Single Event Upsets in SEL Relays

Derrick Haas and Karl Zimmerman

Introduction

This paper provides an overview of single event upsets (SEUs), including their causes, mitigation methods, and, most importantly, impact on SEL microprocessor-based protective relays. We quantify at a high level the impact that SEUs have on system protection and close with best practices for ensuring reliability.

Background

As early as 1979, the computing industry knew of transient memory failures (or soft memory errors) resulting from high-energy particles [1]. Later research and review of data from the Cray-1 mainframe computer in Los Alamos, New Mexico, revealed evidence that an SEU (a type of soft memory error) occurred on that machine in 1976 [2]. Another publication in 1979 documented an SEU that occurred in space in 1975 [3]. These references show that for decades, this phenomenon has been known and has been documented extensively in the computing industry, the aviation industry, and space exploration.

Soft memory errors are defined as “random, nonrecurring, single bit errors in memory devices” [1]. A soft error is not permanent, and the memory device recovers completely by the following write cycle with statistically no greater chance of error recurrence at that location than at any other bit location in any other memory component in the device. Soft memory errors do not damage the components themselves. SEU is nearly synonymous with soft memory error, but an SEU is not specific to a memory component. Similar soft error phenomena can occur with other digital components that make up modern microprocessor-based relays. In this paper, we refer to these errors as SEUs.

Causes

SEUs are caused by high-energy particles, which come from two primary sources: cosmic rays radiating particles that interact with the earth’s atmosphere and trace elements in semiconductor packaging material that emit particles. We provide an overview of each source below.

As high-energy particles from cosmic rays collide with atoms in the earth’s atmosphere, other particles are released as a result. These subsequent particles can then go on to collide with other atoms, and some particles may eventually reach the earth’s surface. Figure 1 illustrates the collision process.

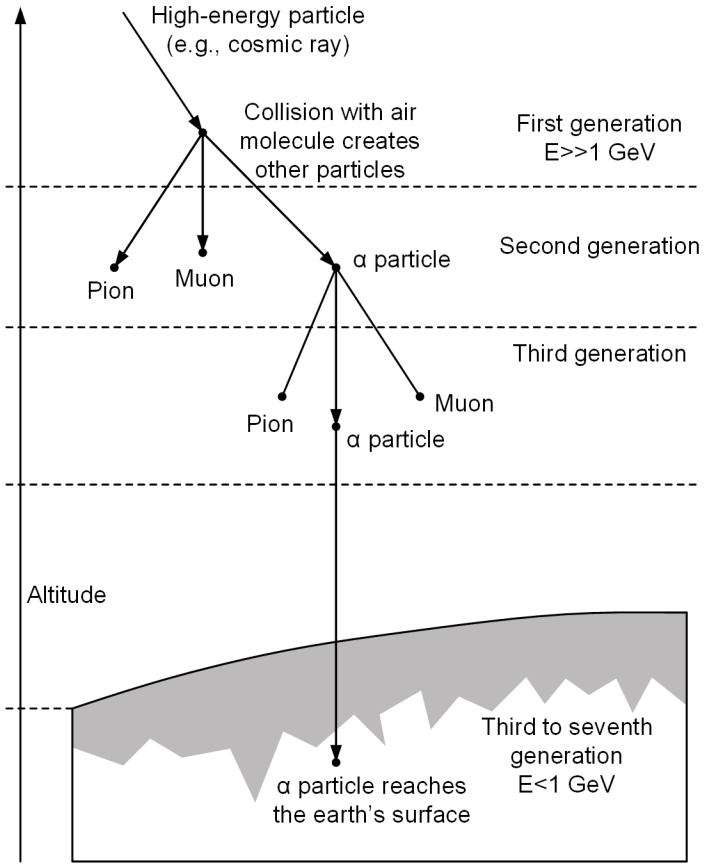

Figure 1 Diagram of particle collisions in the earth's atmosphere

Of particular interest are collisions with nitrogen and oxygen molecules in the earth's atmosphere because these collisions often result in the creation of high-energy neutrons and of alpha (α) particles, which consist of two protons and two neutrons each. The collisions create other particles as well, such as pions and muons. However, it is the high-energy neutrons and the alpha particles in particular that can cause SEUs. Figure 2 shows a rendering of a cosmic ray bombarding the earth's atmosphere and the numerous collisions and particles that a single cosmic ray can generate.

