENT OF PROTON RADIATION SINGLE EVENT EFFECTS IN ELECTRONIC DEVICES (2013)
	ESCC 25100: SINGLE EVENT EFFECTS TEST METHOD AND GUIDELINES (2014)
SEE Testing for Terrestrial Effects	JEDEC JESD89: Measurement and Reporting of Alpha Particle and Terrestrial Cosmic Ray-Induced Soft Errors in Semiconductor Devices.
Heavy-Ion SEE Testing	ASTM F1192-11: Standard Guide for the Measurement of Single Event Phenomena (SEP) Induced by Heavy Ion Irradiation of Semiconductor Devices.
	JEDEC JESD57: Test Procedures for the Measurement of Single Event Effects in Semiconductor Devices from Heavy Ion Irradiation (under revision)
	ESCC 25100: SINGLE EVENT EFFECTS TEST METHOD AND GUIDELINES (2014)
Heavy-Ion Testing for SEB/SEGR	MIL-STD-750, Method 1080, Single-Event Burnout and Single-Event Gate Rupture.

Outlining best practices for SEE testing is beyond the scope of this short-course segment. Not only do the test guidelines do a much better job of providing advice than would a short section in a short course, SEE testing could merit an entire short course in and of itself. In fact, it has...actually several in recent years.

Table V: Recent Short Course Segments on SEE Testing

Year	Author	Short-Course Segment Title
2014	P. Roche and G. Gasiot	Facilities and Methods for Radiation Testing
2013	H. Quinn	Challenges of Testing Complex Systems
2012	J. Pellish	Single-Event and Total Dose Testing for Advanced Electronics
2008	C. Hafer	Ground-Based Testing and Evaluation of Soft Errors

The reader is advised to consult these short course segments for more information about test methods. Instead, we look first at the goals of SEE testing for various types of SEE testing and rate estimation options.

5.3.1 SEE Testing Goals

The purpose of SEE testing is to bound the risk of using the device under test. As such, ideally, the test should reveal the characteristics of the device’s SEE modes and enough information to estimate SEE rates. Consequences may be summarized by a SET amplitude vs. duration plot like 3-5, a distribution number of bits upset for MBU/MCU or a qualitative note on whether a SEL mode was destructive or whether a SEFI mode required a power cycle for recovery. The data on consequences feeds into assessments of system-level effects of the observed SEE mode.

In some cases, the precise rate is less important than whether the part is susceptible at all and/or whether the consequences of the mode are acceptable. For example, if mitigation of SETs

is to be done via capacitive filtering, how often the SET occurs is less important than the duration/amplitude distribution. If destructive SEE are unacceptable, observation of such a mode precludes use of the part. Or, as for SEB or SEGR, rate estimation may be impractical or impossible.

For SEB and SEGR, the goal is to determine application conditions in which the device is susceptible to the SEE failure mode so that they can be avoided in the mission applications. This type of testing can be accomplished with relatively little beam time. In some cases, a single high-fluence irradiation with a high-Z, high LET ion that observes no threatening SEE modes may be sufficient to conclude risks are negligible for those modes. And if SEE are observed...? One of the risks of a go/no-go test is that sometimes the answer is “no-go”. With measurements of cross section for only one LET, there is no way to compare response for susceptible devices. Finally, although beam time for an SEE test is expensive, the other aspects of the test—particularly salaries for highly trained personnel—account for more than 50% of the total test cost. For this reason, a single-ion go/no-go test is not common. Even for SEB and SEGR, testing is often carried out with several parts, a few ions and over a range of application voltages. For SEE modes where rate estimation is possible, adding even a few additional data points to the σ vs. LET allows at least rudimentary rate estimation.

5.3.2 SEE Rate Estimation Methods and Testing Goals

SEE rates can be estimated in a variety of ways, all of which involve measurement of how SEE cross sections change with respect to proton energy, heavy-ion LET or for detailed Monte Carlo approaches, combinations of ion species, energy, angle, etc. In general, the more detailed the model, the more data—and the more test time, cost and schedule are required for the rate estimation.

The simplest rate estimation technique is the Figure of Merit (FOM) approach [69]. The FOM rate for a given radiation environment ($R(\text{ENV})$) is given by the equation:

$$R(\text{ENV}) = C(\text{ENV}) \times \frac{\sigma_{\text{lim}}}{\text{LET}_{0.25}^2} \quad (4)$$

where $C(\text{ENV})$ is a constant for the environment, σ_{lim} is the limiting cross section for the SEE mode and $\text{LET}_{0.25}$ is the LET at which the cross section reaches a quarter of its value. The FOM approach can be used for proton or heavy-ion SEE rates, and it can yield accuracies within an order of magnitude or so of the on-orbit rate. Nominally, the method requires at least two parameters, σ_{lim} and $\text{LET}_{0.25}$, so cross section should be measured for >4 different LETs, at least two of which are on the saturated portion of the curve. However, the FOM rate will improve with more data.

The most common rate estimation approach is the Integral Rectangular Parallelepiped (IRPP) model of the type used in the CREME96 package. (Note: In Europe, the OMERE code from TRAD performs SEE rate estimations via methods similar to CREME96 [70].) For this approach, the σ vs. LET curve is fit to a Weibull (or sometimes a lognormal) form, and the fit parameters serve as input for the rate calculation. The fit for CREME96 requires at least 4 fit parameters, including onset LET, LET_0 , σ_{sat} and the Weibull width and shape parameters, w and s . However, one can fit to as many as 7 parameters including the SV depth, a funnel parameter, and the rectangular dimensions in the plane of the device, x and y . Thus, the σ vs. LET curve should include cross sections for >6 and preferably >9 LET values. JESD57 [39] discussed the appropriate distribution of LETs.

The RPP model presumes that several conditions apply. It presumes that LET is a valid metric for parameterizing the dependence of SEE cross section, and that ion LET is roughly constant as the ion traverses the device SV. It presumes that the SV can be modeled as a single RPP. It presumes that nuclear interactions are negligible. If any of these conditions do not hold, the RPP model will not yield accurate rates and a more sophisticated rate estimation such as the CRÈME-MC Monte Carlo Rate Estimation package may be a more appropriate choice.

CRÈME-MC is a physics-based model that follows the propagation of ions through one or more simplified sensitive volumes [71]. The sensitive volumes may be nested or disjoint. As long as they reflect the basic features of the part, and the calculation uses the correct physics models, the rate estimation for complicated problems should be improved over CREME96. Thus, the emphasis is on defining SV geometry for a sufficient number of LETs, angles, energies/depth, etc., that the tally of ions depositing sufficient charge to cause an SEE is reasonably accurate.

The term “reasonably accurate” is intentionally vague. How accurately the SV must be modeled depends on how accurate a rate estimation is required. Mapping out the SV using ion beams over 4π steradians is impractical. One approach is to use the Computer Aided Design (CAD) capabilities in MRED, a more sophisticated Monte Carlo program [72]. If one has such a model, it becomes a matter of simulating the response of the part to ions with a range of Z, energy and angle and then carrying out heavy-ion testing to validate the model by looking for diagnostic responses shown in the simulations—e.g., angular dependence, σ vs. LET or Energy/Z, and so on. This method is useful for validating SEE hardening approaches in deep submicron CMOS. Unfortunately, one is unlikely to have such a detailed model for a COTS part, as such a model would be considered highly proprietary by the vendor and most COTS vendors are unlikely to assist in obtaining one. Reverse engineering the structures in the part may help, but assembling a sufficiently detailed model would be a Herculean task for any but the simplest structures.




Figure 5-7 CAD model of a single latch with dielectric materials rendered transparent to show metallization. (Courtesy of K. Warren. Adapted from [73].)

The use of MRED and CAD models differs from the conventional approach for SEE rate estimation in that SEE testing plays no role in the rate estimation or the parameterization of the device model. The model is determined from the CAD drawing. Based on simulated irradiation of this model under various conditions, the analyst keeps track of notable distinguishing behaviors over angle, energy, Z, etc. Then, during the irradiation of the parts, the analyst looks whether the behaviors predicted by the model are reproduced in the data. If the modeling and experimental results being compared are quantitative, a goodness of fit criterion can be applied to the data to ensure the model accurately reflects the device geometry. If there are adjustable parameters in the model, multiple series of simulations can be carried out across the ranges of the parameters, and if a model gives a significantly better fit, it can be selected. Otherwise, the worst-case model for the environment that is consistent with the experimental data can be selected or results can be averaged across the models.

Another approach that has proven useful, and which uses the much simpler CRÈME-MC package, is the nested volume approach [74]. In this approach, the increase in the sensitive volume with LET is approximated as a series of RPP nested volumes. The rectangular dimensions of the volume and the efficiency of charge collection in each incremental shell are determined from the σ vs. LET curve (which roughly follows the usual Weibull form in a stair-step fashion (see Figure 5-8)). These nested volumes are then exposed to the radiation environment of interest, and if the sum of the charge deposited in each volume weighted by the efficiency for that volume exceeds the critical charge, an SEE is tallied. This approach allows one to consider situations where LET varies along the track, nuclear interactions are important, etc. while determining the sensitive volume, efficiencies, etc. purely from quantities measurable during a heavy-ion SEE test. In this case, the improvement in SEE rate obtained by adding measurements of σ at more LET values depends on the shape of the σ vs. LET curve.

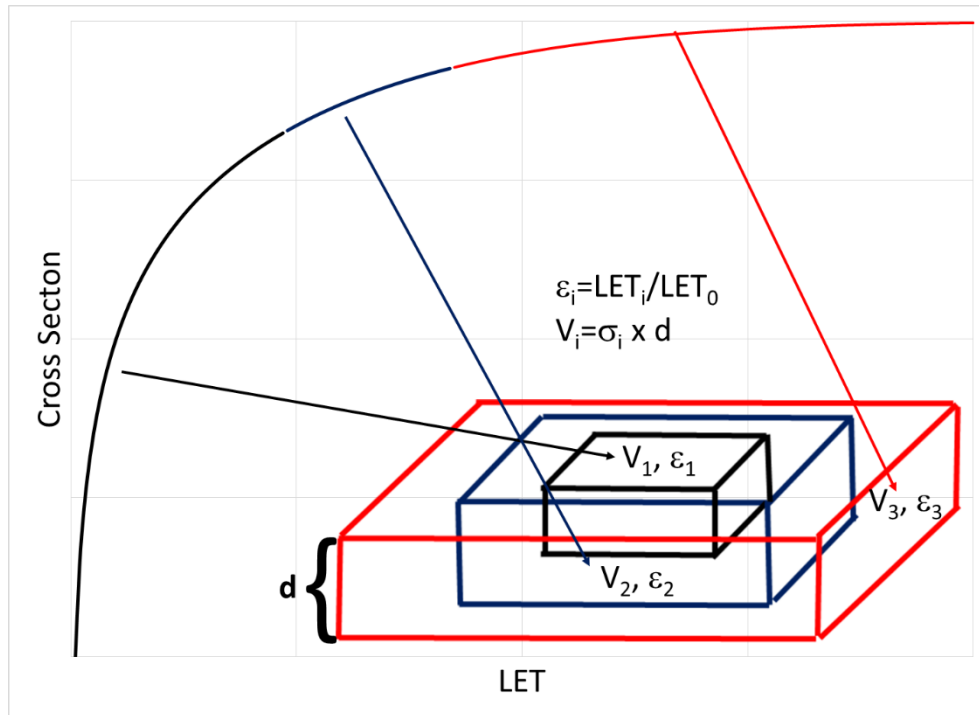


Figure 5-8: For the nested volume rate estimation, dimensions of each volume the the charge collection efficiency in that volume are determined from the σ vs. LET curve.

5.3.3 Proton SEE Testing

Although we discussed use of proton SEE test data to bound heavy-ion SEE susceptibilities in section 5.2.2, we briefly look at the goals of proton testing for much less controversial task of bounding risks of proton-induced SEE. In general, whether one seeks to bound heavy-ion or proton SEE rates, proton testing is much less complicated than heavy-ion testing. This is because high-energy protons cause SEE via recoil ions, and proton beam energy, angle and other characteristics only weakly influence the recoil ion. The distribution of the recoil ions in terms of ion energy, angle, Z and LET varies relatively little whether the proton energy is 100 MeV or 200 MeV. As with heavy-ion testing, the goals of proton SEE testing are mainly two-fold—to reveal susceptibilities and consequences of SEE that can be caused by protons and to estimate on-orbit rates due to protons in the mission environment. There are several ways in which this can be done. The FOM method has also been applied to proton-induced SEE rates, and environmental constants C(ENV) have been determined for trapped proton belts in Earth orbit.[69] A more accurate approach is to measure the increase in SEE cross section with proton energy. Then, based on the LET spectra and other characteristics of proton-recoil ions, a rate can be estimated. CREME96 includes the PUP module, which has options for estimating rates assuming the cross section follows the Bendel one-parameter[75], Bendel two-parameter[76] or Weibull forms. Since the fits to SEE cross section vs. proton energy may include several parameters, a meaningful fit requires measuring the cross section at several proton energies—3-4 for the Bendel methods and 6 for the Weibull fit.

5.3.4 SEE Risk Evaluation: From Data to Rates to Risk

The previous two sections made clear that the data during SEE testing—or archival data to be used—must be tailored to the planned rate estimation methods. The rate estimation method may be as simple as a FOM algebraic calculation or as complicated as a full blown MRED Monte Carlo simulation. Ultimately the data must be sufficient to specify the “correct” model.

Unfortunately, errors in the SEE data may prevent it from fully specifying a single model: Poisson errors are inherent to SEE counts, and other errors (e.g., fluence measurement [77]) may also contribute.

When event counts are small, as is often the case when destructive or disruptive SEE modes prevent accumulation of statistics, Poisson errors can dominate. For the CREME96 model, Reference [78] proposed that in this situation, the data could be fit using a variation of Maximum Likelihood methods called a Generalized Linear Model (GLM). Assume x_{pi} and x_{oi} are the predicted and observed counts for the i^{th} LET, where the x_{pi} are predicted by the model (e.g. Weibull form of the cross section with a specified onset LET, LET_0 , limiting cross section σ_{lim} and Weibull parameters w and s) and x_{oi} were observed during the experiment. Then we construct the likelihood $L(\{x_{oi}\}, \{x_{pi}\})$ that the model generated the observed data:

$$L(\{x_{pi}\}, \{x_{oi}\}) = \prod_{i=1}^{\text{ALLET}} P(x_{pi}, x_{oi}), \quad (5)$$

where $P(x_{pi}, x_{oi})$ is just the Poisson probability that an expected event count of x_{pi} fluctuated to x_{oi} . One calculates the likelihood for all parametric values in the model, and the likelihood that is maximum is the most probable model in a Maximum Likelihood sense. Moreover, because likelihoods tend to be normally distributed as one moves away from the maximum likelihood values of the parameters, any model will agree with the observed data within a given confidence C if it satisfies the following condition:

$$\ln(L(C)/L_{\text{MAX}}) \geq -0.5 * \text{INV}\chi^2(1-C, \# \text{parameters}), \quad (6)$$

where $\text{INV}\chi^2(1-C, \# \text{parameters})$ is the inverse χ^2 distribution for the number of parameters in the fit (4 for the typical Weibull cross section fit). We can then select model that yields the highest SEE rate of all those satisfying the above inequality and this rate is the worst-case rate for confidence level C . Reference [78] noted that there is nothing about this method that precludes use of models other than CRÈME-MC. As long as the model predicts expected event counts for each observed count. Moreover, there is nothing that requires use of LET for parameterization of the cross section—one could even use a multi-dimensional variable space, e.g. (Z , energy, angle). Thus, this method would also be well suited to comparing Monte Carlo models. However, care should be taken that the models being compared have similar complexity to avoid overfitting the data.

