Radiation and Fault Tolerance for Neuromorphic Computing

Rick Alena, Rev A. 5/18/2022

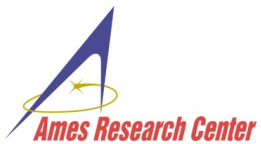

Space Radiation Summary LEO

- The solar wind emanating from the sun is a stream of energetic electrons and protons and ions that are the primary ionizing radiation contributions
- The Earth's magnetosphere significantly reduces radiation exposure and mutes the effects of solar flares out to 65K Km orbital altitude
- The radiation environment is highly dependent upon mission profile: for example geosynchronous orbits can see 50 times the radiation dose of a polar orbit due to reduction of atmospheric/magnetospheric shielding
- In LEO, the magnetosphere shields the Earth by creating trapped Van Allen radiation belts consisting of electrons and protons, which create a challenging radiation environment confined to specific regions
- Even for LEO, radiation dose can vary by a factor of 30 or more from a quiet sun to a solar flare
- Galactic Cosmic Rays (GCRs) are also present at very high energy, but at low rates
- Radiation effects on semiconductors are highly dependent upon particles species, energy and flux
- Therefore, radiation tolerance requirements will vary with mission profile and duration.

Space Radiation Summary (Cruise/Deep Space)

- Cruise or deep space is considered the region beyond 65K Km orbital altitude where the full force of the solar wind and solar flares create a much greater challenge for avionics
- For cruise phase, spacecraft see the full effect of the solar wind and solar flares – as a result, deep space missions require hardened avionics, not COTS
- Beyond the Earth and Moon, radiation environment varies significantly, such as in the vicinity of Jupiter due to trapped radiation
- Any material shields the avionics by blocking low energy particles, but also causes secondary particle emissions with energies that directly affect semiconductors – this effect is much greater in deep space
- Spacecraft avionics failures can occur in deep space due to a single high-energy solar flare

Annual Mission Radiation Dose (Rads Si) with 1mm AL Shielding

Mission	Solar Protons	RAD	RAD	Total Dose
	Total Protons	Rad/Day	Tot/Yr	Rads (Si)
LEO ISS SPENVIS	6735			6735
ISS-RAD		24	8760	8760
CRUISE SPENVIS	18125			18125
MSL-RAD Mars Cruise		45	16425	16470
MSL-RAD Mars Surface		22	8030	8052

Table Comparing LEO vs MARS Radiation Dose from SPENVIS Simulations and Mars Dosimeter

Space Radiation Effects on Avionics – Total Dose

- Total Integrated Dose (TID) is the tolerance in terms of delivered dose (Rads/Gys) before avionics performance characteristics are affected
  - TID degrades the gain and speed of semiconductor devices leading to functional failure.
  - TID can also result in complete failure due to Si crystal lattice damage.
  - TID changes threshold voltages for active devices
- Radiation-hard devices can tolerate 100K to 1 MRad of TID
- Radiation-tolerant devices can tolerate at least 30 KRad of TID
- Commercial devices generally can tolerate between 15 – 50 KRad TID
- Certain susceptible devices can fail at TID < 5 KRad
  - This is the level of radiation dose seen inside ISS over one year
- Therefore, EEE parts selection for TID is required for space missions

Space Radiation Effects on Avionics – Single Event Effects

- Single Event Effects (SEE) are caused by energetic particles creating unintended conductance paths in semiconductors leading to supply current spikes, device latch-up, program execution anomalies or data errors
- SEE rate is determined by incident radiation acting through the Linear Energy Transfer (LET) parameter of a given semiconductor device
  - Typical LET for latch-up should be >50 MeV-cm^2/mg for critical space avionics
  - Typical LET for upset should be >20 MeV-cm^2/mg
- Single Event Latch-up (SEL) can be either destructive or transient with different mitigation
  - Destructive latch-up destroys the semiconductor and results in functional failure
  - Transient latch-up can be cleared by power cycling the affected component
- Single Event Functional Interrupts (SEFI) occur when radiation results in a loss of function
- Single Event Upset (SEU) errors can result in anomalous program execution
  - Programs often crash, requiring reboot of the system
  - Programs can also produce transient execution errors, cleared by executing program again
- Single Event Upset (SEU) result in data errors in memories or processors
  - Data errors can be permanent or transient
  - Mitigation required for most memory types
- SEEs will occur in avionics and require appropriate mitigation techniques

Radiation Requirements and Fault Tolerance for Avionics

- Avionics shall meet the TID requirements for TID dose rate over entire mission duration
  - Integrate dose over mission profile and duration to determine TID requirement
  - Add shielding to reduce TID requirements
  - Select semiconductors with TID rating sufficient to meet mission requirements plus margin
- Avionics shall meet computational reliability requirements for the mission function performed
  - For human space flight-critical functions, double fault tolerance is required.
  - The same requirement is levied for aircraft
  - For life-critical functions, double fault tolerance is also required.
  - For mission-critical functions, single fault tolerance is usually required.
  - For non-critical functions, fault tolerance may not be required.

