Terrestrial-Based Radiation Upsets: A Cautionary Tale

Heather Quinn, Paul Graham

hquinn@lanl.gov

Overview

- **A Quick Introduction to Space Physics**
  - The Earth's Magnetosphere
  - Cosmic Rays
  - Solar Cycles, Solar Flares, and Coronal Mass Ejections
  - A Few of My Favorite Particles
  - Airplane Environments
- **Cosmic Rays and Electronics**
  - The History of Radiation-Induced Errors
  - Current reliability issues with Cosmic Rays and Modern Electronics
  - Radiation-Induced Failure Modes
- **Soft Error Estimates for FPGA-based systems**
  - FPGAs, memories, processors
  - What do these numbers mean?
- **Mitigation Methods**
  - Masking Errors
  - Repairing Errors
- **Testing Methods**
  - Modeling
  - Fault injection
  - Accelerated testing
- **Summary and Conclusions**

The Earth Magnetosphere

- **The earth is a giant magnet**
  - Charged particles will 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 actually penetrate the magnetosphere, but flow around the magnetic field causing the non-sun-facing side to become elongated

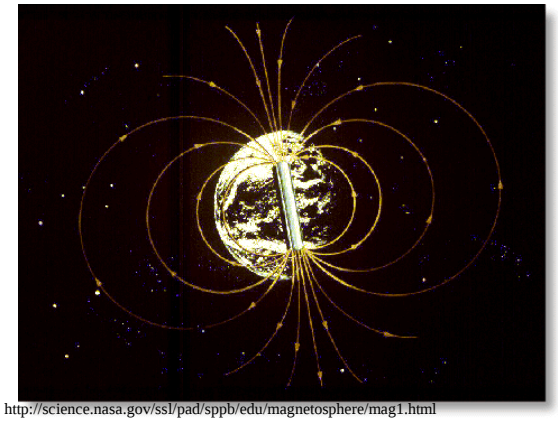

http://science.nasa.gov/ssl/pad/sppb/edu/magnetosphere/mag1.html

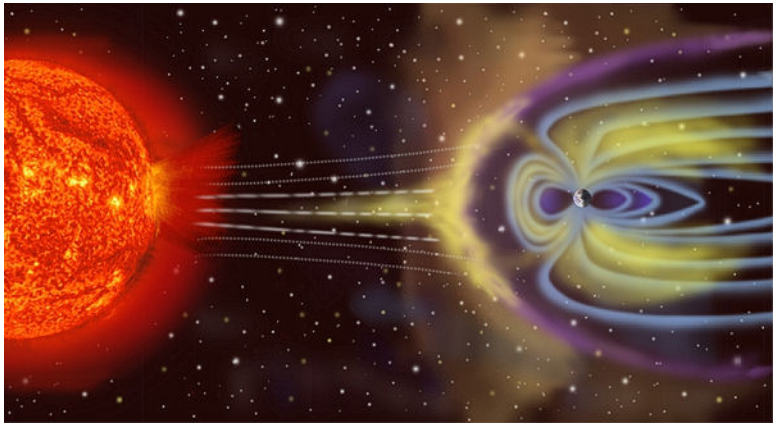

http://www.eskimo.com/~nanook/science/2007_07_01_archive.html

Cosmic Rays and the Magnetosphere

- **Our initial understanding of cosmic rays predates our concepts of sub-atomic particles**
  - Scientists knew charged particles were coming from the atmosphere, but didn't know what a neutron was for another 30 years
  - Original definitions were "particles that rain down from the sky but does not make me wet"
- **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