Figure 2 Artwork of cosmic rays hitting Earth (credit: Mark Garlick/Science Photo Library)

Of note is the amount of energy that a particle possesses, which is measured in electron-volts (eV). The energy required to accelerate one electron through a potential difference of one volt is equivalent to 1 eV. A particle must have a sufficient energy level to cause an SEU, and certain particles will not interact with silicon to the same degree (beta and gamma particles have very low energy loss rates in silicon). Generally, in SEU studies, testing, and literature, only alpha particles and neutrons with energies of 1 MeV (one million electron-volts) and higher are considered [4]. The amount of energy required for a particle to cause an SEU is very dependent on the design of the digital component (e.g., processor or memory component), including aspects such as geometry and critical charge required for state change.

The rate at which these particles pass through an area is called the particle flux. This is given as the number of particles passing through an area over an amount of time, with units of particles per cm^2 per hour. The particle flux gives us an idea of how many of these particles are present and can help evaluate the likelihood of a particle colliding with a digital component and causing an SEU. Because the earth's magnetic field and atmosphere impact many of these particles, the flux or frequency of the particles observed is higher at high altitudes and near the earth's magnetic poles. SEU occurrence rates are listed with the assumption that the component or equipment is at sea level at the latitude and longitude of New York City. Normalization factors can be used to convert SEU rates based on different latitudes, longitudes, and altitudes.

Figure 3 shows the neutron flux levels at various altitudes. We can see that the neutron flux peaks at an altitude of approximately 60,000 ft above sea level and is several hundred times greater than the neutron flux at sea level. Similar data provide flux levels based on latitude [5]. Because of the higher levels of neutron flux at high altitudes, aeronautics and space exploration industries have an added interest in the impacts of high-energy particles on computing systems, including SEUs.

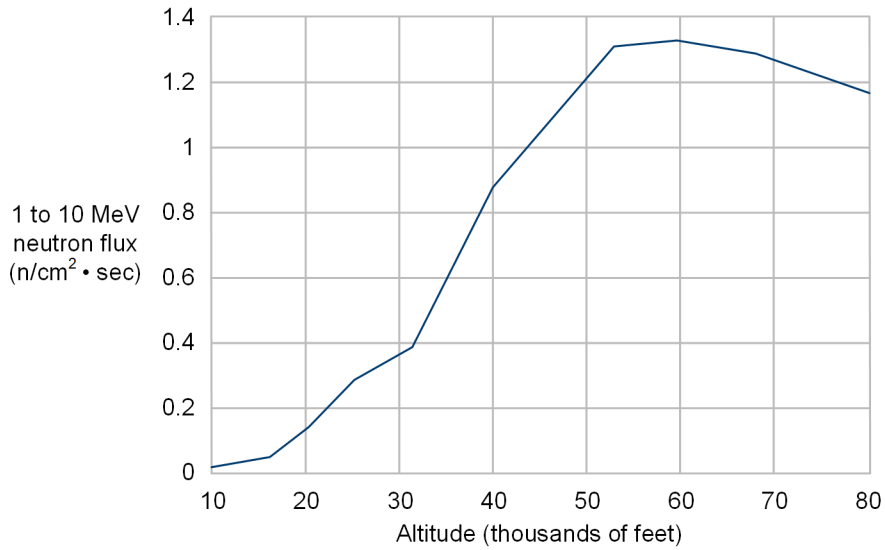

Figure 3 Neutron flux versus altitude [5]

The second source of high-energy particles that can cause an SEU was documented in 1979 [1]. Essentially every material has uranium, thorium, and other heavy radioactive elements present in small quantities. Digital component packaging material can therefore contain traces of these heavy elements. As the radioactive elements in the packing material decay, they often emit alpha particles. For clarity, the packaging or packing material in a microprocessor, memory chip, or an integrated circuit in general is the material (e.g., plastic) that encapsulates the semiconductor material that makes up the microprocessor. Figure 4a shows a simplified diagram of a semiconductor device and its packaging material. Figure 4b shows a microprocessor with a portion of the packaging material removed.