With the selection of the model (or models), one can then estimate the rate for each SEE mode observed in testing. The consequences of the mode can be traced from the part level to the system and the consequences to mission requirements assessed. With the rate and consequences, we can specify risk at the part, circuit, system and mission levels, and based on the risk, we will decide whether mitigation is required.

6 Mitigation of SEE Risk

Since the goal of any risk mitigation is to reduce the risk, and risk is the product of failure probability and failure consequences, it is not surprising that risk mitigations usually focus on reducing probability of failure or on decreasing the consequences of failures that occur or, in some lucky cases both. We next consider the options for mitigating SEE, starting first with strategies for reducing SEE probabilities and then for reducing SEE consequences.

6.1 Reducing SEE Probabilities

As the very concept of reducing SEE probabilities is predicated on knowing something about what those probabilities are, it is safe to assume that test data are available for the problematic parts. Probability reduction is a threat avoidance strategy, and most strategies have to do with part substitution or avoiding conditions where the threat is likely to occur. Part substitution has probably been the preferred strategy for avoiding SEE, but it relies on a part existing that is SEE hard and can meet requirements in the application. Since there was probably a reason why the designers turned to a non-SEE hardened part in the first place, the question is whether anything has changed that would make reduced performance due to the substitution more acceptable.

If part substitution is not acceptable, it still may be possible to avoid the error/failure mode by avoiding applications where the mode is likely to occur. This is essentially the strategy behind the safe-operating voltages to ensure power MOSFETs are not susceptible to SEB and SEGR. Similarly, we know that lowering the supply voltage and application temperature are possible mitigations for reducing SEL rates. SETs are usually less likely to cause error if the system operates at lower frequency. However, unless we have archival or test data, we cannot know whether the deratings or guidelines are adequate to reduce SEE rates to acceptable levels. For example, latchup in some parts occurs so rapidly that current limitation, event detections and power cycling are ineffective for avoiding destructive failure.

The mitigations discussed above rely on nothing more than basic knowledge of the SEE mechanisms—whenever we can control the factors that contribute to the susceptibility, we adjust them to minimize the susceptibility. Other strategies may be specific to a particular part or part type. For example, the version of the LM139 available in the 1990s was found to have much lower susceptibility to SETs when the voltage difference on the two pins was higher.[79] Another example is that some DRAMs exhibit much lower SEFI rates if the mode registers are refreshed regularly [80].

Thus, the options for reducing SEE probability are limited to part substitution, tailoring applications to minimize variability and opportunistic strategies based on test results specific to the part in question. The reader may have noticed that one mitigation is specifically missing from this list: redundancy. This is because redundancy does not reduce SEE rates, but rather SEE consequences.

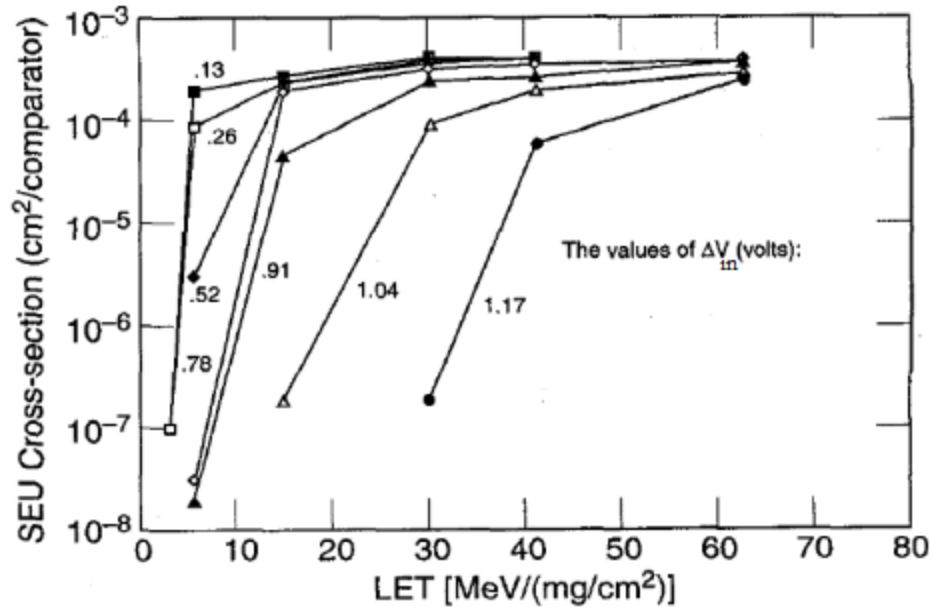


Figure 6-1: The SET threshold of the National Semiconductor LM139 comparator exhibits a strong dependence on the input voltage difference. (Supply voltage was 13 volts, adapted from [79].)

6.2 Reducing SEE Consequences

Most SEE risk reduction strategies seek to reduce the impact of an SEE if it occurs. This can be done in several ways, some of which are specific to the part type and some of which are general strategies. We first look at some of the less general strategies and then look at the most general strategy for mitigating SEE—use of redundant elements to preserve information or functionality in the event of an SEE induced error or failure. In conjunction with the discussion of redundancy, we also discuss some strategies that are employed to supplement other techniques (e.g., memory scrubbing, bit interleaving).

6.2.1 Specific Strategies

If we are mitigating the consequences of an SEE, we are resigned to the fact that it is going to occur. Where we hope to make a difference is by diminishing the consequences of the SEE. Two ways of doing this are directly limiting the effects of the SEE at the part, circuit or system level, or by facilitating recovery from the SEE consequences.

An example of limiting SEE consequences is the use of a SEL protection circuit. This circuit senses the rise in current as the SEL begins and rapidly cuts off power to the device, stopping (hopefully) the SEL before it can damage the circuit. There are several challenges to implementing such a circuit. The first is ensuring the circuit is effective—not only at sensing the overcurrent quickly enough to avoid run-away current, but also in proving that it is rapid enough to avoid latent damage to the parts. (See [38] for an account of successful implementation of this strategy.) Another challenge arises from the need to keep the trigger sufficiently sensitive to protect the parts while avoiding spurious triggers that can cause outages.

Another circuit-level approach mitigates SETs by capacitive filtering, preventing them from affecting downstream electronics. Again, the challenge is to ensure that the filtering is effective at stopping the transients while not slowing down the circuit unacceptably. As a general rule,

transients are less likely to propagate through the system if it is not required to perform at its top speed. SETs are more likely to be captured by downstream devices if they are operating at high frequency. Also, some analog-to-digital or digital-to-analog converters exhibit more transients corresponding to low-order bits, and the rate can be reduced by masking these. However, in some cases it may not be possible to prevent the SEE mode affecting the system, or in some cases, the cost to performance may be too high.

If one cannot prevent the SEE from having consequences, the best approach may be to facilitate recovery from those consequences. As can be seen in Figure 6-2, which illustrates the recovery process, there are several times that can be minimized to speed recovery. Event detection can be facilitated by watchdog timers; error checking can be implemented on calculation results; telemetry monitored for anomalous readings, and so on. Once the error is detected, the recovery can begin. The mitigation for the recovery process consists of ensuring the resources are in place and ready to do what is needed to restore normal operations. This may involve ensuring that the SEE susceptibilities are well characterized so that procedures are in place for an efficient recovery. If the failure mode is very disruptive, it may involve ensuring resources can be dedicated to restoring the system configuration. For the final time increment—from recovery to restoration of full operations, the main issue is ensuring that verification of system status, health and safety are carried out efficiently—which again is facilitated by understanding the error/failure modes and systems for the hardware.

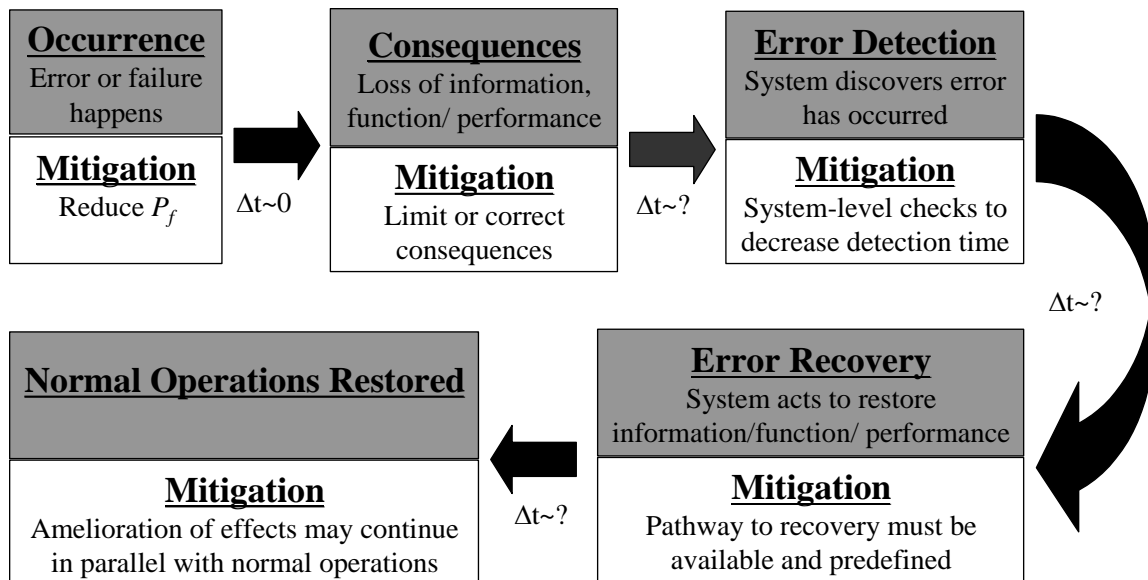


Figure 6-2: Steps in the process of recovery from an anomaly.

6.2.2 Redundancy and Supporting Strategies

Because the direct consequences of a SEE are confined to a single die, redundancy—in one form or another—can be used to mitigate any SEE mode in a system. Unfortunately, using redundancy to harden against SEE is a strategy that can fail if basic assumptions of the method are violated:

- 1) The error/failures must be independent and occur at a constant rate.
- 2) The probability of an error in the interval of concern must be sufficiently low

Although these assumptions seem simple in practice, there are many ways in which small changes can seriously reduce the efficacy of the mitigation. As an example, a DICE latch, which relies on redundant, interlocking nodes to provide SEU hardening, may work fine in 90 nm CMOS, but fail when implemented at 45 nm CMOS because the increased proximity of the redundant nodes in the latter process makes simultaneous strikes more probable (violating assumption 1). Similarly, upsets may be correctable with EDAC as long as the memory is scrubbed more than once per day, but may accumulate and cause system-level errors for less frequent scrubbing—the error probability in the scrubbing interval is too high for the redundancy to help at the system level (violating assumption 2).

Redundancy can be used to mitigate any SEE that occurs, provided that the above criteria are met and the cost to the system in terms of size, weight, power, performance and complexity are acceptable. However, one must carefully consider the consequences of the SEE one is trying to mitigate in designing the mitigation strategy. For instance, a SEFI can result in loss of both functionality and of data integrity for large blocks of data. EDAC or triplicate voting may be sufficient to restore the data. However, if restoring functionality of the affected device requires a power cycle, this could result in the loss of even more data. Restoring both functionality and the data requires a much more complicated intervention and system architecture than merely maintaining functionality. Similarly, redundant devices can provide protection against destructive failure, but if they are also used as part of a triplicate voting system, when one of the parts fails, all error correction capability is lost, despite the fact that there is still one redundant part. Not only that, but the triplicate system has roughly three times the probability that a failure will occur, since all the parts must be biased to provide error correction capability. The morale of these examples is three-fold:

- 1) The goal of the mitigation needs to be defined carefully, and one needs to pay in the coin of the realm—extra samples to mitigate samples corrupted by an SET, extra bits (in EDAC or voted in triplicate) to correct SEU, extra devices stored “cold” (unbiased) to replace failed devices—or—stored hot to maintain functionality when one device suffers a SEFI).
- 2) The likely performance of the devices must be sufficiently well understood, that it can be verified that they will meet the assumptions required for effective redundancy above.
- 3) For redundant systems, it is important that the environments for which device performance is analyzed should include the worst-case environment. This is because in a system that uses n parts where m parts are required ($n > m$, termed n -for- m redundant), the system failure rate scales nonlinearly with the part-level rate. Thus, if part-level rate rises by a factor of 10 in a moderate solar particle event, the system-level error rate can rise by 100x or more.

Below, we give some examples of mitigation relying on redundancy that illustrate these points.

The thing about SETs is that they are transient—the output of the device eventually returns to normal. Thus, sampling the output of the device three times on a timescale significantly longer than the transients the device exhibits and combining the three results (note that the combination could be by voting the results for digital outputs, taking the median or the average, depending on the characteristics of the errors). Alternatively, one can compare the outputs of three separate devices and vote them the same way. In this case, the redundancy relies on the fact that simultaneous errors in two different devices will be rare, so the system sampling rate is not reduced. The triple-device voting scheme in Figure 6-2 also works for SEU, although the voting logic must correct the erroneous device.




Figure 6-2: SETs can be mitigated either by resampling the output of the device and voting or voting multiple devices. Actually, depending on the nature of the transient, the noise on the signal, it may be most advantageous to take the mean, median or mode (winner of the majority vote) of the signal.

For digital data stored as data words or blocks, error detection and correction algorithms (EDAC) [81], such as the Hamming code depicted in Figure 6-3 [82], may be a more efficient error correction solution. Such algorithms allow a finite number of bits or groups of bits (nibbles, bytes, symbols...) to be corrected. For instance, the Hamming code with 4 data bits and 3 error-correction or parity bits can correct any single-bit error. In general, a Hamming code can correct a single bit in a word of $2^n - 1$ bits by adding n additional parity bits. Adding one more bit allows detection but not correction of any two-bit errors. The problem with Hamming codes is that they can generally only correct a single bit in error, making MBU, block errors and SEFI uncorrectable. More sophisticated EDAC schemes, such as the widely used Reed-Solomon code [83], detect and correct errors in blocks of code (e.g. nibbles, bytes, symbols...) and can correct any number of errors in these blocks.