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

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**

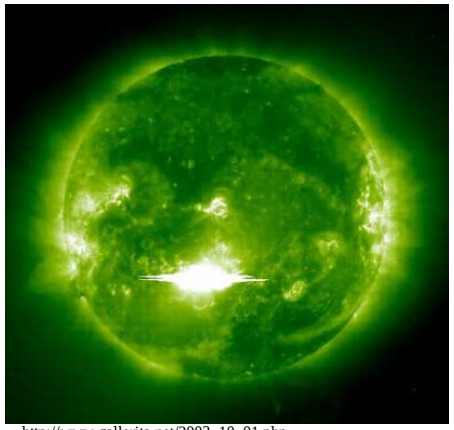

http://www.gallerita.net/2003_10_01.php

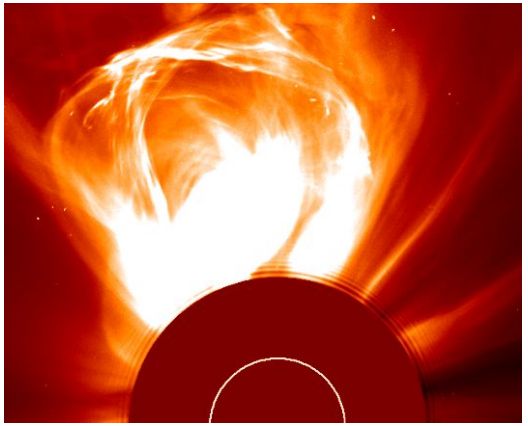

http://www.arm.ac.uk/climate/images/febcme_sohoc2_big.gif

The Halloween 2003 Coronal Mass Ejections

- **Three active solar spot groups (10484, 10486, 10488):**
  - All three were “remarkable in size and magnetic complexity”
  - One sun spot group (10486) was on over 13 times the size of the Earth and was the largest sun spot group observed since Nov 1990
- **17 CME ejections from mid-October to early November**
  - 12 events from 10486 alone
  - Three major events: the X17 on Oct 28, X10 on Oct 29, X28e on Nov 4
  - X28e event occurred while the GOES detector and was likely an X40 event

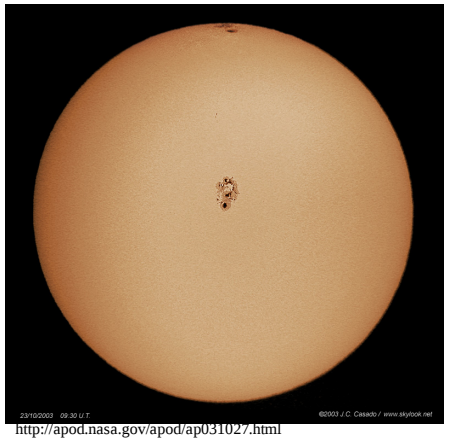

http://apod.nasa.gov/apod/ap031027.html

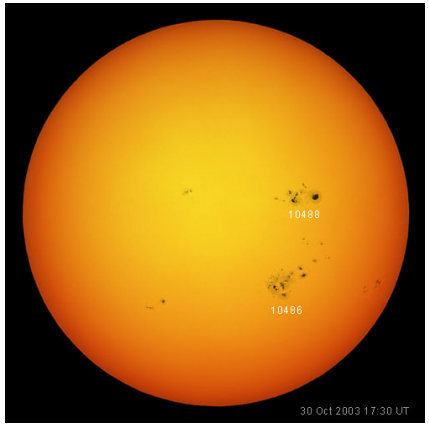

http://www.astropix.com/HTML/G_SUN/SS486488.HTM

T, Gombosi. “Comprehensive Solar-Terrestrial Environment Model for Space Weather Predictions. DoD MURI Project Report 2001-2004

Effects from The Halloween 2003 Coronal Mass Ejection

- **Auroras seen as low as CO, CA, NM, AZ**
- **Damage caused by the Halloween storms:**
  - 28 satellites (overt) damaged, 2 unrecoverably damaged
  - Diverted airplanes
  - Power failure in Sweden
- **…and two supercomputers came up in late October hoping to get onto the yearly Supercomputing list on Nov 10th**
  - At Los Alamos, the “Q” cluster had 26.1 errors a week and one unfortunate cluster topology
  - At Virginia Tech, System X architect joked they “felt like [they] had not only built the world's third fastest supercomputer, but also one of the world's best cosmic ray detectors.”
    - VT processed at night while in the magnetosphere tail
    - VT replaced all of their processors within the next 6 months