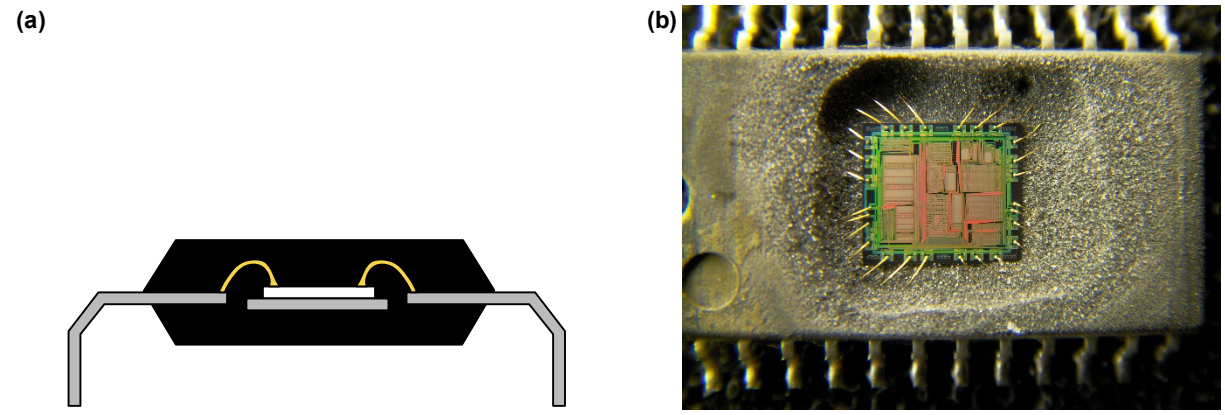

Figure 4 Simplified diagram of semiconductor packaging material (a) and a microprocessor with some packaging material removed to expose the semiconductor device (b) [6]

Not much can be done to eliminate the SEU-causing sources like alpha particles and other high-energy particles coming from space. Fortunately, the earth's atmosphere does an excellent job of shielding us and our electronic devices from high-energy particles, making the statistical likelihood of an SEU resulting from a space particle relatively low for devices installed at low altitudes. Of course, at high altitudes or at an installation in the northern latitudes, the probability of a cosmic ray (or derivative particle) causing an SEU is higher. For alpha particles resulting from integrated circuit packaging material, microprocessor manufacturers are working to limit the impact of trace elements in packaging material. Integrated circuit suppliers have made significant improvements

on packaging material quality and the number of impurities present. However, we cannot practically remove the sources of alpha particles or prevent the exposure of protective relays to them.

Bit Flip Mechanism

Now that we have established the sources of high-energy particles responsible for SEUs, we can share an example of how a high-energy particle causes a bit to flip. When an alpha particle collides with semiconductor material, it creates electron-hole pairs. This is theoretically possible in nearly every type and variety of memory element, processor, or gate. All digital components, from static RAM (SRAM) to dynamic RAM (DRAM) to field-programmable gate arrays (FPGAs) and more, have a non-zero SEU occurrence rate. However, certain digital components and their designs make SEUs more likely. Figure 5 shows the sequence of events that leads to a bit flip from a 0 to a 1 in a single dynamic memory cell.