However, even for such powerful algorithms, block errors and SEFIs can overwhelm the code. For this reason, EDAC is usually paired with additional measures to prevent errors from exceeding the capacity of the EDAC. Scrubbing the memory to detect and correct errors before they accumulate over time is effective provided the scrubbing rate is significantly faster than the error rate. However, because the error rate can increase by orders of magnitude for a large SPE, it is important that the scrubbing rate is sufficient to handle such conditions if the system is required to operate through them. Another measure often combined with EDAC is interleaving of bits in each word across multiple die. Because the effects of a SEE are confined to a single die, as long as the maximum number of bits in a word, block or symbol on a given die are fewer than the correction capability of the EDAC, any error mode will be recoverable in principle. In practice, if the memory is volatile and if a power cycle of the system is required to recover functionality, all data will be lost. It is often impractical to implement the ability to cycle power to individual memory die, memory modules or even single boards. As such, if a part requires such a power cycle to recover from a SEFI mode, this can often drive the error rate and/or the design of the system. As such, it is important to avoid such susceptibilities. In some SDRAMs, for example, frequent rewriting the mode registers of the device is sufficient to reduce the rate for SEFI requiring a power cycle to undetectable levels [80]. Alternatively, large memory blocks can be implemented with nonvolatile memory if the lower write and access speeds are tolerable. This allows power to be cycled—or even turned off when not in use—without loss of data.

Although it may seem complicated to implement scrubbing and interleaving in the memory, the payoff is that one can achieve detection and correction of errors equal to or better than triplicate voting with a fraction of the overhead. A Reed-Solomon algorithm capable of correcting any errors occurring in two 4-bit nibbles can be implemented with 50% overhead,

compared to the 200% overhead of a triplicate voting system. It is thus possible to achieve a high level of data integrity, as long as the cost in system size, weight, power, complexity and performance are acceptable. However, as indicated by the discussion of power cycling to recover functionality, data integrity is only one of the issues that mitigation must address.

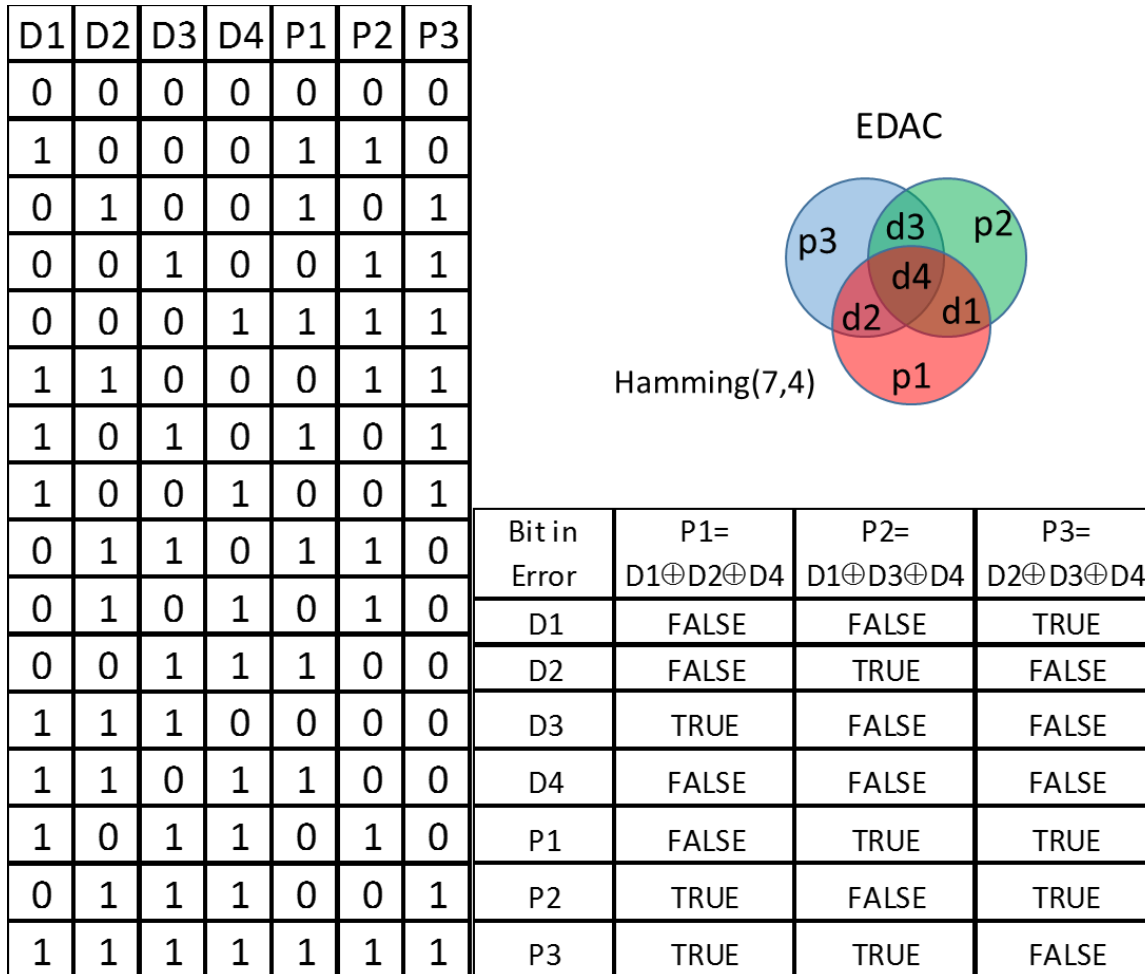


Figure 6-3: A Hamming code (7,4) is among the simplest error corrections codes, but it illustrates the basic ideas underlying such schemes. Three parity bits are determined by combinations of 3 data bits such that if any one bit (parity or data) upsets, a unique parity-bit combination will be in error. Unfortunately, for two upsets, the configuration will still be valid. However, a fourth parity bit, determined from the 3 other parity and 4 data bits, allows detection but not correction of two errors. A single-error correcting (SEC) Hamming code can always become a single-error correcting, double-error detecting (SECDED) code by adding such a single parity bit.

As discussed in Section 3, SEE may have several different consequences other than corruption of data. Destructive SEE can cause catastrophic failures resulting in loss of functionality, system availability or even the mission. SEFI modes can cause loss of system availability and can be result in prolonged and disruptive recovery processes that can themselves affect system availability, data integrity and other performance metrics. Redundancy can be used to meet system reliability and availability requirements, but there are trade-offs between these goals. Figure 6-4 illustrates the reasons for this. If the purpose of the redundant device is to replace a primary that fails due to DSEE, usually the redundant device is unbiased (cold) and so not susceptible to the failure mode. No more than one device is biased at any given time. On

the other hand, for a redundant device to mitigate loss of availability, it must be in the same state and performing the same function as the primary device. As such, it must be powered and operating, and it will be susceptible to the same SEE modes—particularly destructive SEE—as the primary device. Thus, if the primary device were susceptible to DSEE at a rate of λ , the system, including primary and redundant devices, will have an expected rate of 2λ . The effects of the increased system rate depend on the redundancy of the system and on the DSEE rate of the part. For instance, a 2:1 system will always be more reliable than the single-string system if the expected number of failures is less than 1 per mission, while a 4:3 system (e.g. a triplicate voting architecture with a 4th string to replace a failed primary) becomes less reliable than an individual string (which, admittedly, cannot correct errors) if the expected number of DSEE in the mission is as low as 0.184. As such, in a 4:3 system, keeping the fourth device as a cold spare would yield a more reliable system, while a hot spare provides a system with greater availability.

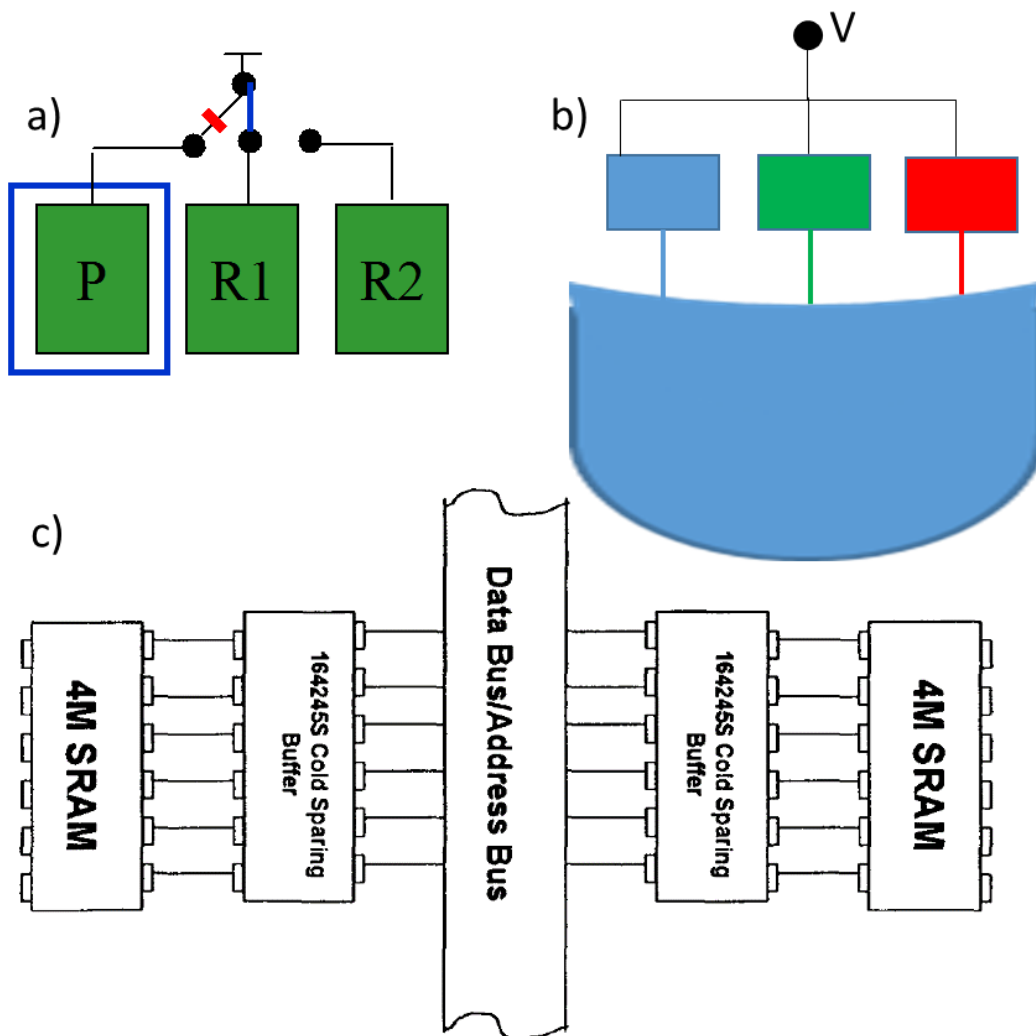


Figure 6-4: Reliability and availability are competing metrics with respect to redundancy. a) If a triple redundant scheme is configured with cold spares, only one device is susceptible to failure (with rate λ) at a time, and 3 independent failures are required for system failure. However, if the primary fails, service becomes unavailable until the primary can be isolated and the redundant unit mapped in. b) For a triplicate system configured to maximize availability (and/or data integrity), all three devices are biased and so vulnerable to failure (rate= 3λ). c) A cold spare system requires isolation of the failed device in this case using a buffer [84].

6.3 Putting It All Together

As discussed above, although a single mitigation technique may not provide the SEE hardness desired, combining techniques can produce a very robust system, provided the conditions for redundancy to be effective are met and the cost, performance, power consumption, size and complexity of the system are acceptable. Figure 6-5 illustrates this process, including the multiple ways in which redundancy is used.

First, memory hardening is implemented. The primary mitigation for the memory is EDAC. However, because EDAC can only correct a certain number of bits in error, scrubbing must be implemented, and bits in data words may must be interleaved across die so that no single SEE (or no two... depending on desired hardness) can corrupt more bits than can be corrected by the EDAC. If the parts are susceptible to destructive SEE, additional memories may be needed so the system meets its end-of-life requirements. These extra memories may be cold spares or hot if the failure rate allows. The scrub rate for the memory must be fast enough that the probability of errors accumulating is negligible. Ultimately, the success or failure of these mitigations depends on the underlying performance of the memory chips. If their error or failure rates are too high, the requirements for EDAC, system architecture, scrubbing and redundancy will not be feasible.

Mitigating options for processors and similar parts are more limited. Data are used actively and change too rapidly for EDAC to be an option, so triplicate voting is often the only mitigation feasible. Each processor performs identical processes on data fed synchronously from memory, and then the output is voted bit by bit. Depending on the susceptibilities of the processors, the redundant processor (PR in Figure 6-5) may be configured as a cold spare to mitigate a processor failure (due to DSEE or any other cause), or it can be “hot”—biased, processing data and either voting in a quad configuration or “silent” until another processors fails or becomes unavailable.

The example in Figure 6-5 illustrates the point made in Section 6.2.2—mitigation costs must be paid in the coin of the realm. For the memories, the coin is data, and we are paying in redundant bits for EDAC (or for triplicate voting of memories; we would still be providing extra bits...just many more of them). For the processors, SEE can manifest in multiple ways:

- 1) Data errors may occur while executing programming, giving rise to errors in the output.
- 2) Programming of the part may become corrupted, producing streams of erroneous data
- 3) The part may simply cease to function, but recover when reprogrammed.
- 4) The part may cease to function, but recover when power cycled and reprogrammed.
- 5) The part may fail catastrophically.

Triplicate voting mitigates the data errors. If one processor stops functioning, the remaining two continue to provide an accurate data stream. However, the system has lost error correction capability. The redundant processor preserves that capability—either keeping it available if the spare is hot or preserving system reliability against destructive failures for a cold processor.

The above simplified example has deliberately omitted some of the most challenging aspects of implementing such a system. For example, ensuring timing is coordinated for all the processors is challenging. Even if a power cycle were not required to restore functionality to a processor that had suffered a SEFI, one is likely required to ensure the processors come into synchronization.

