FPGA-Based High Altitude Aircraft Systems

Heather Quinn*, Paul Graham
*hquinn@lanl.gov

ABSTRACT: This poster presents soft error rate information on FPGA-based aircraft systems, including proton and neutron data on the Xilinx Virtex line of FPGAs, commercial SRAM technology and Intel microprocessors.

The Earth's Magnetosphere

- Charged particles flow within the magnetic field lines
- Solar wind from the Sun affects the shape of the magnetic field lines
  - The Sun-facing magnetic fields are shaped by the "bow shock"
  - Most of the particles in the solar wind do not penetrate the magnetosphere, but flow around the magnetic field causing the non-Sun-facing side to become elongated

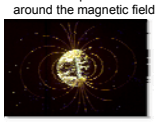

http://science.nasa.gov/ssl/pad/sppb/edu/magnetosphere/mag1.html

http://www.eskimo.com/~nanook/science/2007_07_01_archive.html

Cosmic Rays and the Magnetosphere

- Our initial understanding of cosmic rays predates our concepts of sub-atomic particles by 30 years
- Galactic flux:
  - "Debatable Origins": extra-galactic, or intra-galactic origins
  - Very energetic (10^{23} eV), very dense flux (100,000/m^2-s)
- Solar flux:
  - Not energetic, not likely to make it to sea level
  - During "active" periods
    - 10^6 increase in flux over quiet sun, 10x denser than galactic flux
    - Increased solar wind, distorts magnetosphere, increases "earth shielding", decreases galactic cosmic rays
  - During "quiet" periods
    - Very low fluxes
    - Magnetosphere more likely to allow galactic cosmic rays to penetrate into atmosphere

Solar Flares and Coronal Mass Ejections

- Coronal mass ejections (CME) release solar atmosphere
  - Often in conjunction with solar flares, but not necessarily
  - X-Ray, gamma-rays, electrons, protons, and heavy ions released at near speed of light
- CME/solar flares can filter to earth for a few hours after the event
  - Auroras
  - X-Ray-induced communication problems
  - Increased soft errors
- Every solar cycle seems to have one unusually large CME
  - Damage caused by the Halloween 2003 storms:
    - 28 satellites damaged, diverted airplanes, power failure in Sweden

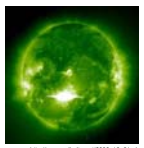

http://www.gallerita.net/2003_10_01.php

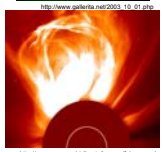

http://www.arm.ac.uk/climate/images/lebcme_sohoc2_big.gif

Cosmic Rays and the Atmosphere

- Cosmic rays that make it through the magnetosphere to the atmosphere cause a cascade of particles
  - Neutrons, protons, pions, and muons
- These particles can cause problems with electronics:
  - Memory upsets
  - Transient charge changes
  - Latch-up
  - Functional interrupts

J. F. Ziegler, "Terrestrial Cosmic Ray Intesities," IBM Journal of Research and Development, Vol 42 (1), 1998

Neutrons and Protons

- Flux dependent on longitude, latitude, altitude, geomagnetic rigidity, solar cycles, time of day, and time of the year
  - Radiation peaks at high altitudes and near poles
  - Some reduced affects at night or in winter months
- Flux sensitive to surroundings
  - Seventh transition (thermals) happens close to the electronics: either using nearby humans or building materials
  - Ship effect can increase flux by an order of magnitude
  - Can shield with water or concrete, but will need a lot of it
- Flux measurements often model-based (JEDEC89)
  - Seutest.com provides a flux calculator

J. F. Ziegler, "Terrestrial Cosmic Ray Intesities," IBM Journal of Research and Development, Vol 42 (1), 1998

- Since most of the neutron studies are done for land-based situations, protons are often ignored as they are 3-5% of the entire particle flux at sea level
  - JEDEC89 states that the proton flux is 5-20% of the neutron flux with 20% more likely in the peak neutron belts
  - IEC62396 states that the proton flux is 20-30% of the neutron flux up to 500 MeV and equivalent to neutron flux above 500 MeV
  - Will conservatively estimate the proton flux to be 25% of the neutron flux

Airplane Environments

Latitude	Sun Activity	Flux (n/cm^2-hr)	Flux (p/cm^2-hr)
Equator	Active	1,649-2,352	412-588
	50%	1,770-2,579	443-645
	Quiet	1,892-2,805	473-701
45°	Active	5,908-13,943	1,477-3,486
	50%	6,920-18,236	1,730-4,559
	Quiet	7,931-22,528	1,983-5,362
Polar	Active	14,194	3,549
	50%	18,639	4,660
	Quiet	23,083	5,771