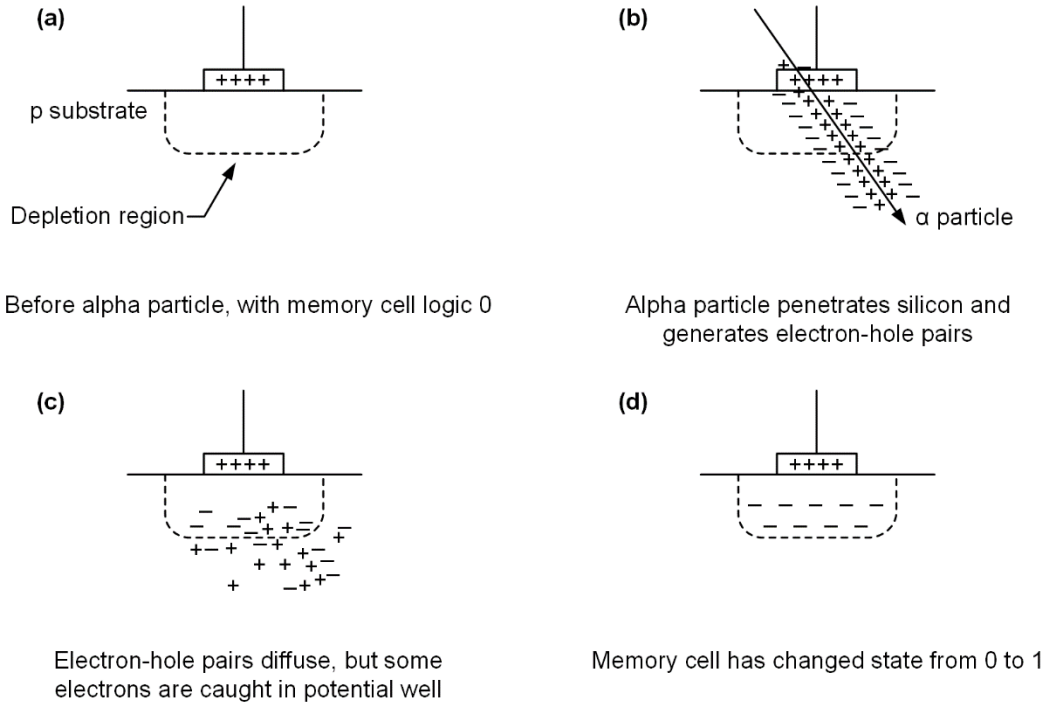

Figure 5 Process of memory change from 0 to 1 because of alpha particle collision

The high-energy particle creates electron-hole pairs as it passes through the semiconductor material (Figure 5b). An alpha particle with an energy of 5 MeV can create approximately 1.4 x 10^6 electron-hole pairs and typically penetrates 25 µm in silicon [1]. Most of the electron-hole pairs diffuse through the substrate material, as shown in Figure 5c. However, the potential well captures some of these electrons and repels the holes. It is these captured electrons in the depletion region that result in a state change from a 0 to a 1 in this dynamic memory cell (Figure 5d). The electrons trapped in the potential well diffuse over time. However, in certain systems, if a clock edge occurs before the electrons diffuse, the errant memory or bit flip is made permanent.

The location, geometry, and arrangement of the semiconductor device, the amount of critical charge, and other factors impact how a particular device experiences an SEU. In addition, many of the same factors impact what type of SEU is generated and whether there is a bit flip from 0 to 1 or from 1 to 0.

Statistical Likelihood of SEUs and Measurement and Testing of SEU Rates

The estimated statistical likelihood of an SEU occurring is expressed as units of failures in time (FIT). The FIT rate is typically measured in failures per billion hours. It is now common for digital component manufacturers to provide an estimated FIT rate specification for the component, be it a microprocessor, FPGA, or RAM variety such as synchronous DRAM (SDRAM). The total FIT rate for a protective relay is the combined FIT rate of all of the digital components needed for a relay to perform its required function. For example, if a process critical to the functioning of a protective relay relied on three different components, each with a FIT rate of 100 failures per billion hours, then the expected FIT rate for the relay is 300 failures per billion hours.

In addition to evaluating a component FIT rate, both component manufacturers and SEL are interested in determining the likelihood of an SEU occurring. Statistical models that predict the FIT rate of a memory cell have been around since SEUs were first discovered [1]. Testing the validity of such models is as important now as it was then. If a protective relay has a FIT rate of 400 failures per billion hours, that would equate to one failure every 285 years on average. However, waiting that long for an SEU to occur in order to validate the estimated FIT rate is beyond impractical.

Many components allow error injection, a way to simulate a bit flip without high-energy particle exposure. Another way to attempt to measure FIT rates is to place components or products in an environment with a higher exposure to alpha particles (or similar high-energy particles) than ground level. There are several high-energy particle sources where, statistically, the particle flux is significantly higher than what is observed naturally in a substation environment. These include nuclear reactors, particle accelerators, or similar energy sources that can generate high-energy particles. Figure 6 shows an SEL relay at a testing facility.