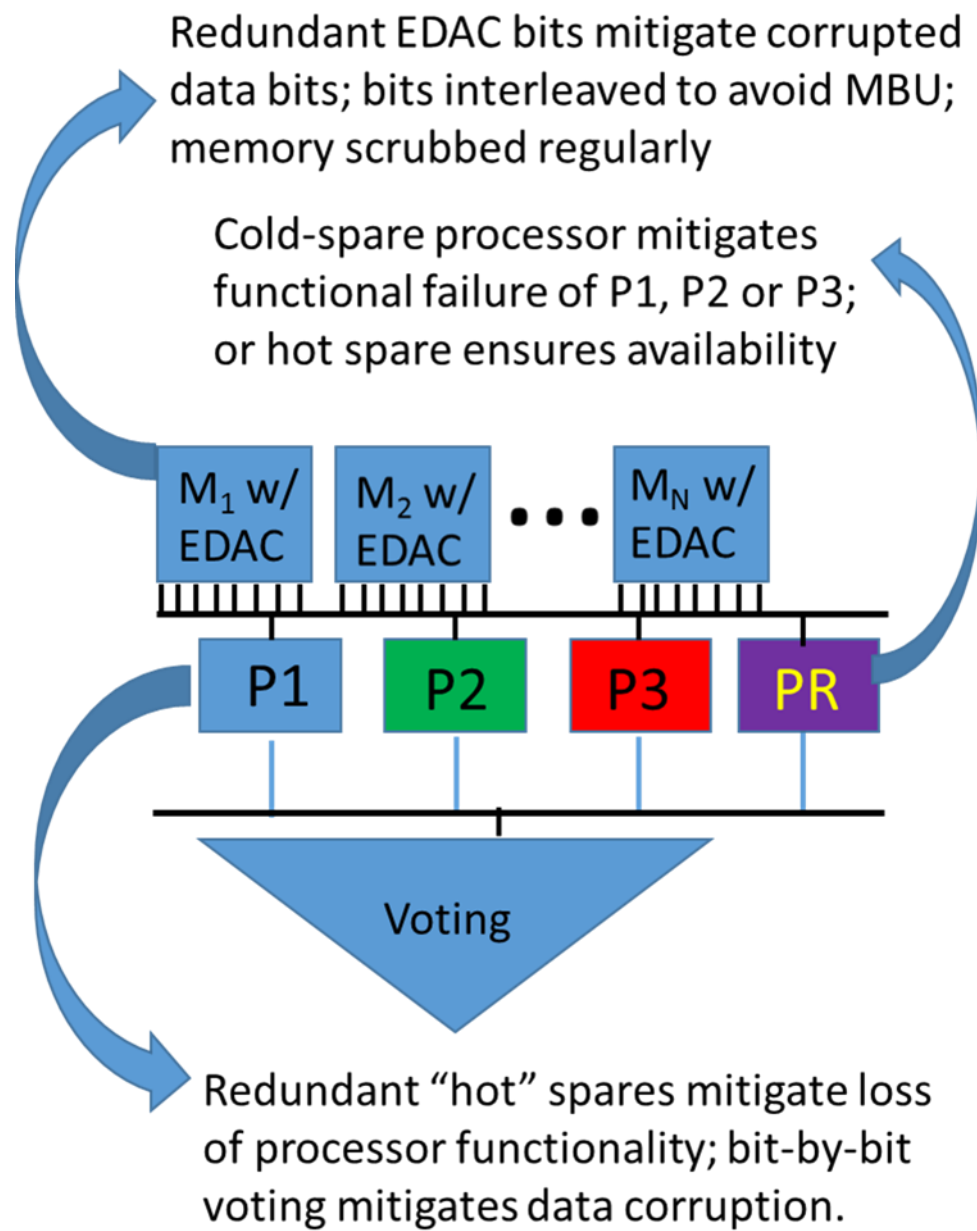


Figure 6-5: For a radiation tolerant system, multiple SEE mitigation techniques are combined. EDAC, scrubbing, bit interleaving and other mitigations can be applied in the memory elements. Processors are usually voted in triplicate, and redundant processors and memories (hot or cold) can be used to ensure availability or reliability.

7 SEE Hardness Assurance for SmallSats

The thing that sets SmallSats apart is not their size (Sputnik was small, after all), but the capabilities that can now be fit into such a small package and the speed with which they can be developed. Indeed, the fact that many SmallSats are secondary payloads makes rapid development essential, because if they miss their ride or if they don't fit in their allotted space, they stay home. This elevates the importance of schedule and of packing performance into a small size. Moreover, not having the cost of a launch is not in the equation changes the scale of the budget—a \$50000 SEE test assumes a greater significance when the project does not face a \$10 million launch cost. It is not just that SmallSats have lowered costs for spaceflight, but rather they have made missions possible that were previously inconceivable and allowed satellites to be built by designers for whom it would have previously been unthinkable. A conventional satellite that flew only one instrument or that had a mission length less than two years would probably not be viable. A conventional satellite would be beyond the aspirations of a university.

However, to make these new missions possible, SmallSats have to play by different rules. They are bound to be more heavily reliant on COTS parts and systems. Their lower budgets mean smaller staffs, so a full-time radiation specialist is not feasible. Of necessity, a lower budget will also mean less SEE testing. On the other hand, the lower costs of the mission are conducive to taking more risks, which could in theory lower SEE HA costs. In this section, we look at why it has proved difficult to realize such savings. We begin with a look at why hardening SmallSats is desirable. We next look at the how SmallSat reliance on COTS parts limits options for evaluating and mitigating SEE risk, how the cost structure of SEE testing make it difficult to trade risk tolerance for cost savings and then discussing how the concept of risk applies to SmallSats.

7.1 Why Harden SmallSats

To date, most SmallSats have been sponsored by universities, and their goals have been as much about pedagogy as about science or exploration. The short development time and mission duration have been well suited to student's timelines for graduation, and if the mission failed, the consequences to the pedagogical goals were minimal, and to science, not much worse. However, just as SmallSats have made possible missions that could not be contemplated previously, so too, radiation tolerant SmallSats could open up further opportunities, allowing SmallSat missions not just in more severe radiation environments, but also longer SmallSat missions that could produce time-series data of significant value.

Swarms of SmallSats could deploy instruments in a range of different orbits, allowing measurements of data (e.g., planetary magnetic fields) over much of a planet and potentially producing valuable time series of data. Unfortunately, the value of a time series of data often increases nonlinearly with its duration, and splicing data series over multiple missions and multiple instruments is a tedious and error-prone task. Thus, SmallSats capable of missions of 5 years or longer could be far more valuable for science than those capable only of shorter missions.

There could also be significant value in extending SmallSat missions into harsher space environments. Unfortunately, not only would this increase SEE rates for these mission, it also takes much longer for the orbit of a satellite at higher altitude to decay than it does for one at low

altitude. Current SmallSat protocols call for cubesats to make provisions to decommission or deorbit within a finite time of the end of the mission. However, if a SmallSat were to fail prematurely due to an SEE, it might not be able to carry out the deorbit protocol, and today's SmallSat becomes tomorrow's orbital debris. Thus it is critical that at end of life the SmallSat retains sufficient control to end its mission according to plan.

Finally, because SmallSats have higher risk tolerance, they can fly new technologies that promise significant performance, but which have no flight heritage and would be difficult to test. However, the SmallSat must be reliable enough that the new technology can be assessed over a long enough period of time—and preferably in a severe enough environment—to serve as meaningful heritage for future missions.

7.2 Challenges for SmallSat SEE Hardness Assurance

SmallSats pose special challenges for SEE HA by their very nature. The schedule and cost constraints favor the use of COTS. The size and power constraints limit the mitigation schemes that are acceptable. Finally, the cost and schedule constraints limit the amount of SEE testing that can be done, and the cost breakdown of SEE testing limits the extent to which risk tolerance can be leveraged into saving on SEE testing. We address each of these challenges in turn.

7.2.1 Implications of Using COTS Parts and Systems

The considerations discussed above make it clear that even coming out of the gate, before risk tolerance or cost of the system enters the equation, SmallSats will be predisposed toward use of COTS. The ability to purchase parts off the shelf means that the project will not have to worry about long lead times typical for radiation-hardened components, nor about schedule if those delivery dates slip. The high integration and capabilities of commercial parts are well suited to packing a lot of performance into a small package. Finally, the significantly lower costs of COTS parts are more compatible with budgets of most SmallSat builders.

Unfortunately, the high-performance hardware inside SmallSats must perform in a harsh radiation environment for which it was not designed, and missions are moving into harsher environments all the time. The surest way to ensure reliability would be to test every unfamiliar part. However, this would blow a giant hole in the budget of even a National Asset class mission, and is almost certainly out of the question for a low-budget SmallSat. It is instructive to step through the conventional SEE HA as it would apply to a typical SmallSat, recapitulating Figure 5-2 as a guide.

- 1) The first order of business would be to scrub the parts list for parts with unknown radiation performance or that are likely to pose SEE risks due to their technology. In the conventional SEE HA, each of these parts would be a candidate for testing if the SEE concerns could not be alleviated by other data.
- 2) The next task would be to search the literature to see if anyone else had tested the part. Unfortunately, SmallSat builders who are most likely to use COTS parts do relatively little radiation testing, and there are often so many commercial options that the probability of finding test data on most parts is small. Conventional satellite builders typically do more testing, but use COTS in small numbers, so again, the most likely outcome is that the search will come up empty. Moreover, most commercial parts have a product life cycle less than a couple of years, so the shelf life of any data available is likely short. This step is unlikely to provide significant reductions to the list.

- 3) Having failed to find data for the selected part, the next resort would be to look for data on “similar” parts. The same factors cited in step 2 also suggest that similarity data will be hard to come by. Moreover, it can be very difficult to obtain specific process data for a part from a commercial vendor. Indeed, the vendor may not even know which of several foundries a part was manufactured in. Although automotive grade parts are more tightly regulated, they may be reluctant to spend time with a project that is buying less than 10 parts when their typical customers order 10s of millions.
- 4) If, as is likely, no similarity data are available, the options are to try to bound risk based on part technology, physics of failure or expert opinion. As mentioned above, obtaining part technology data from COTS vendors may be difficult. Not only will the orders by a typical SmallSat project be puny compared to those of their typical customers, but the company may consider process information and even the foundry in which the parts are made to be proprietary. Even if this information can be obtained, the lack of previous SEE testing of the technology may preclude knowledge of the physics of failure. Finally, even relying on expert opinion a fraught proposition. As mentioned above, a small project is unlikely to be able to retain experts, and experts may be reluctant to opine on a technology for which little data are available. (Indeed, the opinion of an expert who was willing to speculate under such circumstances might be of questionable value.)

Thus, it is unlikely that the list of parts of concern for a SmallSat will be winnowed down significantly by a conventional SEE analysis. One approach in the face of such uncertainty would be to implement a very tolerant design, in which devices are assumed to be susceptible *a priori*. However, the size and power constraints may pose limits as to what mitigations are acceptable.

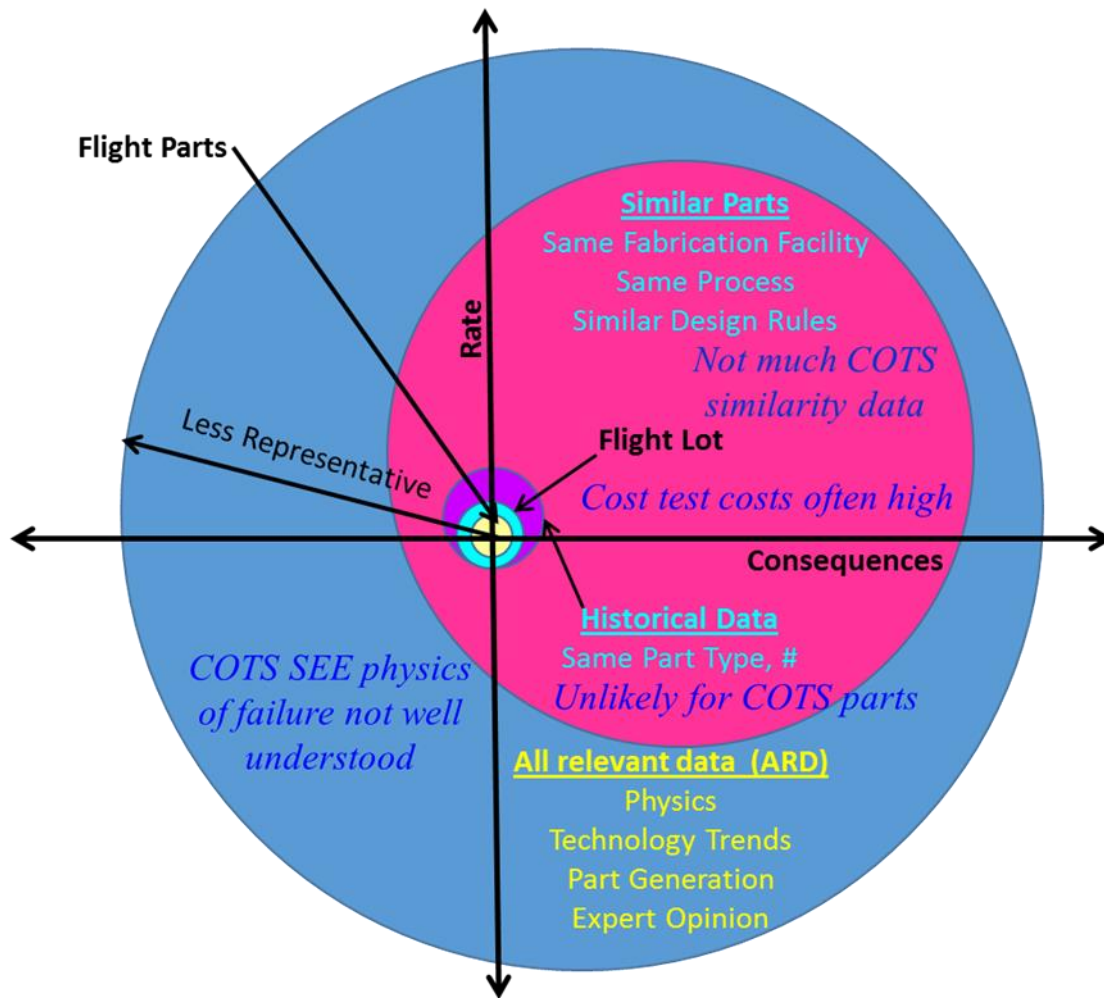


Figure 7-1: Bounding SEE risk for COTS parts is difficult because the lack of radiation test data, profusion of parts and vendors and short product life cycle limit the amount of data available for understanding SEE risks for such parts.

7.2.2 Limitations of Tolerant Design for SmallSats

Although mitigation strategies for SEE are well understood, not all mitigations are equally easy to implement, and the size and power limitations of SmallSats further complicate the process for some strategies. Even with large satellite projects, tolerant design in the absence of device-specific data means, for example:

- 1) Derating power MOSFET VDS to 20% of its rated value if one does not have data on the part, as SEB failures have been seen down to 22% of the rated VDS for commercial parts [48];
- 2) Implementing capacitive filtering sufficient to filter the >1 ms transients seen in the LTC6268 data in Figure 3-5;
- 3) And so on, as without data, designers must implement mitigation based on worst-case parts observed in past testing.

While such defensive design may result in unacceptable performance penalties, the only way to show that they are not needed is via test or proxy data (e.g., similarity). For SmallSats, some

measures are relatively easy to implement—e.g., watchdog timers, error, checking—while others are much more difficult. Cold sparing may strain size limitations, while triplicate voting may bump against limitations of size or power. Memory requirements for the mission may be too small to interleave memory words over the requisite number of die to avoid overwhelming available EDAC. Moreover, while it may be possible to implement triplicate voting among the cores of a multi-core processor, most current processors share common cache amongst the cores, leading to possible common failure modes. The extent to which these failure modes manifest for a particular processor could perhaps be revealed by fault injection, but without an SEE test validation, such a solution could still yield unpleasant surprises.

As noted above, destructive SEE often pose the most significant radiation related risks for short-duration mission. In addition, it is very difficult to predict from process/similarity data or other proxy data which parts are most likely to exhibit significant susceptibilities to DSEE. Mitigation of DSEE by derating to safe operating conditions may pose significant penalties for power efficiency or performance, and cold sparing can significantly complicate board design and take up precious board space.

Next in line in terms of difficulty in mitigation are very disruptive SEE, particularly SEFI and nondestructive SEL. These SEE modes likely require a power cycle to recover part functionality, and implementing the ability to cycle power only to likely affected devices (e.g., memories, processors, etc.) would complicate circuit design unacceptably for most SmallSats. As such, recovery usually implies cycling power to the affected board, with the requisite loss of data that implies, and then restoring the configuration of the board from an image stored in nonvolatile memory. Whether such a disruptive recovery process is acceptable depends on mission requirements. Whether it can be implemented successfully, depends on how tight board space is in the design. Certainly, it would be advantageous to know whether such mitigation is needed before committing to it. Unfortunately, such knowledge requires SEE testing.