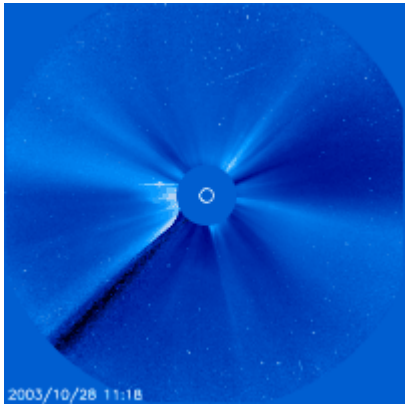

http://apod.nasa.gov/apod/ap031029.html

http://apod.nasa.gov/apod/image/0310/aurora031029b_westlake_full.jpg

T, Gombosi. “Comprehensive Solar-Terrestrial Environment Model for Space Weather Predictions. DoD MURI Project Report 2001-2004

Cosmic Rays and the Atmosphere

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

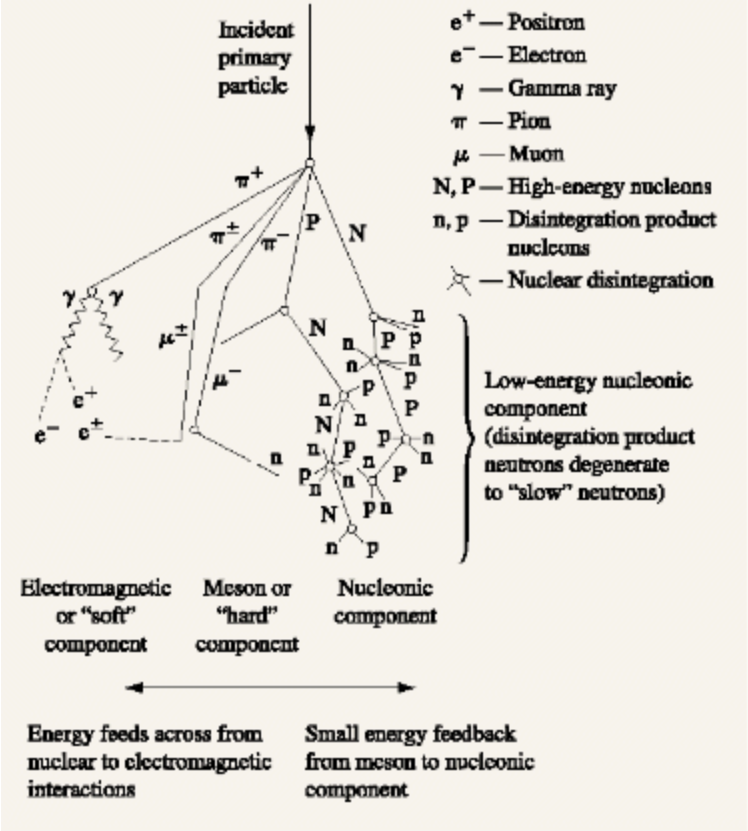

J. F. Ziegler, "Terrestrial Cosmic Ray Intesities," IBM Journal of Research and Development, Vol 42 (1), 1998

Muons and Pions

- **Pions:**
  - Unstable particles
  - Lifetime of ~26 ns
  - Very low flux at sea level (~450 pions/cm^2-year), more common at 50,000 ft
  - Rarely interact with Silicon, except for rare pion capture events
  - Cause 0.0003 fails/chip-year, considered minimum error rate possible for radiation-induced failures
- **Muons:**
  - Relativistic particles
  - Lifetime of ~2 μs
  - 100x more muons at sea level then any other sub-atomic particle
  - Very rarely interact with Silicon, except for muon capture events
  - Cause 0.0006 fails/chip-year