TID and Destructive SEL Mitigation Techniques

- Critical avionics must survive the radiation dose for a specific mission
  - Parts must be selected from components characterized for TID tolerance
  - For cruise/deep space missions, extra radiation tolerance is needed to cover solar flares
- Add shielding to increase TID rating
  - Shielding is VERY effective for reducing alpha and beta particle flux with low energy < 1MeV
  - Certain chips can be individually shielded with high-Z tantalum foil
- Shielding does not work for high energy particles due to creation of secondary particles
- Must eliminate destructive latch-up modes for all critical avionics
  - Must qualify critical parts for destructive latch-up modes under all environmental conditions
  - SOI fabrication uses dielectric insulation to increase TID and LET, increasing chip-level tolerance
  - Guard rings and substrate enhancement can produce similar benefits
  - De-rate components in terms of temperature range, timing margin and power dissipation
- The avionics hardware must tolerate the full range of radiation dose expected during the given mission without destructive latch-up failure modes

Radiation Fault Hardening and Detection Techniques

- Use radiation tolerant hardware platform adequate for desired mission environment
  - Use Error Detection and Correction (EDAC) in all memory
  - Triplicate certain memory such as Flash and FPGA circuits
  - All avionics components selected based on known radiation tolerance requirements
  - Significant derating of components and increased design margins are used
  - Incorporate power cycle functions to clear transient latch-ups and functional interrupts
- Add architectural features to DETECT transient SEU errors
  - Self-Checking Pair (SCP) as computational building blocks: disagreement triggers recovery
  - Watch-dog timer(s): initiate recovery if not inhibited by software within timeout interval
  - Voting architectures: multiple computational strings with software running same calculations, with the results compared in external logic; disagreement means that string has a fault
  - Memory scrubbing: background process ensures all memory is accessed and corrected within a given time interval

Radiation Fault Recovery Techniques

- Memory Recovery
  - Use Error Detection and Correction (EDAC) circuits for all memory – corrects single bit errors and can be expanded to multiple bit errors
  - Triplicate memory elements and use parallel voting architectures
  - Apply memory scrubbing working with EDAC to clean up entire array
- Architectural features to address SEU errors
  - Voting architectures: multiple computational strings with software running same calculations, with the results compared in external logic – requires 3 strings for single fault tolerance
  - Full power cycle: the entire computer system is shutdown and restarted cold
  - Reset – after fault is detected, the computer is reset and restarts the software
  - Rollback – after computational error is detected, the software redoes just the faulty computation
  - Checkpointing – multiple strings perform the same computations in parallel, but comparison is at specific points in software execution
  - Hot backup: independent computer running different software follows the main computer and takes over upon fault detection (also used for covering software faults)
  - Selective power cycling – after a fault is detected, only the affected component or subsystem is power-cycled, which can reduce recovery time significantly

Example Radiation Tolerance for Specific Components

Component	TID	LET SEL	LET SEU	Comments
Virtex6 FPGA	100 Krad	37 MeV-cm2/mg	< 1 for config mem	40 nm CMOS
NOR Flash	15 Krad	47.5	2.8	CMOS, tolerant flash
GaAs RF	1 Mrad	immune	N/A	GaAs
SRAM	50	80	< 1	SEU LET can be low
Regulators	300	68	N/A	SEGR
Kintex Rad-Tol FPGA	120	80	6	20 nm CMOS
DDR3		35	1.3	Micron

- Radiation tolerance of commercial components varies widely, so actual radiation testing is needed for missions in higher radiation environments
- However, SEU susceptibility is generally very high, particularly for memory chips or embedded memory elements

Total Integrated Dose (TID) Mitigation

- TID effects are mitigated by shielding, device hardening and architecture
- Lowering the TID using 5 mm of aluminum shielding is standard practice for space avionics by eliminating low energy particles
- Tantalum foil is often used for additional shielding using a high-Z material applied directly to chip packages
- Chips can be hardened against TID by enhanced fabrication techniques:
  - Silicon-on-Insulator (SOI) substrate
  - Buried guard rings to control stray charge
  - Enhanced conductivity substrates
- Circuits can be hardened against TID by design
  - Using larger devices
  - Increased design margins wrt gain, threshold voltages
  - Increased spacing between circuit elements
  - Slower operation with increased timing margins