- The airplane itself creates thermal neutrons when struck by protons
  - Dyer paper suggests the thermal flux could be on the order of 1/7th to 1/9th the flux of the fast neutrons in airplanes

Soft Error Rates (SER)

- SER_{device} = flux * σ_{device} = flux * σ_{bit} * (bits/device)
- Mean time to upset (MTTU) = 1/SER

LA-UR-08-05710

Neutron-Induced SEEs in Xilinx FPGAs

- SEE mechanisms:
  - Predominant mechanism is the SEU
  - SEU-induced SEFIs very, very rare in terrestrial-based systems
  - Xilinx has no SEL problems
  - SETs impossible to see through the SEUs
- Failure Modes:
  - SEUs in user memory (flip-flops, BlockRAM) can change intermediate data
  - SEUs in routing can either short or open your routing
  - SEUs in lookup tables (LUTs) can change logic values
  - SEUs in half-latches (Virtex-I) or power network (Virtex-II) can change logical constants for a region of the design
  - SEFIs in control logic can reprogram or de-program device
- Mitigation Methods:
  - Mask through redundancy methods
  - Repair through scrubbing

FPGA Estimates

- Neutron estimates were determined from Xilinx tests
  - Rosetta test: Atmospheric testing of a 100 device system
  - Accelerated radiation tests: using LANL's neutron beam
- Proton estimates were determined from LANL tests
  - Accelerated tests using 63.3 MeV protons at UC-Davis or 65 MeV protons at IUCF
- No thermal cross-section
  - No ^{10}B or reflow glass used in the manufacturing process
- No SEL cross-section

MTTU for the Virtex-I to Virtex-5

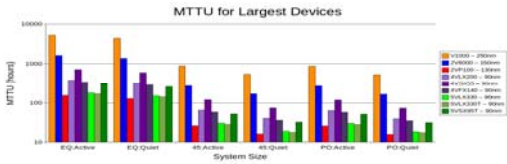

Neutron-Induced SEEs in Memories

- SEE mechanisms:
  - Predominant mechanism is SEU
    - Thermal neutrons continue to be a problem
  - SEL from neutrons becoming more commonplace
  - Micro-latching seen in some SRAM devices
- Failure modes:
  - Data corruption
- Mitigation Methods
  - Mask errors through redundancy, bit interleaving, Hamming codes
  - Repair errors through scrubbing
- Data used for calculations from Dyer TNS Oct 2004

MTTU for Neutrons: 512 Mb

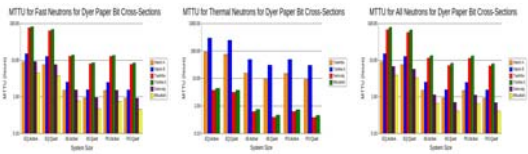

Other Factors to Consider….

- Part selection is important in the memory subsystem
- The layout of the memory is important
  - Triple-well SRAM layouts have higher SER rates due to multiple-cell upsets (Gasiot, TNS Dec 2007)
  - Trench-in-Channel SRAM layouts have very low SER rates (Ziegler, "SER – History, Trend, and Challenges: A Guide for Designing with Memory ICs)
- Using ECC-protected memory can help
  - Typical ECC protection is single error correct and double error detect
    - Protection from all but the MBUs SEUs

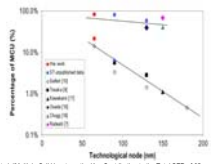

Gasiot et al, "Multiple Cell Upsets as the Key Contribution to the Total SER of 65 nm CMOS SRAMs and Its Dependence on Well Engineering, Transactions on Nuclear Science, 54(6), December 2007, 2468–2473

Feature Size	MBU% (educated guesses)
65 nm	12%
90 nm	5%
130 nm	2%
180 nm	0.5%

MTTU for All Neutrons: 512 Mb with ECC

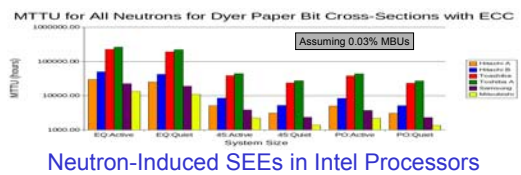

Neutron-Induced SEEs in Intel Processors

- SEE mechanisms:
  - The joy of owning your own fab: caches are rock hard
    - ECC + bit interleaving => no visible SEUs
  - SEUs in the registers and SETs in the logic are the predominant mechanisms
  - No known SEL
- Failure modes:
  - Data corruption by SEUs in registers
  - Data corruption by SETs in gates
  - Unrepeatable crashes
- Mitigation methods:
  - Mask errors through redundancy
  - Clever uses of duplicate computation and checkpoints using multiple cores