7.2.3 SEE Test Challenges for SmallSat Missions

Although the limited budgets of most SmallSats pose significant constraints on the amount of SEE testing, they are not the only and sometimes not even the most serious constraints. In addition to being expensive, SEE testing is also time consuming, especially for complex, highly integrated COTS parts. Complex parts require complicated test hardware to measure SEE susceptibilities. Even if all one seeks to determine in an SEE test is whether the part is susceptible to SEL, the test hardware must be sufficient to verify whether the part remains fully functional during and after the test. As mentioned above, such a susceptibility test can be carried out in a single high-fluence, high-LET test run at a heavy-ion accelerator. Unfortunately, beam time is only one of many costs incurred during an SEE test. Figure 7-2 illustrates the breakdown of the cost (~\$80,000) a typical test of a complex device by cost category and phase of the test.




Figure 7-2: Although the high cost of SEE testing is driven by many factors, direct costs for beam time are among the less significant drivers. Nearly 70% of test costs are for highly skilled labor, and over 50% of the cost is spent in the development phase. This makes it difficult to realize savings by “simplifying” the test.

As can be seen from this breakdown, 69% of the cost is for labor, with most of this being during the test development phase. Typical beam costs account for only about 15% of the test cost, so even if beam time were eliminated entirely, this would reduce the cost by only ~\$12 K. While a simpler test would also save money on test development and the savings would exceed ~50%.

Although a proton test requires less part preparation to ensure the beam reached the device SV, more time may be spent by the radiation analyst finding out about the part technology so that the maximum information can be gleaned from the results. This is particularly important if one is trying to constrain heavy-ion susceptibility using proton data. Moreover, while a proton test is simpler (there being only one ion), it is useful to irradiate with several proton energies to constrain proton upset rates. Proton testing at the board or box level can yield significant savings over part-level testing. Moreover, it saves on test hardware, as the test object is usually a flight board or box, and sophisticated test equipment is superfluous for a test that yields little data except whether test-object performance remains nominal or not. Also, a board-level proton test yields no information on part-to-part variation, little constraint on heavy-ion DSEE risk, and results may be challenging to interpret in a way that yields understanding of the SEE modes of the parts in the system. Even accepting these limitations, a board- or box-level proton test of complex hardware is difficult to carry out for less than \$30-40K. Costs for proton beam time, travel and analysis can easily reach \$20 K.

Using lasers for SEE testing can save significantly on beam costs. However, part preparation costs will be as high as if not higher than those for heavy-ion testing, and costs will be similar for test equipment, development, execution and analysis. Moreover, because a laser beam illuminates only a small area on a test part for each run, will likely require refocusing several times during the test and because results must be tallied for each run, it is unlikely that the tester will be able to cover the full die area of a large, complex microcircuit in even a few days of testing. However, laser testing must be done at the level of an individual part (or a portion thereof) rather than a board or box, and it may yield limited information on DSEE susceptibility.

Thus, limited heavy-ion testing, proton testing, board-level proton testing and laser testing can realize some savings at a cost of yielding less information that can be used to infer in-flight performance. However, the savings are limited due to the cost breakdown of SEE testing.

7.3 Summary

By their very nature as secondary payloads, many SmallSats are driven to use of COTS parts to ensure schedule, size, power and performance constraints and requirements are met. The fact that conventional satellite missions, which do most of the radiation testing, use COTS in limited numbers, the short product life cycles and the sheer number of COTS parts available make it unlikely that historical data for that part—or even similar parts—will be available.

Size, power and performance requirements also limit the sorts of mitigations that can be included a priori for SmallSats in an SEE tolerant design. Use of redundancy to ensure reliability and availability of satellite services can be difficult to incorporate widely on a board with limited space. Ensuring sufficient derating to avoid failures/errors from worst-case parts can impose unacceptable penalties on performance, and monitoring currents and cycling power to avoid failures and recover functionality are likely to be unacceptable on a wide scale.

Finally, the cost structure of SEE testing makes it very difficult to leverage risk tolerance into significant savings for a test that constrains SEE susceptibility in any meaningful way. Most of the costs for SEE testing occur during the development phase, and most go to highly trained people developing and carrying out the test.

The problem of making SmallSats SEE tolerant is a work in progress. If there were a way of using commercial parts, even with well managed risk, the incentives of using COTS would have conventional satellite already doing it. We discuss some of the partial solutions SmallSat builders have been using next.

8 SmallSat SEE Hardness Assurance

In the previous section, we have discussed some of the constraints on SmallSat builders that make it difficult to implement a rigorous SEE Hardness Assurance program. In particular the emphasis on schedule and size restrictions strongly incentivizes SmallSat builders to use COTS parts, and although conventional SEE HA can handle COTS parts in small quantities where they promise a performance advantage worth the cost of qualification, large quantities of COTS pose significant SEE risks and are very costly to qualify. SmallSat builders have responded to the SEE challenge with a range of strategies, depending on their resources and their level of understanding of SEE. We outline a few of these below.

8.1 Buy-It-And-Fly-It

If a SmallSat's budget does not permit radiation testing, then the mission itself winds up being the radiation test—a strategy somewhat glibly referred to as buy-it-and fly it. Most commercial microelectronics and COTS systems will be susceptible to SEU. If the electronics has complex control and state circuitry, SEFI are a strong possibility, and any time commercial CMOS is flown in a radiation environment, SEL is a significant threat. However, sometimes a builder gets lucky. The probabilities of a lucky break can be enhanced significantly if the hardware has prior flight heritage, if the voltages are derated significantly from rated values, and if temperatures are on the low side (room temperature down to liquid nitrogen temperatures). If the mission duration is short, the probabilities may work out. The probability of success can be further increased if the mission plans for failures—that is, facilitating error/failure recovery through use of watchdog timers, error checking and other techniques discussed in section 6.2. Judicious use on diverse redundancy—e.g., two different radio systems or processors—can also significantly improve chances of success, albeit at a cost in size, weight and power that may not be acceptable for a SmallSat.

If critical deployment and communications components are made sufficiently reliable and telemetry is available for the early phases of the mission, even if the satellite fails, it may be possible to determine the likely cause of the failure, enabling gradual improvement and greater probability of mission success over time. However, if building a satellite out of commercial hardware, there is no guarantee that two purchases of the same system will use the same components. Components that are functionally the same may have significantly different radiation behavior. Moreover, if one chooses two different components (e.g., radio, on-board computer, etc.), it is still possible that the systems may be using some of the same components, increasing the risks of common failure modes for these systems. In short, drawing reliable conclusions about observed failures requires knowledge of each part in the design.

8.2 The Swarm—Safety in Numbers?

A Swarm mission is a large number of satellites flying in formation. Such a mission may provide global communications coverage or imaging or scientific data. If the number of satellites in the swarm significantly exceeds the number needed for mission success, then the mission success probability can exceed the survival probability a single satellite. However, such a rosy conclusion depends on the same assumptions that underlie all redundancy-based mitigations—namely that the failures of each satellite are independent and that the probability of failure for any one satellite during the mission is $\ll 1$. If the satellites all have the same design

and utilize the same components, they may have a common failure mode. Even though SEE failures are Poisson, if, for instance, each satellite contains the same highly SEL susceptible SRAM, that common failure mode may doom the swarm as surely as it dooms a single satellite. Also, since Poisson failure rates are additive, the expected number of failures increases linearly with the satellites in the swarm. It is only when the number of failures required for the mission to fail significantly exceeds the expected number that the swarm survivability significantly exceeds that for an individual satellite.

8.3 Board- and Box-Level Testing

As discussed previously in Section 5.2.2, board- or box-level SEE testing with high-energy protons can be one of the least expensive options for gleaning at least some knowledge of SEE susceptibility of a system. Testing with protons with energy ≥ 200 MeV will reveal most of the proton-induced SEE that can occur in the space environment. It will also reveal some information about susceptibilities to GCR and SPE light ions—particularly susceptibilities with low onset LET, rapidly rising σ vs. LET and shallow SV. When a SEE tolerant design survives a proton SEE test to a relatively high fluence, this does provide some assurance that system is not exceptionally susceptible. However, quantifying how susceptible and the degree of assurance can be difficult, especially when testing is conducted at a high level of integration—e.g. at the board or box level.

We have already discussed some of the limitations of proton testing in general as a proxy for heavy-ion SEE susceptibility—mainly the limitations on recoil ion coverage posed by TID and the likelihood that protons may not reveal all destructive SEE susceptibilities. As illustrated in Figure 5-6, even inferring the equivalent LET spectrum due to a given proton fluence is challenging, and may vary from device to device and even within a device for different SEE modes. In addition, the more complicated the system under test, the more difficult it is to draw general conclusions from the test. In current microprocessors, for example, the number of SEUs that give rise to observable effects (errors, crashes, etc.). The unobserved—or silent—faults have no observable effect during the test. However, if the same fault occurred when the processor was performing a different set of operations the consequences of the same SEU could be different. Fault injection studies can provide some measure of understanding of the proportion of faults that are silent, but exhaustive simulation of all faults at all stages of operation is not possible. For a board or box containing processors, FPGAs, DDR3 SDRAMs and Flash memory, the problem of silent corruption is likely to be even more significant. Thus, to the issue of spatial coverage over the die surface(s), one must also add the issue of temporal coverage over the state space of the system under test. Finally, as mentioned above, if components used in the test item exhibit significant part-to-part or lot-to-lot variation, this will not be revealed by a test on a single board or box. Even tests of several test units may not reveal the extent of variability or the effects that can manifest when these variable parts interact with each other.

Moreover, lest the reader conclude that the problems all arise from use of proton recoils as a proxy for space-environment ions, figure 8-1 illustrates the difficulty of defining the LET of an ion in a board-level heavy-ion test. The Variable Depth Bragg Peak test method [85] has been developed specifically for ultrahigh-energy ion beams like those available at the NASA Space Radiation Laboratory. In this method, the test item is exposed to an ultrahigh-energy ion beam, and the upsets are tallied for a small ion fluence. Then the exposure is repeated with various thicknesses of energy degrader, increasing the LET as the ion energy decreases toward the Bragg Peak. When the SEE cross section reaches its maximum and then decreases as more degrader is

added, one can conclude this degrader thickness corresponds to the Bragg peak for the ion overlapping the device SV. The method allows SEE cross sections to be measured over a range of LETs, from (near) minimum ionizing up to the Bragg peak without changing ions. Unfortunately, when testing at the board level, each device on the board may have different amounts of overburden over the SV, leading to slightly different LET for each device for a given degrader thickness. In addition, silent data corruption and coverage of the state space can also be an issue for ultrahigh-energy heavy ion testing, although ion fluence is less constricted by TID.

As mentioned above, the key to board-level testing with either protons or heavy ions is knowing as much as possible about the parts on the board prior to the test. Knowing the overburden or SV thickness is key to understanding the LET(s) to which they are exposed. Knowing the process of each die and its complexity may assist in developing a better SEE test regime, with the entire board exposed to a light fluence of particles sufficient to provide adequate coverage for simple devices, and complex devices fabricated in deep submicron (and hopefully TID hard) CMOS an additional higher fluence to provide better coverage for the device.

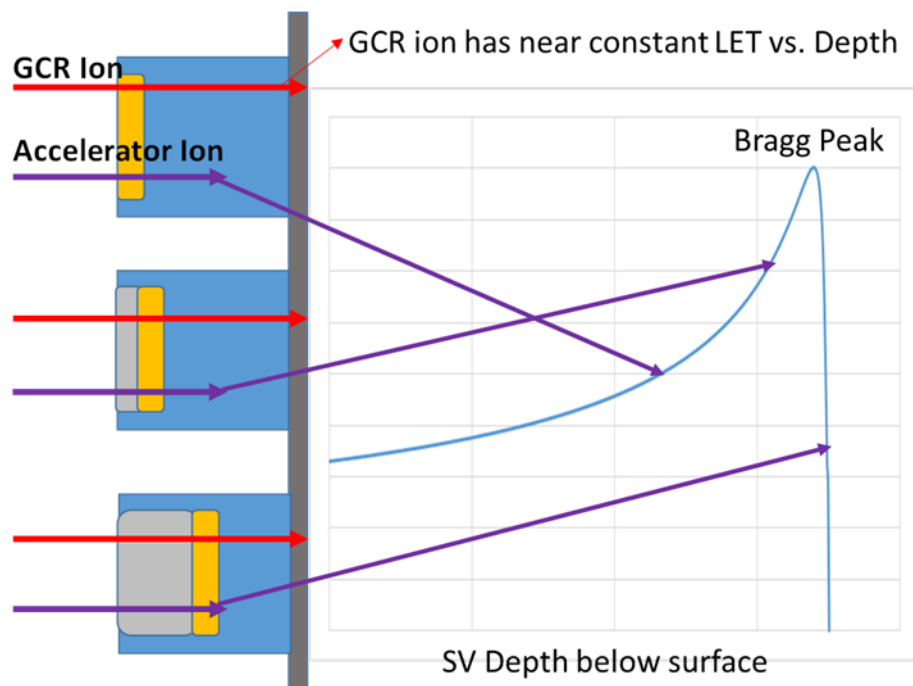


Figure 8-1: Varying amounts above the device sensitive volume can result in ions having different LET in different devices during ultra-high energy board level testing. The rapid changes in energy as one nears the Bragg peak exacerbate this effect for lower energy beams or thicker overburden.

8.4 Partnering—Relying on the Kindness of Strangers

As will undoubtedly be discussed throughout all segments of this short course, SmallSat projects labor under significant restrictions not just of budget, but of other resources as well—particularly schedule (and so to order long-lead-time radiation-hardened parts) and the ability to consult with experts. If the SmallSat builder can partner with a conventional satellite builder, this can relieve some of these constraints. A conventional builder will retain experts who can provide limited consultation on radiation issues the SmallSat builder faces. There may even be “leftover” radiation-hardened parts in stock, saving scarce funds—and more important the risk of schedule

slip due to a late delivery. In return, the SmallSat builder may be able to fly technology of interest to both parties, but which is too risky for conventional satellites. Thus, the SmallSat becomes a technology demonstration mission for the conventional builder, who has an incentive to ensure that the mission is at least sufficiently successful that the performance of the new technology can be assessed. Often, the new technology will be employed in a flight computer, large-scale data storage or a payload sensor and the associated data conversion and processing. The conventional builder may provide a redundant radiation hardened processor for the flight computer that can assess system health and performance and send the data to ground.

To date, most such missions have been flown by organizations associated with NASA centers, Department of Energy Laboratories and other conventional satellite builders. However, there is no reason partnerships could not be formed between the same conventional builders and universities or other SmallSat builders.

8.5 Risk-Informed Testing

If a project has a limited, but nonzero budget for SEE HA, the goal should be to maximize risk reduction for the funds that are available. In the absence of data on the SEE susceptibilities of the parts being considered, one cannot determine what risks they pose to the application. However, one can work backward from the system level to determine the system consequences of various SEE modes, and evaluate the technologies of the parts to assess the likelihood of those SEE modes. Based on these considerations, one can assign a risk for each SEE mode deemed credible according to the consequences assuming that mode is realized. In other words, it is the risk assuming the probability of the SEE mode is 1.