Single Event Effects – SEL Mitigation

- SEE effects create anomalous currents in various circuits of the chip due to particle-induced charge paths
  - These anomalous currents are often the first detectable symptom of radiation effects
  - Charge paths consist of electrons and hole pairs created in the Silicon due to passage of a radiation particle
  - A high-energy particle can induce multiple secondary particle emissions from any material surrounding the active circuit – shielding, chip package and bulk silicon
- Anomalous charge paths lead to Single-Event-Latchup, either transient or permanent
  - Permanent SEL is caused by forward bias of a circuit to either the supply or substrate resulting in local overheating and permanent device damage
  - Permanent SEL fault modes should be avoided for any space avionics
  - Transient SEL is caused by similar conditions, but does not result in permanent damage
  - Transient SEL often requires a power cycle to clear the latchup condition
- SEL is countered by substrate enhancement, guard rings and circuit design
- SEL rate is determined by the Linear Energy Transfer (LET) parameter:
  - Typical LET for latch-up should be >50 MeV-cm^2/mg for critical space avionics
  - SEL can occur in relatively benign orbits due to GCRs and solar flares

Single Event Effects – SEFI Mitigation

- SEE effects can also alter the state of any register or flip-flop used on a chip
  - This effect is transient
- When a bit upset changes a state machine register, the state machine will malfunction
  - The state machine can produce anomalous output, but only for one cycle
  - The state machine can hang and requires a power cycle to clear
- Typical example is the controller for DRAM
  - Documented SEFI modes for certain DRAM chips
  - Controller hangs, preventing access to entire memory chip
  - Memory chip has to be power-cycled for recovery
  - Complex detection and power-cycling schemes can be employed
- SEFI can be countered by using DICE flip-flops or TMR for state machines
- SEFI rate is determined by the Linear Energy Transfer (LET) parameter:
  - Typical LET for SEFI should be >50 MeV-cm^2/mg for critical space avionics
  - SEFI can occur in relatively benign orbits due to GCRs and solar flares but at very low rates

Single Event Effects – Memory SEU Mitigation

- Generally, the most radiation sensitive component in computational systems is memory
  - There are many memory caches in computational systems and all are vulnerable to changes
  - DRAM requires the use of Error Detection and Correction (EDAC) in hardware for space use
  - EDAC needs to cover multiple bit flips, particularly for high-density DRAM
  - Processor cache memory is most vulnerable and can only be protected within the chip. This is seldom done
  - Configuration memory in FPGAs can have SEU LET < 1 Mev-cm2/mg, a critical vulnerability
- Non-volatile memory (NVM) is also sensitive to loss of data and bit flips
  - Memory element itself can be directly affected – floating gate of EEPROM Flash
  - Industry is developing radiation hard NVM like MRAM or memristors
- Interleaving of rows and columns helps to cover multiple-bit flips
- For soft memory, EDAC and TMR is often used – i.e. flash memory used for boot
- Radiation-hard memories remain an elusive goal and memory radiation tolerance often lowers rating of entire subsystem

Single Event Effects – Processors and FPGA Mitigation

- Processors can suffer SEUs in caches, registers or buffer memory
- For critical computational functions, basic self-checking pairs (SCP) are used as the computational building block
  - If the two do not agree on output, then entire SCP is considered as producing a fault
  - The SCP architecture can cover on-chip cache data errors – pair will not agree
- The configuration memory in FPGAs is implemented as SRAM and are the softest elements within the FPGA
  - Either use different FPGA technology for space or move to ASIC.
  - Configuration memory scrubbing is often used
  - Many FPGAs allow use of triple modular redundancy (TMR) for all logic circuits
- FPGAs also have internal Flash and SRAM memory blocks
  - Radiation hardness tends to follow similar types of chips
  - Internal blocks can be covered with EDAC or parity or TMR
- FPGAs can access external DRAM
  - Use EDAC circuitry and good DRAM parts

Architectural Fault Mitigation Techniques for Neuromorphic Computing (NMC)

- Start with a radiation tolerant hardware platform:
  - Memory used for defining weights of neuron connections (synapses) needs to be protected
    - Vulnerability may be highest during learning (programming) of the weights
  - Component selection based on mission requirements, with significant derating and increased margins
- Add architectural features to address transient SEU errors in NMC hardware
  - Parallel processing approach: (triplication on NMC chip)
    - Multiple neural strings with exact same weights and inputs should produce exactly the same outputs
    - Can compare results in external logic – requires 3 or more strings
    - Requires 3-4 times the hardware, but execution speed remains the same
  - Sequential processing approach (supervised by host)
    - Performing the same operation multiple times using the same weights and inputs should produce the same outputs
    - Compare at least 3 runs to determine correctness
    - Does not require more hardware, just 3X computational time
  - Memory scrubbing – background process ensure all NMC weight memory is accessed and corrected within a given time interval
  - Checkpointing – multiple strings perform the same computations in parallel, but comparison is at specific points in NMC execution
  - Rollback – after computational error is detected, the software redoes just the faulty computation
  - Power cycling and reset – when all else fails