Microprocessor Estimates

- What we know
  - A lot of anecdotal evidence, no data
  - Intel publishes without numbers or units on their y-axis
  - They've told NASA that they are having problems with SETs in the combinatorial logic and SEUs in the register files
  - They've told us that their caches have no visible SEUs due to ECC and bit-interleaving, but they have stopped listing whether caches are ECC-protected
  - It's also clear that they are addressing an issue with cosmic rays, since they have become progressively more radiation-hardened over the years
  - The best number we have from them is a server quality microprocessor has a fail rate of once every 25 years, assuming that number is from sea level, that means a failure every 123-1210 hours at 60,000', depending on location.
- In absence of data what do you do?
  - Don't buy the low-end processors – there is no ECC protection on these devices
  - Get data of your own

System Level Estimates

- 1 5VLX330T FPGA, 512 Mb SRAM, 1 microprocessor
- 3 5VLX330T FPGA, 1 Gb SRAM, 3 microprocessor
- Use cases:
  - Case 1: Base System
  - Case 2: ECC-Protected SRAM
  - Case 3: ECC-Protected SRAM and 20% FPGA utilization

Case 1: Base Systems

- MTTU dominated by the SRAM
  - FPGA ~10x larger MTTU than SRAM
  - Processor ~10x larger than the FPGA
- System reliability dominated by the largest piece of unprotected memory
  - In this case, system reliability depends on protecting the SRAM

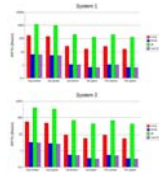

Case 2: ECC-Protected SRAM

- ECC protection increases the MTTU so that it is larger than the FPGA and the microprocessor
  - MTTU of the system is dependent on the FPGA

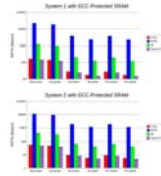

Case 3: ECC-Protected SRAM, 20% FPGA Util

- We have found that often times only 1-20% of the design is sensitive to errors
  - Part of this is due to design sensitivity
  - A large part is due to utilization
- High utilization of the FPGA is impossible
  - 80% of the bits are dedicated to routing
    - There is a multiplicity to the routing to give the tools a choice in how to layout a design
    - Once a route is chosen many of the other routing choices will not be used, nor have an affect on the ability to manifest an error
    - Therefore, very little of 80% of the bits are used
  - Depending on the design, might not use all of the logic bits
- Actual utilization and sensitivity to errors is unique to each design, each system

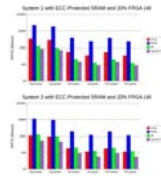

System Level Caveats

- These numbers represent a best attempt at raw estimates of total system error rates
  - The important numbers are the error rates of your actual system in it's natural environment
  - Testing your system will tighten up these numbers
  - Given the error in the particle flux, you might not know until you are in the air
- Some systems are inherently more inclined to faults than other systems
  - Need to find out how errors manifest in your system
  - What does it mean when your system becomes unavailable?
    - Do you miss a unique event?
    - Do you miss yet another chance to take the same picture?
  - How often is data overwritten? Or read?
  - Digital signal processing systems often only see an increase in noise (reasonable within limits)
  - Image processing errors might be bad pixels (reasonable) or a badly compressed image that cannot be restored (bad)
- Staying within the availability needs of the project is most important
  - A fault every 18 operating hours that takes one minute to repair translates to an availability of 0.99907 or 486 minutes of downtime each year. Are 3 9's acceptable?
  - When availability requirements will not be met mitigation will be needed, but how much is needed? Clearly, mitigating at the level at which we do spacecrafts might be overkill, but not many options for "light-weight" mitigation methods.

Summary and Conclusions

- Cosmic rays can cause errors to be introduced into airborne systems
  - The MTTU is directly related to the largest piece of unprotected memory
- Mask errors through error control coding and redundancy
  - Protect FPGAs with redundancy-based methods
  - Protect SRAM with Hamming codes
- Remove errors through scrubbing
- Modeling and fault injection tools available
- Test before the system goes live

This research has been supported through the DOE in the Deployable Adaptive Computing Systems project and the Sensor-Oriented Processing and Networking project; and supported through the DOD from the FPGA Mission Assurance Center.