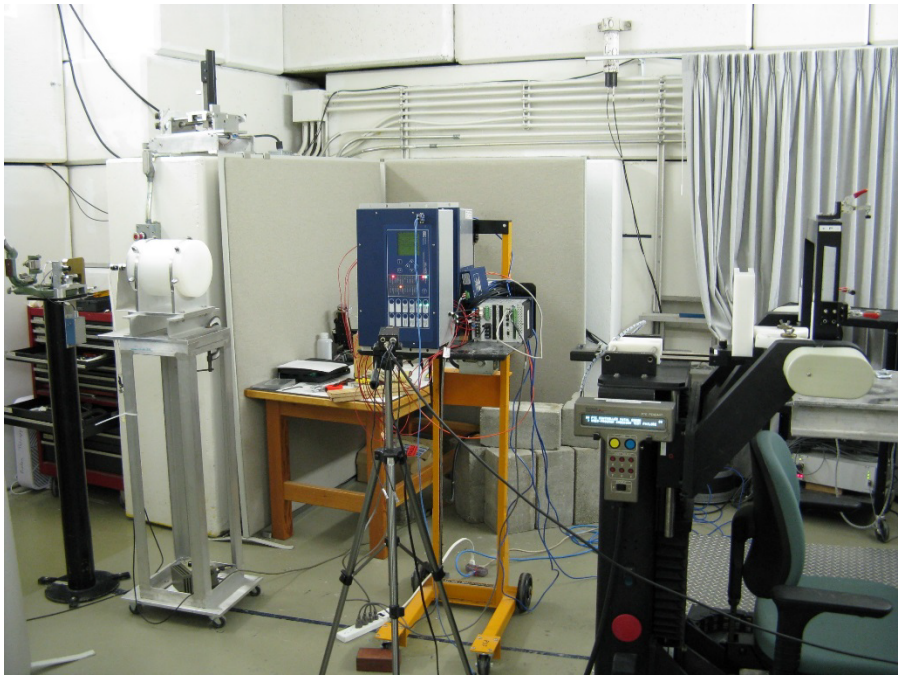

Figure 6 SEU testing of SEL relay

SEL's interest in SEU testing is not only to evaluate the FIT rate, but more importantly to test mitigation techniques, which are discussed later in this paper. By putting the relays in an

environment where SEUs occur much more frequently, we can evaluate the effectiveness of a variety of mitigating techniques.

SEL takes many factors into consideration when designing protective relays. These factors include the overall quality and reliability of a component, as well as a component's features, supplier, price, and more. FIT rate is also considered. SEL is presently implementing a design criterion to limit SEU rates to a mean time between SEUs (MTBSEU) of 500 years, equating to a FIT rate of approximately 228 failures per billion hours. That means not only using FIT rate as a criterion for evaluating individual components but also considering mitigating techniques for SEUs as part of our designs.

Gathering records of SEUs from field-installed relays can be more difficult. One corrective technique for an SEU is for the impacted device to restart or power cycle. Restarting the device overwrites the impacted device memory or instructions in a processor and removes the error. Most relays now log a time-stamped entry in the Sequential Events Recorder (SER) report when a relay restarts.

Impact on Protective Relays

The potential impact of an SEU can vary greatly. Microprocessors, FPGAs, and memory components are part of nearly every aspect of protective relaying, including analog to digital conversions, protection element algorithms and logic, and tripping decisions. An SEU that impacts a memory address related to a communications protocol may only result in a temporary loss of SCADA communication or a report of an errant SCADA point. An SEU that impacts a measurement related to the power system current that is used by a protection algorithm may result in incorrect measurement that could cause a protective element to incorrectly pick up. An SEU that directly impacts a Relay Word bit could result in an undesired operation. So, as these examples illustrate, the impact of an SEU can range from minor to severe.

In Figure 7, an event report shows an undesired trip of an SEL-311L Line Current Differential Protection and Automation System that has been attributed to an SEU [7]. There is no fault on the line, and the differential element and TRIP87 Relay Word bit assert for no clear reason.