An example of the risk-informed testing strategy is the Goal Structured Notation approach adopted by Vanderbilt University for its ... CubeSat experiment. An example of a Goal Structured Notation (GSN) graph is shown in Figure 8-3.

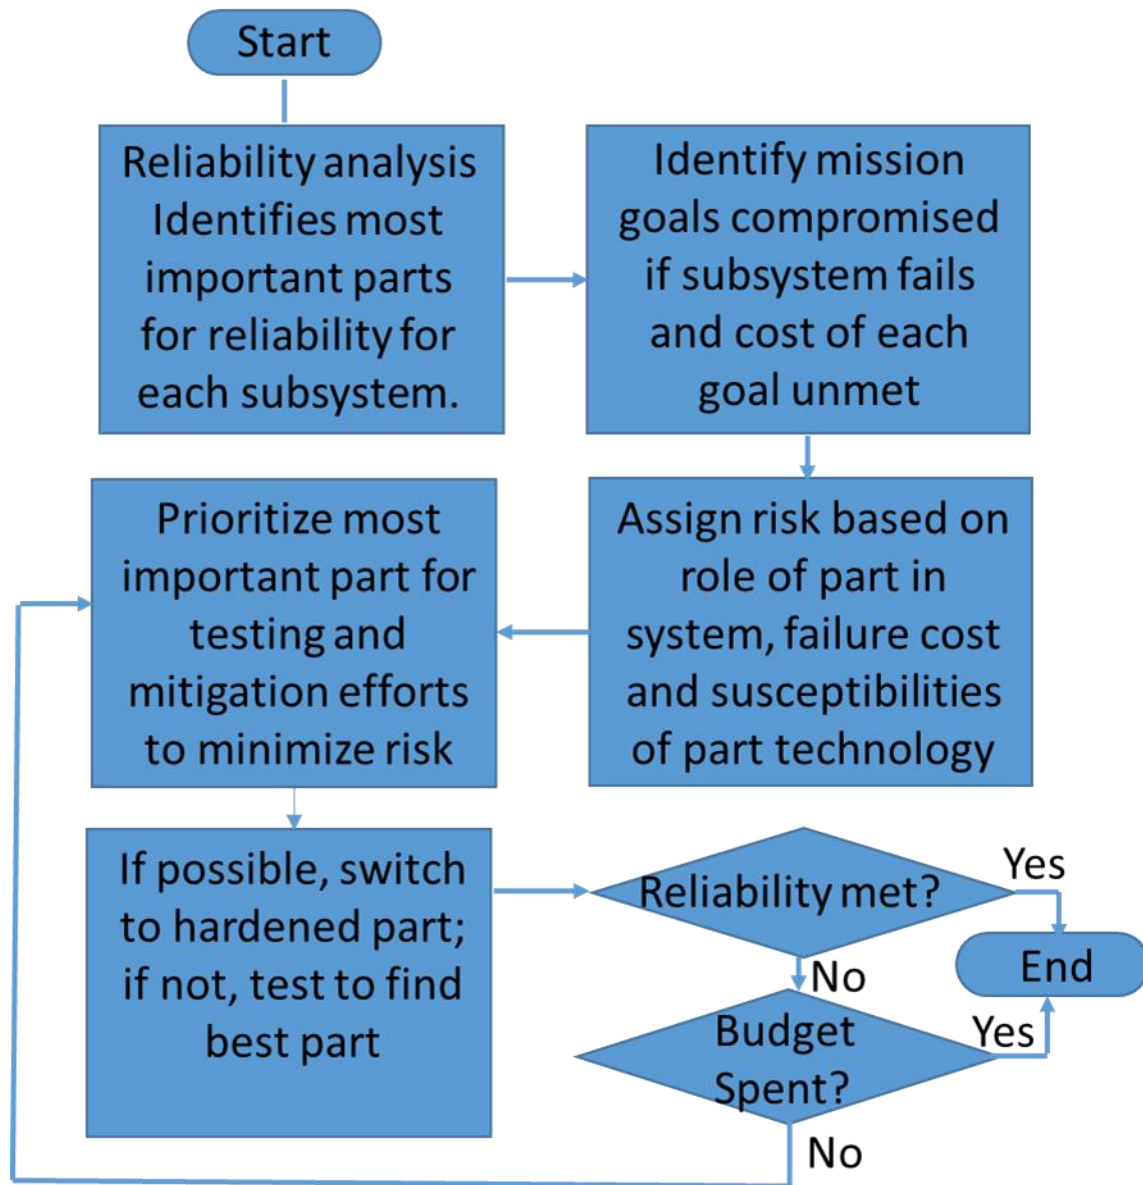


Figure 8-2: For risk-informed SEE testing, the reliability analysis defines the parts most significant for the system. SEE risk is prioritized bases on the consequences of failure and the credible failure modes for the technology.

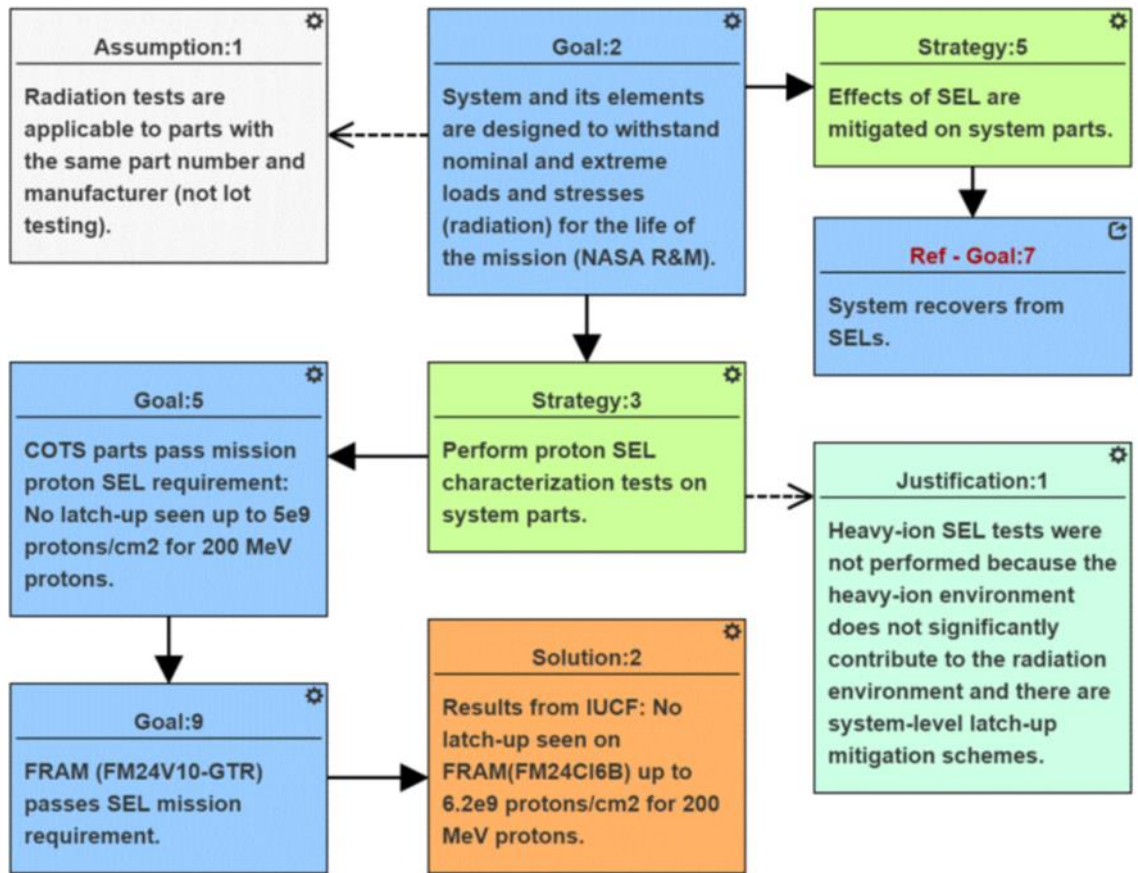


Figure 8-3: Goal-Structured Notation (GSN) is a formalism for identifying both risks and the means for ameliorating those risks. The approach ties risks to the goals they threaten and encourages assumptions and justifications to be stated explicitly. (Courtesy of Rebekah Austin, Vanderbilt University.)

8.5.1 Evaluating Failure Consequences for SmallSats

For the most part, assigning failure consequences is straightforward. Mission goals and requirements can be prioritized as a percentage of the overall mission value. Thus, if a part critical to mission success has credible failure modes, the consequences of its failure could be equal to the mission cost. Indeed, it could exceed the mission cost if it were the cause of a mission failure, its responsibility for the failure went undetected and it went on to cause failures on future missions. As such, in evaluating the priority to assign to a given part, one may need to consider its potential importance not just for the current mission but also for future missions the SmallSat builder may be planning.

There is another way in which failure cost can exceed mission cost, especially if a series of failures are experienced by the same satellite provider. At some point, funding agencies may simply decide not to throw good money after bad. This is especially true for conventional satellite builders. If they experience too many failures, their brand may be damaged. SmallSat builders may be less brand conscious, but they cannot afford to fall too far behind their peers in terms of reliability. The converse of this consideration is also true: If success in a challenging mission brings recognition and helps win future contracts and/or funding, the value to the satellite provider may actually exceed the value of the mission goals in and of themselves.

8.6 Examples

One example of the sort of partnership outlined in Section 8.4 is the CIBOLA Flight Experiment (CFE). CFE was funded by the US Department of Energy and developed as a by Los Alamos National Laboratory as a technology demonstration platform. Although the main technology experiment involved antennae and radios, the payload computer, which processes the payload experiment data is also an experiment in radiation-tolerant, reconfigurable computing. It was one of the first space missions to pioneer the use of commercial reprogrammable FPGAs. The computer implements several types of error checking and watchdog timers to detect anomalies. In addition to multiple Xilinx reconfigurable FPGAs to handle the main computing, there is also a BAE RAD6000 radiation hardened processor, which supervises payload operations. Volatile data storage is in radiation hardened SRAM. However, nonvolatile memory includes a radiation tolerant EEPROM as well as FLASH. Nonvolatile memory is used to store configuration as well as data in the event configuration in the FPGAs is lost due to SEE. CFE is an example of a SmallSat (albeit a very well-funded one) that did a lot of things right. The use of diverse redundancy, triplicate voting of the FPGAs, inclusion of error checking and watchdog timers, selective use of radiation hardened technologies and multiple types of nonvolatile configuration storage significantly reduce SEE risk to the payload, either by reducing SEE probability or by facilitating detection and recovery from the SEE.

Another interesting example from the commercial space sector shows that the concept of a “SmallSat” is relative. The Peregrine lunar lander (see figure 8-4) being developed by Astrobotic Technologies is a mission designed to deliver payloads to the lunar surface at a cost of \$1.2 million per kilogram of payload. Although the lander is projected to weigh 345 kg dry weight, and 1200 kg fueled. Both the cost and mass are quite lean for a lunar mission. Moreover, the mission is planning to launch as a secondary payload, thus reducing the importance of launch costs in the budget. Also, use of redundancy for SEE hardening is not an optimal strategy, since every gram of redundant hardware reduces the payload that can be carried.

Peregrine is being driven to heavy use of COTS parts by cost, schedule, performance and mass considerations. While they plan to leverage off of parts—especially COTS—previously tested by the radiation effects community and similarity if they cannot find data on specific parts they need, they will also probably be forced to test some parts themselves. In order to prioritize scarce testing resources, Astrobotic Technologies is carrying out a risk-informed approach taking into account the criticality of the system, the likely SEE susceptibilities of parts making up the system and the relative cost and impact of testing vs. mitigation vs. changing mission requirements to reduce risk. Even with all of this consideration and flexibility, the radiation approach for Peregrine is still a work in progress. For assistance in development of their approach, they are working with NASA experts under the Lunar Cargo Transportation and Landing by Soft Touchdown (Lunar CATALYST) program—thus adding aspects of partnering, risk-informed testing, board/system-level testing and even some buy-it-and-fly-it to their approach.

Another example of the risk-informed testing strategy is the Goal Structured Notation approach adopted by Vanderbilt University for its ... CubeSat experiment. An example of a GSN graph is shown in figure 8-3. In this notation, the goal is stated in the top level and is followed by strategies for meeting the goal, derivative goals associated with that strategy and assumptions and justifications for the strategies. One can take issue with the assumptions and

justifications—e.g., we have seen that proton testing can underestimate SEL susceptibility—but that is the point. The GSN approach makes a point of explicitly stating the assumptions underlying the strategy so they can be evaluated.

The examples discussed in this section illustrate that the approaches outlined above are more a matter of convenience and a pedagogical aid than a reflection of real SEE HA programs for SmallSats. In reality, a successful program will have to remain flexible and combine elements of all these approaches—and perhaps some yet to be developed.




Figure 8-4: Artist's rendering of the Peregrine Lunar Lander. [Used by permission from Astrobotics.]

9 Conclusions

The previous analysis has emphasized the nature of SEE HA as an example of risk-mitigation approach as a pedagogical aide to understanding the challenges of applying the methodology to SmallSats and other risk-tolerant platforms. In short, the problem with risk-tolerant platforms is that they take risks, and if the risks are not assessed thoroughly, they can easily wind up exceeding their risk budget many times over. In particular, because SmallSats are secondary payloads constrained by schedule, size and power, they are driven to take technical and reliability risks by using predominantly COTS parts. Because COTS parts are designed with no consideration of how they will operate in a space radiation environment, and because the conventional satellite builders who do most of the SEE testing use and test COTS only sparingly, there is little data available upon which to base a SEE risk assessment. This leaves SmallSat builders with a choice of either blowing a big hole in their cost budget by performing extensive, costly SEE testing, blowing a big hole in their risk budget by accepting very high risk bounds—or worst of all, using proxies—e.g. proton test data and/or similarity data—in unrepresentative ways that may yield overly optimistic risk estimates.

Moreover, SEE HA is already tolerant of risk, in the sense that it bases its assessment of threat credibility on known susceptibilities. I noted in the short course in 2007 that this has led to surprises in the past, and it continues to yield surprises now. As such, it is critical to the health of SEE HA that heavy-ion testing continue to be part of the methodology even if SmallSats and other risk-tolerant missions assume a larger role. Semiconductor parts will continue to evolve, and we will be blind to new failure modes if we don't keep looking for them.

Finally, SEE HA approaches need to evolve to better leverage risk tolerance to realize cost savings for missions with tight economic and schedule budgets, while at the same time maintaining a knowledge base of current semiconductor part technologies—from radiation hardened through COTS.

We will address each of these considerations in turn, first summarizing the issues discussed above concerning the difficulty of adapting conventional SEE HA to SmallSats, then revisiting the risks inherent in current SEE HA methods and finally looking to the future to consider how SEE HA methods could evolve to better address threats to all mission types.