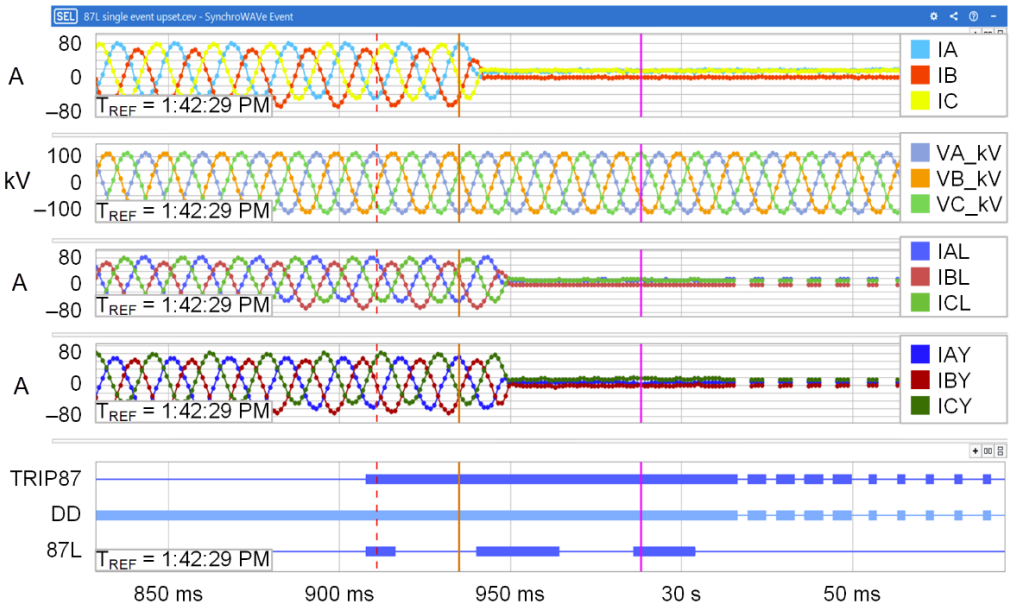

Figure 7 Event report showing SEL-311L trip as a result of an SEU

SEL relays come with mitigation capabilities to prevent undesired operations resulting from SEUs and, in some cases, to help the relay recover from SEUs gracefully and with minimal impact to protection. Furthermore, the statistical likelihood of an SEU causing an undesired operation is small. Based on field data, the mean time between undesired operations (MTBUO) due to SEUs over a five-year period (2012–2016) is greater than 50,000 years. Stated another way, if 50,000 relays are in service for one year, we will see one or fewer undesired operations due to SEUs.

Mitigation and Prevention Techniques

SEL uses several techniques to mitigate the effect of SEUs (see Table 1). Note that SEUs can go undetected and/or result in undesired operations despite mitigation techniques.

Table 1 Mitigation and Prevention Techniques and Considerations

Technique	Considerations
Part selection	Select parts with low FIT rates, and balance build requirements with better SEU tolerance.
Design	Create and use internal relay data (e.g., VECTOR and MEMORY commands) to assist in diagnosing relay memory if an in-service relay fails. Use error-correcting codes to detect and correct bit flips in real time. These require more memory and therefore impact hardware design. This method is used in SEL computers and in some SEL relays [8].
Relay disable	If an error is detected, clear the memory, assert the relay alarm contact or Relay Word bit (e.g., HALARM), and disable the relay. Relay users can cycle power once to see if the relay recovers. If a fault occurs while a relay is disabled, protection is disabled.
Diagnostic restart	If an error is detected, restart to maximize relay availability. Most SEL relay models have automatic diagnostic restart functionality (see Table 2). Many models disable relay and alarm contacts for multiple restarts in a specific time frame. The number of restarts and the time frames vary between relay models. Many SEL relays create a time-stamped entry in the SER report when a diagnostic restart occurs.

It is important to note that bit flips can also be the result of non-SEU events, such as component failures or manufacturing defects. One key distinction is that an SEU is random (the particle emissions are not random, but exactly when and where they hit the earth is random). The statistical likelihood of having a repeated SEU on the same relay is small. Or, put differently, SEUs are very unlikely to be repeating errors. If we consider a relay with a FIT rate of 1,000 failures per billion hours, that rate equates to approximately one failure every 114 years.

In 1996, SEL began to enable devices to automatically restart in the event of a detected error, starting with the SEL-321 Phase and Ground Distance Relay. SEL-321-1 Relays manufactured in 1996 or later included the global setting ERESTART. If a CRAM (CR_RAM) error is detected when ERESTART is set to Y, the relay automatically restarts and resets. In addition, several commands are available to users so they can gather diagnostic information should a relay fail.

Diagnostic restart is also included in other SEL products. Table 2 lists the firmware versions in which automatic restarting functionality was added to various SEL products.

Table 2 Firmware Revision When Automatic Restarting Functionality Was Added to SEL Products

SEL Product	Firmware Revision
SEL-100 series	Diagnostic restart not available
SEL-200 series	Diagnostic restart not available
SEL-500 series	Diagnostic restart not available
SEL-321	Available by setting in SEL-321-1 Relays with 1996 firmware versions or later: • R413—60 Hz, 5 A, 1 I/O board • R463—50 Hz, 5 A, 1 I/O board • R513—60 Hz, 5 A, 2 I/O boards • R563—50 Hz, 5 A, 2 I/O boards • R614—60 Hz, 5 A, 1 I/O board, ACB rotation • R714—60 Hz, 5 A, 2 I/O boards, ACB rotation • R813—60 Hz, 1 A, 1 I/O board • R863—50 Hz, 1 A, 1 I/O board • R913—60 Hz, 1 A, 2 I/O boards • R964—50 Hz, 1 A, 2 I/O boards
SEL-351 Protection System family	All R500 firmware
SEL-351-5,-6,-7 legacy	R405 and later
SEL-351S legacy	R405 and later
SEL-351A legacy	R405 and later
SEL-311C Protection System family	All R500 firmware
SEL-311A legacy	R110 and later
SEL-311B legacy	R110 and later
SEL-311C legacy	R113 and later
SEL-311L-0,-6	R163 and later
SEL-311L-1,-7	All R500 and all R300 firmware R414 and later with Ethernet R215 and later without Ethernet
SEL-421-0,-1	R200 and later
SEL-421-2,-3	R121 and later
SEL-421-4,-5	R310 and later
SEL-411L	R106 and later
SEL-487B	R112 and later
SEL-487B-1	R303 and later
SEL-487E-0,-2	R113 and later

SEL Product	Firmware Revision
SEL-487E-3,-4	R303 and later
SEL-487V	R104 and later
SEL-751A	R406 and later
SEL-751	All firmware versions
SEL-710	R403 and later
SEL-700G	R102 and later
SEL-849	All firmware versions
SEL-T400L	All firmware versions

Even the most robust systems can have some risk of an SEU causing an undesired operation. For dependability concerns, standard practices such as having both a primary and a secondary protective relay can ensure protection still operates for a fault in the event of a relay failure to operate. Recently, much of the focus from the power industry has been on security and undesired trips when there are no faults or disturbances. We offer some recommendations in the following section to address the impact of SEUs on protection systems.

SEL Recommendations

It is important to emphasize that SEUs are not the only cause of protective relay system failures or undesired operations. An SEU is only one failure mode. Additionally, we need to recognize that, statistically, SEUs cause only a small percentage of undesired operations. However, there are practical actions that users can take to reduce vulnerability to SEUs including the following:

- Always monitor relay alarm contacts.
- Always collect SER data when available.
- Always act on SEL service bulletins.
- Keep firmware and hardware updated to the latest versions when possible, especially when the upgrade provides automatic diagnostic restart functionality.
- Cycle power once if the relay is disabled and does not have automatic diagnostic restart functionality to see if the relay recovers, and call SEL Technical Support to report the results.
- Ensure that the protective relay and protection system are secure during a diagnostic restart or power cycle [9].
- Use best practices and data to assess risk and to improve security and dependability [10]. For example, consider adding security to the tripping control logic. This can be achieved by using a voting scheme as shown in Figure 8a (three relays—A, B and C), an interdependent tripping scheme as shown in Figure 8b (two relays—A and B), or other schemes designed to enhance security.

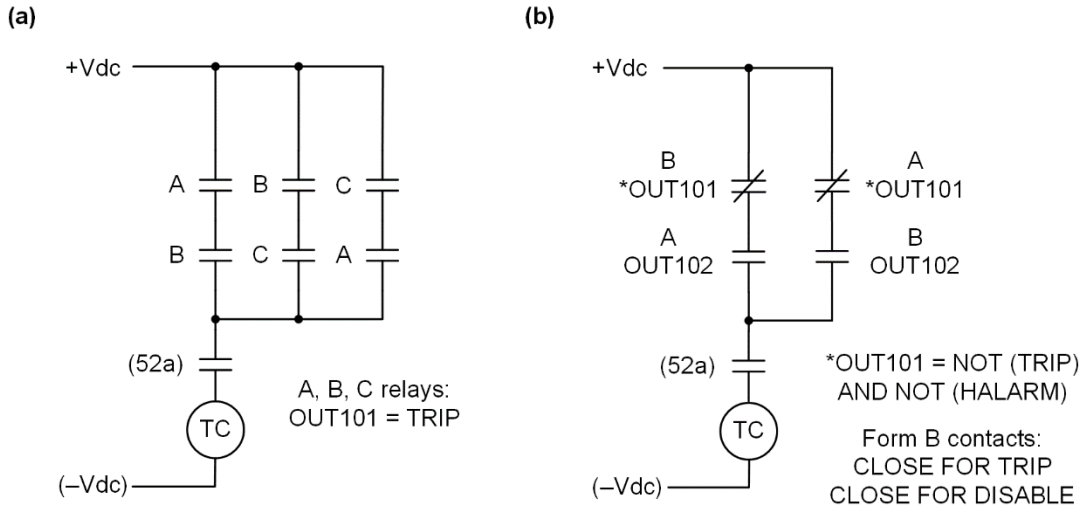

Figure 8 Simplified control logic for a relay voting scheme (a) and interdependent tripping scheme (b)

Automatic restarting clears the SEU. After restart, if the relay is still in failure mode, it is likely not the result of an SEU. Users should send the device back to SEL for evaluation and repair.