9.1 Risk Tolerance and SEE Risk Management

The challenges for SmallSat SEE risk management occur at all levels of the process. Novel technologies may have unknown SEE susceptibilities, and while one can look at “similar” technologies as a guide, the conclusions will be uncertain, and in some cases, even the determination of what constitutes a similar technology may be uncertain. In some cases, the part chosen may exhibit much greater sensitivity than previous parts [48], or the part-to-part variability may exceed previously tested “similar” parts [49].

Probably the most salient challenges for SmallSat SEE HA arise due to the difficulty of using COTS in space radiation environments. Conventional satellites face many of the same issues regarding COTS use. However, SmallSats are driven to use COTS parts in large numbers by considerations of schedule, size, performance and at least some cases by cost, and SmallSats have far fewer resources to apply toward solving them. COTS use is always challenging for

several reasons. First, commercial electronics is designed to operate in a terrestrial environment where radiation is considered to be a minor issue for designers, if it is a consideration at all. Second, because conventional satellite builders use COTS only sparingly and because there are often many COTS parts that can fulfill a particular function, it is unlikely that a SmallSat builder will find prior SEE test data on that part. Additionally, the complexity of many COTS parts means that much of the test data available will be application specific, and the short product life cycle for most COTS parts means that SEE test results are nearly worthless after only a few years. The same considerations also ensure that similarity data will be virtually nonexistent. Although an analysis may be attempted based on part technology, such analyses rarely yield meaningful limits on SEE risk, and the fact that most COTS vendors consider process data highly proprietary will limit the information available to the analyst.

The complexity, performance and degree of integration of COTS parts also often make them more difficult to test than their military and radiation-hardened counterparts. Packaging and the presence of multiple stacked die may prevent ions from a conventional heavy-ion accelerator from reaching the device sensitive volumes. Although proton SEE testing can produce recoil ions in device SV, the low recoil ion fluences, lack of knowledge of which ion caused an SEE and differences in SV depths make complicate the interpretation of test results. Indeed, in a complex System-In-a-Package (SIP) with multiple die, the same proton fluence may correspond to a different fluence vs. equivalent LET for each die. If the projects opts to test at the board or box level, interpretation of the test results becomes even more complicated.

Finally, although the approaches to mitigating SEE risks are well understood, the consequences of these mitigations in terms of increased size and power consumption or decreased performance may not be acceptable for the SmallSat. Some mitigations, such as simple EDAC, memory scrubbing, watchdog timers and error checking, are straightforward to implement, and these may be implemented even in the absence of test data for the parts. On the other hand, triplicate voting, hot and cold sparing, interleaving of bits from memory words across multiple die and other mitigations may require too much board space to be implemented easily for a SmallSat.

The problems of implementing economical, risk-tolerant SEE HA for SmallSats seem nearly insurmountable. Indeed, they have yet to be resolved in a satisfactory, comprehensive fashion. However, in section 9.4, we will recap some of the approaches developed to date and suggest some additional improvements. Next, we look at some of the new challenges facing SEE HA that affect both SmallSat and conventional platforms.

9.2 Another Decade of SEE Surprises

Even for conventional satellite platforms, SEE hardness assurance remains a dynamic field. To some extent, this is inevitable when the parts used in satellites continue to evolve at an exponential pace. However, the field has also experienced challenges to many of its fundamental assumptions over the past decade. I noted several of these challenges in my 2007 short course on system level radiation hardening. Table VI notes some of the SEE related surprises of the past decade.

Table VI: The Latest Surprises in SEE Hardness Assurance

Rule of Thumb	Oops!
SEL gets better with decreasing temperature	SEL observed in Read-out IC at cryogenic temperatures <40 K (<i>Marshall-2010, [42]</i>)
SEE are not important for diodes	Schottky diodes observed to fail catastrophically due to previously unknown SEE mechanism. (<i>Casey-NEPP Electronic Technology Workshop 2012</i>)
If part has no W plugs, no need to worry about proton destructive SEE if onset LET >20 MeVcm ² /mg.	p+Au fission in packages w/ Au plated lids produces ions w/ LET ~40 MeV-cm ² /mg (<i>Turflinger-2015, [34]</i>)
No need to worry about proton-induced fission if package has no high-Z materials	σ for p + Pd (Z=46) fission nearly as high as for p + Au (Z=79). Fission seen even for p + Ni (<i>Turflinger-2016, [35]</i>)

As can be seen from the table, surprises that challenge fundamental assumptions of SEE HA occur roughly every 2-3 years. In some cases, these affect only niche applications (e.g. there are relatively few CMOS applications running at temperatures below 40 K, so cryogenic SEL is not a major concern in most missions). In others, the threat is significant for a few parts when particular conditions apply (e.g. the fission events in high-Z materials). However, in the case of the failures observed in Schottky diodes and similar parts, the implications can be profound. These observations introduce an entirely new class of parts into SEE HA that could previously be ignored. To date, the failure mode has been seen only in parts biased at >50% of their rated reverse voltage, so as long as such a derating is acceptable for the project, mitigation via threat avoidance is straightforward. However, such a derating may require use of less efficient and larger components, which could pose challenges for SmallSats. More important, the discovery of this new failure mechanism demonstrates the importance of SmallSat builders remaining current with developments in SEE HA. We detail this and other recommendations below.

9.3 Improving SEE HA for SmallSats... and For All Sats

Although some of the challenges facing SmallSat SEE HA are unique—e.g. the emphasis on schedule, low budgets for SEE testing—many are common to any mission seeking to capitalize on the performance conferred by COTS. As such, development of more SmallSat friendly approaches to SEE HA could have benefits far beyond those platforms. Below, we sketch out some possible suggestions for the SmallSat approaches discussed in Section 8.

- 1) Buy-it-and-fly-it—Although the name suggests a cavalier attitude toward SEE HA and indeed toward reliability in general, using COTS as is can be implemented to achieve increasing reliability over time. However, for this to occur the overall SmallSat must be sufficiently reliable that root cause of failure can be determined. This suggests that a

hybrid approach could be useful—a basic, reliable architecture supervising commercial hardware that supplies the needed performance. Chances of success could be improved if diverse redundancy of the commercial hardware could be used for critical systems, although this is costly in scarce board space and system complexity. In general, the better one understands the hardware being flown, the less likely one is to be surprised by error/failure modes that could occur on orbit. For instance, if one is able to fly two different programmable radios, it is essential that the components inside the radios be assessed to ensure that critical functions are fulfilled by different components so that common failure modes can be avoided.

- 2) Shelter from the Swarm—Swarm missions are one type of mission where diverse redundancy could be both a realistic and helpful mitigation, because the diversity could be implemented over the entire swarm without taking up valuable board space in individual satellites. Such a scheme would require additional development, integration and testing, which would be a strain for a small team. However, skillfully implementing diverse systems across the constellation of satellites and monitoring differential success/survival rates could allow the project could select the most reliable combinations and improve reliability over time. This would be of great value not just for SmallSats, but for conventional builders as well.
- 3) Don't Get Board—Although board-level testing can save on test costs, it means giving up a lot of information about device performance and a lot of generality in understanding how a given device on the board would perform in another system. Also, unless one is willing to test many boards/boxes, any significant amount of part-to-part variability in any of the parts on the board could invalidate the test results as predictors of on-orbit performance. Nonetheless, there are suggestions that could improve results. One key to better proton testing is ensuring that complex parts receive adequate coverage. This might be achieved by supplementing an initial, low-fluence board test with high-fluence tests of complex, deep-submicron parts that require higher fluences to yield representative error/failure mode samples. The rastering capabilities of proton accelerators at many of the medical facilities would be well suited to delivering such focused proton exposures.
Moreover, the excellent penetration of proton beams means that one can test multiple boards back to back. Not only does this increase the total fluence to which the test boards are exposed, it also can alert the tester to any serious variability from one test board to the next that might render the tests unrepresentative of flight boards.
- 4) Although partnering SmallSat producers with producers of conventional satellites has the potential to benefit both parties, the strategy has yet to be exploited to any significant degree. Most such collaborations have been between one arm of an organization that builds conventional satellites and another arm that builds or is seeking to start building SmallSats. However, as the success of the Cibola collaboration shows, such partnerships can be beneficial to both parties and to the radiation community generally.

Risk-informed testing is another promising approach. In some ways, this approach is not new. The SEE Criticality Approach proposed by NASA in the 1990s also advocated prioritizing testing by the criticality of the application.[36] However, the issue here is that the resources for testing are more limited, and the options considered for reducing risk bounds are more diverse.

9.4 Final Thoughts

As one would expect of a discipline tasked with managing risk for components that change at an exponential rate, SEE hardness assurance has always been dynamic. The current challenges posed by a new, more complicated technologies, new mission platforms and ever tightening budgets seem daunting until we remember that when SEE were discovered, the community had no idea how to assess part susceptibilities, estimate rates or how the threat would evolve. The path we have trod sheds light and provides perspective on the path we are treading.

10 References

- [1] Binder, D., Smith, E. C., Holman, A. B., "Satellite Anomalies from Galactic Cosmic Rays," IEEE Trans. Nucl. Sci., Vol. 22, No. 6, pp. 2675-2680.
- [2] Kolasinski, W. A., Blake, J. B., Anthony, J. K., Price, W. E., Smith, E. C., "Simulation of Cosmic-Ray Induced Soft Errors and Latchup in Integrated-Circuit Computer Memories," IEEE Trans. Nucl. Sci., Vol. 26, No. 6, pp. 5087-5091.
- [3] Wyatt, R. C.; McNulty, P. J.; Toumbas, P.; Rothwell, P. L.; Filz, R. C., "Soft Errors Induced by Energetic Protons," IEEE Trans. Nucl. Sci. Vol. 25, No. 6, 4905-4910.
- [4] Soliman, K. A.; Nichols, Donald K., "Latchup in CMOS Devices from Heavy Ions," IEEE Trans. Nucl. Sci., Vol. 30, No. 6, pp. 4514-4519.
- [5] Waskiewicz, A. E., Groninger, J. W., Strahan, V. H., Long, D. M., "Burnout of Power MOS Transistors with Heavy Ions of Californium-252," IEEE Trans. Nucl. Sci., Vol. 33, No. 6, pp. 1710-1713.
- [6] Fischer, T. A., "HEAVY-ION-INDUCED, GATE-RUPTURE IN POWER MOSFETs," IEEE Trans. Nucl. Sci., Vol. 34, No. 6, pp. 1786-1791.
- [7] Puchner, H.; Tausch, H. J.; Koga, R., "Proton-Induced Single Event Upsets in 90nm Technology High Performance SRAM Memories," IEEE 2011 Radiation Effects Data Workshop (REDW), Summerlin, NV, 27 July 2011, pp. 161-163.
- [8] "History of UltraCMOS® Technology Platform," <http://www.psemi.com/technologies/ultracmos-technology-platform/history-of-ultracmos-technology-platform>, 4/14/17.
- [9] Sierawski, B. D. *et al.*, "Muon-induced single event upsets in deep-submicron technology," IEEE Trans. Nucl. Sci., vol. 57, no. 6, pp. 3273–3278, 2010.
- [10] Trippe, J. M. *et al.*, "Electron-Induced Single Event Upsets in 28 nm and 45 nm Bulk SRAMs," IEEE Trans. Nucl. Sci., vol. 62, no. 6, pp. 2709–2716, 2015.
- [11] Normand, E. and Baker, T. J., "Altitude and latitude variations in avionics SEU and atmospheric neutron flux," IEEE Trans. Nucl. Sci., Vol. 40, No. 6, pp. 1484-1490.
- [12] Schroeder, B. Pinheiro, E. and Weber, W.-D., "DRAM Errors in the Wild: A Large-Scale Field Study," Presented at SIGMETRICS/Performance'09, June 15–19, 2009, Seattle, WA, USA. <https://www.cs.toronto.edu/~bianca/papers/sigmetrics09.pdf>
- [13] Bradley, P. D. and Normand, E. "Single event upsets in implantable cardioverter defibrillators," IEEE Trans. Nucl. Sci., Vol. 45, No. 6, pp. 2929-2940.
- [14] Siefert, N., Jahinuzzaman, S, Velamala, J., Patel, N., "Susceptibility of planar and 3D tri-gate technologies to muon-induced single event upsets," 2015 IEEE International Reliability Physics Symposium.

- [15] Bagatin, M. *et al.*, "Muon-induced soft errors in 16-nm NAND flash memories," 2015 IEEE International Reliability Physics Symposium.
- [16] Hands, A., Lei, F., Ryden, K., Dyer, C., Underwood, C., Mertens, C., "New Data and Modelling for Single Event Effects in the Stratospheric Radiation Environment," IEEE Trans. Nucl. Sci., Vol. 64, No. 1, pp. 587-595.
- [17] Baumann, R. C., "Landmarks in Terrestrial Single-Event Effects," Part B of the Short Course presented at the 2013 Nuclear and Space Radiation Effects Conference, San Francisco, CA, 8 July 2013
- [18] Dodds, N. A. *et al.*, "The Contribution of Low-Energy Protons to the Total On-Orbit SEU Rate," IEEE Trans. Nucl. Sci., Vol. 62, No. 6, pp. 2440-2451.
- [19] Garrett, H. B., Kim, W. and Evans, R. W., "Updating the Jovian Plasma and Radiation Environments: The Latest Results for 2015," JOURNAL OF SPACECRAFT AND ROCKETS Vol. 53, No. 4, July-August 2016, pp. 693-707.
- [20] Tylka, A. J. *et al.*, "CREME96: A Revision of the Cosmic Ray Effects on Micro-Electronics Code," IEEE Trans. Nucl. Sci., Vol. 44, No. 6, pp. 2150-2160.
- [21] Xapsos, M. A., Stauffer, C., Gee, G. B., Barth, J. L., Stassinopoulos, E. G., and McGuire, R. E., "Model for Solar Proton Risk Assessment," IEEE Trans. Nucl. Sci., Vol. 51, No. 6, pp. 3394-3398.
- [22] McDonald, F. B., "Cosmic ray modulation in the heliosphere. A phenomenological study," Space Sci. Rev., 83, 33-50.
- [23] Garrett, H. B., Ratliff, J. M., and Evans, R. W., "SATURN RADIATION (SATRAD) MODEL," JPL Publication 05-9, October 2005.
- [24] Plainaki, C. *et al.*, "Planetary space weather: scientific aspects and future perspectives," J. Space Weather Space Clim., 6, A31 (2016) DOI: 10.1051/swsc/2016024
- [25] N. L. Grigorov, et al., "Evidence for Anomalous Cosmic Ray Oxygen Ions in the Inner Magnetosphere," Geo. Res. Ltrs., Vol. 18, pp. 1959 (1991).
- [26] Roth, C., "AE9/AP9/SPM Radiation Environment Model Users Guide Version 1.35.001," AFRL Virtual Distributed Laboratory (VDL), released 1/3/2017, https://www.vdl.afrl.af.mil/programs/ae9ap9/files/package/Ae9_Ap9_SPM_Users_Guide.pdf.
- [27] AE9/AP9/SPM Development Team, "AE9/AP9 V1.30.001 Model Comparison Summary Report: Spatial Comparisons," Air Force Research Laboratory, 5 January 2016.
- [28] Pellish, J. A. *et al.*, "Impact of Spacecraft Shielding on Direct Ionization Soft Error Rates for Sub-130 nm Technologies," IEEE Trans. Nucl. Sci., Vol. 57, No. 6, pp. 3183-3189.
- [29] O'Neill, P., Badhwar, G., and Culpepper, W., "Risk Assessment for Heavy Ions of Parts Tested with Protons," IEEE Trans. Nucl. Sci., Vol. 44, No. 6, p. 2311-2314.
- [30] O'Neill, P. M., Badhwar G. D., and Culpepper, W. X., "Internuclear cascade-evaporation model for LET spectra of 200 MeV protons used for parts testing," IEEE Trans. Nucl. Sci., vol. 45, pp. 2467 - 2474.
- [31] Hiemstra, D. M. and Blackmore, E. W., "LET Spectra of Proton Energy Levels From 50 to 500 MeV and Their Effectiveness for Single Event Effects Characterization of Microelectronics," IEEE Trans. Nucl. Sci., Vol. 50, No. 6, page 2245 - 2250.
- [32] Ladbury, R. L., Lauenstein, J.-M. and Hayes, K. P., "Use of Proton SEE Data as a Proxy for Bounding Heavy-Ion SEE Susceptibility," IEEE Trans. Nucl. Sci., Vol. 62, No. 6, p. 2505 - 2510.

- [33] Schwank, J. R. *et al.*, “Effects of Particle Energy on Proton-Induced Single-Event Latchup,” IEEE Trans. Nucl. Sci. Vol. 52, No. 6, pp. 2622 – 2629.
- [34] Turflinger, T. L. *et al.*, “RHA Implications of Proton on Gold-Plated Package Structures in SEE Evaluations,” IEEE Trans. Nucl. Sci. Vol. 62, No. 6, pp. 2468 – 2475.
- [35] Turflinger, T. L. *et al.*, “Proton on Metal Fission Environments in an IC Package: An RHA Evaluation Method,” IEEE Trans. Nucl. Sci. Vol. 64, No. 1, pp. 309 – 316.
- [36] 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>.
- [37] 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.
- [38] Layton, P. *et al.*, “SEL Induced Latent Damage, Testing, and Evaluation,” IEEE Trans. Nucl. Sci., Vol. 53, No. 6, pp. 3153-3157, 2006.
- [39] JESD57, “TEST PROCEDURE FOR THE MANAGEMENT OF SINGLE-EVENT EFFECTS IN SEMICONDUCTOR DEVICES FROM HEAVY ION IRRADIATION,” JEDEC, <https://www.jedec.org/standards-documents/results/jesd57> (1996)
- [40] Johnston, A. H., Hughlock, B.W., Baze, M. P., and Plaag, R. E., “The Effect of Temperature on Single-Particle Latchup,” IEEE Trans. Nucl. Sci. Vol. 38, No. 6, pp. 1435-1441.
- [41] Levinson, J., Even, O., Adler E., Hass, M., Ilberg D., and Lifshitz Y., “Single Event Latchup (SEL) in IDT 7187 SRAMs-dependence on ion penetration depth,” Second European Conference on Radiation and its Effects on Components and Systems (Cat. No.93TH0616-3).
- [42] C. J. Marshall *et al.*, “Mechanisms and temperature dependence of single event latchup observed in a CMOS readout integrated circuit from 16–300 K,” IEEE Trans. Nucl. Sci., vol. 60, no. 6, pp. 3078–3086.
- [43] Ladbury, R. L. and Campola, M. J., “Bayesian Methods for Bounding Single-Event Related Risk in Low-Cost Satellite Missions,” IEEE Trans. Nucl. Sci., vol. 60, no. 6, pp. 4464–4469.
- [44] Sexton, F. W., “Destructive Single-Event Effects in Semiconductor Devices and ICs,” IEEE Trans. Nucl. Sci, Vol. 50, No.3, pp. 603-621, 2003.
- [45] Allenspach, M. *et al.* “Evaluation of SEGR Threshold in Power MOSFETs,” IEEE Trans Nucl. Sci., Vol. 41, No. 6, pp. 2160-2166, 1994.
- [46] Lauenstein, J.-M., Ladbury, R. L., Goldsman, N., Kim, H. S., Batchelor, D. A., Phan, A. M., “Interpreting Space-Mission LET Requirements for SEGR in Power MOSFETs,” IEEE Trans Nucl. Sci., Vol. 57, No. 6, pp. 3443 - 3449, 2010.
- [47] S. Liu *et al.*, “Effects of ion species on SEB failure voltage of power DMOSFET,” IEEE Trans. Nucl. Sci., vol. 58, no. 6, pp. 2991–2997.
- [48] 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.
- [49] George, J. S. *et al.*, “Response Variability in Commercial MOSFET SEE Qualification,” IEEE Trans. Nucl. Sci., vol. 64, no. 1, pp. 317–324.
- [50] 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.

- [51] Casey, M. C. *et al.*, “Destructive Single-Event Failures in Schottky Diodes,” NEPP Electronic Technology Workshop, NASA Goddard Space Flight Center, Greenbelt, MD. June 11-12, 2012, https://nepp.nasa.gov/workshops/etw2013/talks/Wed_June12_2013/1430_Casey_Schottky_Diode_Radiation_Failures.pdf
- [52] Casey, M. C., Lauenstein, J.-M., Ladbury R. L., Wilcox, E. P., Topper A. D., and LaBel, K. A., “Schottky Diode Derating for Survivability in a Heavy Ion Environment,” IEEE Trans. Nucl. Sci., vol. 62, No. 6, pp. 2482-2489.
- [53] Xapsos, M. A.; O’Neill, P. M.; O’Brien, T. P., “Near-Earth Space Radiation Models,” IEEE Trans. Nucl. Sci. Vol. 60, No. 3, pp. 1691 – 1705.
- [54] Mewaldt, R. A. *et al.*, “RECORD-SETTING COSMIC-RAY INTENSITIES IN 2009 AND 2010,” Ap. J. Lett. 723:L1–L6, 2010 November 1.
- [55] King, J. H., “Solar proton fluences for 1977–1983 space missions,” J. Spacecraft, vol. 11, pp. 401–408, 1974.
- [56] Feynman, J., Armstrong, T. P., Dao-Gibner, L. and Silverman, S. M., “New interplanetary proton fluence model,” J. Spacecraft, vol. 27, pp. 403–410, 1990.
- [57] Feynman, J., Spitale, G., Wang, J., and Gabriel, S., “Interplanetary fluence model: JPL 1991,” J. Geophys. Res., vol. 98, pp. 13281–13294, 1993.
- [58] Stassinopoulos, E. G. and King, J. H., “Empirical solar proton models for orbiting spacecraft applications,” IEEE Trans. Aerospace and Elect. Sys., vol. 10, pp. 442–450, 1974.
- [59] Xapsos, M. A., Summers, G. P., Barth, J. L., Stassinopoulos, E. G. and Burke, E. A., “Probability model for cumulative solar proton event fluences,” IEEE Trans. Nucl. Sci., vol. 47, no. 3, pp. 486–490, Jun. 2000.
- [60] M. A. Xapsos, C. Stauffer, T. Jordan, J. L. Barth, and R. Mewaldt, “Model for cumulative solar heavy ion energy and linear energy transfer spectra,” IEEE Trans. Nucl. Sci., vol. 54, pp. 1985–1989.
- [61] Xapsos, M. A., Summers, G. P., J. L. Barth, J. L., Stassinopoulos, E. G. and Burke, E. A., “Probability model for worst case solar proton event fluences,” IEEE Trans. Nucl. Sci., vol. 46, No. 6, pp. 1481–1485.
- [62] M. A. Xapsos, C. A. Stauffer, T. M. Jordan, J. H. Adams Jr., and W. F. Dietrich, “Periods of high intensity solar proton flux,” IEEE Trans. Nucl. Sci., vol. 59, no. 4, pp. 1054–1059.
- [63] Dyer, C. S. *et al.*, “Observation of Solar Particle Events from CREDO and MPTB During the Current Solar Maximum,” IEEE Trans. Nucl. Sci., vol. 49, No.6, pp. 2771–2775.
- [64] Ladbury, R. L. *et al.*, “Use of Commercial FPGA-Based Evaluation Boards for Single-Event Testing of DDR2 and DDR3 SDRAMs,” IEEE Trans. Nucl. Sci., vol. 60, no. 6, pp. 4457–4463.
- [65] Foster, C. C., O’Neill P. M., and Kouba, C. K., “Risk Assessment Based on Upset Rates from High Energy Proton Tests and Monte Carlo Simulations,” IEEE Trans. on Nucl. Sci. Vol. 55, No. 6, pp. 2962 – 2969.
- [66] Ladbury, R. L., Lauenstein, J.-M., “Evaluating Constraints on Heavy-Ion SEE Susceptibility Imposed by Proton SEE Testing and Other Mixed Environments,” IEEE Trans. on Nucl. Sci. Vol. 64, No. 1, pp. 301 – 308.
- [67] Hale, J. M. *et al.*, “Strong Correlation Between Experiment and Simulation for Two-Photon Absorption Induced Carrier Generation,” IEEE Trans. on Nucl. Sci. Vol. 64, No. 5, pp. 1133 – 1136.

- [68] D. Cardoza, *et al.*, "Comparison of Single Event Transients Generated by Short Pulsed X-Rays, Lasers and Heavy Ions," *IEEE Trans. on Nucl. Sci.* Vol. 61, No. 6, pp. 3154 – 3162.
- [69] Petersen, E. L., "The SEU Figure of Merit and Proton Upset Rate Calculations," *IEEE Trans. Nucl. Sci.*, Vol. 45, No. 6, p. 2550-2562, 1998.
- [70] Peyrard, P.-F., Beutier, T., Serres, O., Chatry, C., Ecoffet, R., Rolland, G., Radiation and its Effects on Components and Systems, RADECS 2003, Proceedings of the 7th European Conference, held 15-19 September 2003 in Noordwijk, The Netherlands. Edited by K. Fletcher. ESA SP-536, ESA/ESTEC, 2004., p.639
- [71] Adams, J. H., Jr. *et al.*, "CRÈME: The 2011 Revision of the Cosmic Ray Effects on Micro-Electronics Code," *IEEE Trans. on Nucl. Sci.* Vol. 59, No. 6, pp. 3141 – 3147.
- [72] Reed, R. A. *et al.*, "Physical Processes and Applications of the Monte Carlo Radiative Energy Deposition (MRED) Code," *IEEE Trans. Nucl. Sci.*, vol. 62, no. 4, pp. 1141–1161.
- [73] Warren, K. M. *et al.*, "Monte-Carlo Based On-Orbit Single Event Upset Rate Prediction for a Radiation Hardened by Design Latch," *IEEE Trans. Nucl. Sci.*, vol. 54, no. 6, pp. 2419–2425, 2007.
- [74] Warren, K. M. *et al.*, "Application of RADSAFE to Model the Single Event Upset Response of a 0.25 μm CMOS SRAM," *IEEE Trans. Nucl. Sci.*, vol. 54, no. 4, pp. 898–903, 2007.
- [75] Bendel, W. L. and Petersen, E.L., "Predicting Single Event Upsets in the Earth's Proton Belts," *IEEE Trans. Nucl. Sci.*, vol. 31, 1201-1206, 1984.
- [76] Stapor, W. J.; Meyers, J. P.; Langworthy, J. B.; Petersen, E. L., "Two Parameter Bendel Model Calculations for Predicting Proton Induced Upset," *IEEE Trans. Nucl. Sci.*, vol. 37, 1966-1973, 1984.
- [77] Morris, R. D. and Foster, C. C., "Cross-Section Estimation in the Presence of Uncertain Dosimetry," *IEEE Trans. Nucl. Sci.*, vol. 57, no. 6, pp. 3528–3536.
- [78] Ladbury, R. L., "Statistical Properties of SEE Rate Calculation in the Limits of Large and Small Event Counts," *IEEE Trans. Nucl. Sci.*, vol. 54, no. 6, pp. 2113–2119.
- [79] Koga, R. *et al.* "Single Event Transient (SET) Sensitivity of Radiation Hardened and COTS Voltage Comparators," 2000 Radiation Effects Data Workshop, 24-28 July 2000 pp. 53 - 60.
- [80] Guertin, S. M., Allen, G. R., Sheldon, D. J., "Programmatic Impact of SDRAM SEFI," 2012 IEEE Radiation Effects Data Workshop (REDW), Miami, FL, 19 July 2012, pp. 87-94.
- [81] Shannon, C. E., "A mathematical Theory of Communication," *Bell System Tech. Jour.*, **27**, pp. 379-423 and 623-656, 1948.
- [82] R. Hamming, "Error Detecting and Error Correcting Codes," *Bell System Tech. Jour.*, **29**, pp.147-160, 1950.
- [83] 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.
- [84] 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.
- [85] Buchner, S., Kanyogoro, N., McMorrow, D., Foster, C. C., O'Neill, P. M., and Nguyen, K. V., "Variable Depth Bragg Peak Method for Single Event Effects Testing," *IEEE Trans. Nucl. Sci.*, Vol. 58, No. 6, pp. 2976-2982, 2011.

11 Acronyms

ADC—Analog to Digital Converter
ADI—Analog Devices, Incorporated
AU—Astronomical Unit
BJT—Bipolar Junction Transistor
CAD—Computer Aided Design
CFE—Cibola Flight Experiment
CME—Coronal Mass Ejection
CMOS—Complementary Metal Oxide Semiconductor
COTS—Commercial Off The Shelf
DDR—Double-Data Rate
DRAM—Dynamic Random Access Memory
DSEE—Destructive Single-Event Effect
ECC—Error Correction Code
EDAC—Error Detection and Correction
FET—Field Effect Transistor
FOM—Figure of Merit
FMECA—Failure Modes and Effects Criticality Analysis
FPGA—Field Programmable Gate Array
GCR—Galactic Cosmic Ray
GEO—Geostationary Equatorial Orbit
GLM—Generalized Linear Model
GSN—Global Structured Network
IRPP—Integral Rectangular Parallelepiped
ISS—International Space Station
IEEE—Institute for Electrical and Electronics Engineers
LET—Linear Energy Transfer
LET_{EQ}—Equivalent Linear Energy Transfer
MBU—Multibit upset
MCU—Multi-cell upset
MEO—Medium Earth Orbit
MOSFET-- Metal Oxide Semiconductor Field Effect Transistor
NSREC—Nuclear and Space Radiation Effects Conference
REDW—Radiation Effects Data Workshop
SCR—Silicon Controlled Register
SDRAM—Synchronous Dynamic Random Access Memory
SEC—Single-Error Correct
SECEDED—Single-Error Correct, Double-Error Detect
SEE—Single-Event Effect
SEB—Single-Event Burnout
SEDR—Single-Event Dielectric Rupture
SEEHA—SEE Hardness Assurance
SEFI—Single-Event Functional Interrupt
SEGR—Single-event gate rupture
SEL—Single-Event Latchup

SESB—Single-event Stuck Bit
SET—Single-Event Transient
SEU—Single-Event Upset
SIP—System In a Package
SPE—Solar Particle Event
SRAM—Static Random Access Memory
SV—Sensitive Volume
TNS—Transactons on Nuclear Science