For relays or devices that do not have automatic restarting or do not have it enabled, we recommend gathering data from the device if possible. Contact SEL Technical Support for detailed instructions on how to collect the appropriate data.

Conclusions

SEUs and their impacts on electronic devices have been known for decades but have become of increasing importance to SEL customers. SEL has been aware of these phenomena and has applied mitigation techniques since 1996. The impact of most SEUs is causing a relay to disable or produce a diagnostic restart. In rare cases, an SEU can cause an undesired operation. SEL provides recommendations for reducing the risk of an SEU causing an undesired operation. We continuously monitor quality and work toward improved product design.

References

[1] T. C. May and M. H. Woods, “Alpha-Particle-Induced Soft Errors in Dynamic Memory,” *IEEE Transactions on Electron Devices*, Vol. 26, Issue 1, January 1979, pp. 2–9.

[2] E. Normand, J. L. Wert, H. Quinn, T. D. Fairbanks, S. Michalak, G. Grider, P. Iwanchuk, J. Morrison, S. Wender, and S. Johnson, “First Record of Single-Event Upset on Ground, Cray-1 Computer at Los Alamos in 1976,” *IEEE Transactions on Nuclear Science*, Vol. 57, Issue 6, December 2010, pp. 3114–3120.

[3] J. F. Ziegler and W. A. Lanford, “Effect of Cosmic Rays on Computer Memories,” *Science*, Vol. 206, Issue 4420, November 1979, pp. 776–788.

[4] J. L. Wert, E. Normand, D. L. Oberg, D. C. Underwood, M. Vallejo, C. Kouba, T. E. Page, and W. M. Perry, “Single Event Effects Test and Analysis Results From the Boeing Radiation Effects Laboratory (BREL),” proceedings of the IEEE Radiation Effects Data Workshop, Seattle, WA, July 2005.

[5] A. Taber and E. Normand, “Single Event Upset in Avionics,” *IEEE Transactions on Nuclear Science*, Vol. 40, Issue 2, April 1993, pp. 120–126.

[6] O. Niemitalo, “Yamaha YMF262 audio IC decapsulated,” Wikimedia Commons, May 2007. Digital image. Available: https://commons.wikimedia.org/wiki/File: Yamaha_YMF262_audio_IC_decapsulated.jpg.

[7] K. Zimmerman and D. Costello, “A Practical Approach to Line Current Differential Testing,” proceedings of the 66th Annual Conference for Protective Relay Engineers, College Station, TX, April 2013.

[8] J. Harrell, “The Importance of ECC Memory in Your Substation Computer,” July 2010. Available: https://selinc.com.

[9] K. Zimmerman and D. Costello, “How Disruptions in DC Power and Communications Circuits Can Affect Protection,” proceedings of the 68th Annual Conference for Protective Relay Engineers, College Station, TX, March 2015.

[10] R. Sandoval, C. A. Ventura Santana, H. J. Altuve Ferrer, R. A. Schwartz, D. A. Costello, D. A. Tziouvaras, and D. Sánchez Escobedo, “Using Fault Tree Analysis to Evaluate Protection Scheme Redundancy,” proceedings of the 37th Annual Western Protective Relay Conference, Spokane, WA, October 2010.

Biographies

Derrick Haas graduated from Texas A&M University with a BSEE. He worked as a distribution engineer for CenterPoint Energy in Houston, Texas, until 2006 when he joined Schweitzer Engineering Laboratories, Inc. Derrick has held several titles including field application engineer, senior application engineer, team lead, and his current role of regional technical manager. He is a senior member of the IEEE and involved in the IEEE Power System Relaying Committee.

Karl Zimmerman is a Principal Engineer with Schweitzer Engineering Laboratories in Saint Louis. He is an active member of the IEEE Power System Relaying Committee and Chairman of the Line Protection Subcommittee. Karl received his BSEE degree from the University of Illinois at Urbana-Champaign and is a registered Professional Engineer in the State of Wisconsin. Karl received the 2008 Walter A. Elmore Best Paper Award from the Georgia Tech Protective Relaying Conference and the best presentation award at the 2016 PowerTest Conference. He has authored over 40 technical papers and application guides on protective relaying.