J. F. Ziegler. "SER – History, Trends, and Challenges: A Guide for Designing with Memory ICS", Cypress Press, 2004

Neutrons

- **Lifetime of 11-12 minutes**
- **Tend to "drill through" most substances**
  - Often just loses energy as it bounces off atoms
  - A direct strike with a Silicon atom releases a heavy ion, called "nuclear recoil reaction"
  - The heavy ion causes "soft errors"
  - Protons also cause a nuclear recoil and the sensitivities to both neutrons and protons often similar
- **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)**

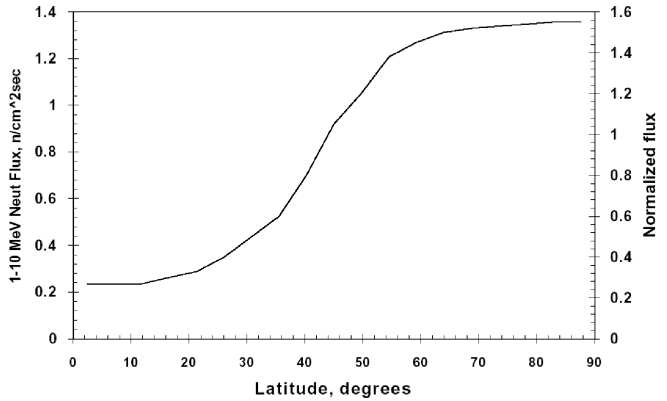

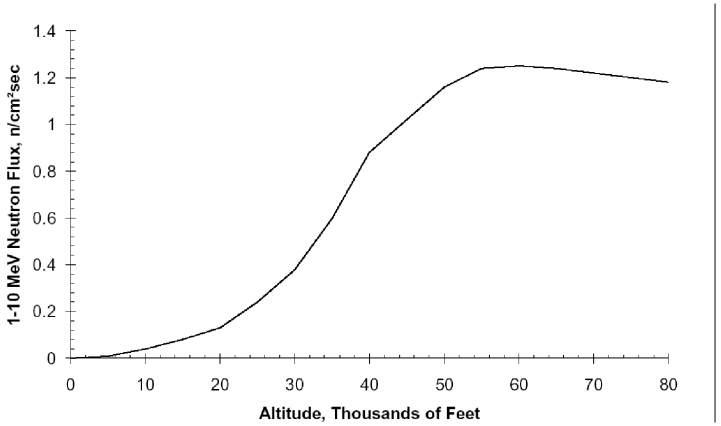

J. F. Ziegler, "Terrestrial Cosmic Ray Intesities," IBM Journal of Research and Development, Vol 42 (1), 1998

Protons

- 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
- JEDEC89 states that protons are 5-20% of the neutron flux with 20% more likely in the peak neutron belts
- IEC62396 states that protons are 20-30% of the neutron flux up to 500 MeV and equivalent to neutron flux above 500 MeV
- Not including proton calculations could introduce significant error into the calculation
- Including proton calculations is not simple either, since there is not a published proton spectrum
- Will conservatively estimate the proton flux to be 25% of the neutron flux with a monotonic spectrum so that 65 MeV bit cross-sections can be used

Airplane Environments

- **In the 30,000' to 60,000' range, flying through peak neutron belts**
  - Recent research shows that many of these neutrons are actually protons
  - Nuclear components are much more energetic than at sea level
- **In the 50,000'+ range, hadrons are more common**
  - Pions are much more common in this range
  - Could see up to 2 chips fail per year due to pion capture events
- **Disagreement in literature about the amount of peak flux**
  - Further complicated by airplane moving around
  - Estimate flux at seutest.com
    - Not particularly good at airplanes
    - Able to reproduce one estimate commonly used as peak flux

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

Airplane Environments (Cont)

- **The airplane itself creates neutrons**
  - Protons hitting the metal shell release low energetic neutrons (“thermal neutrons”)
  - Models of thermal neutrons are bad, detecting thermal neutrons difficult, no one knows how bad the thermal neutron flux is inside of airplanes
    - Most people hedge a number
    - 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
    - IEC62396 suggests that the thermal flux is 10x thermal flux outside of the airplane
  - Thermal neutrons are only a problem for some electronics:
    - Devices with ^{10}B will have increased error rates
    - Very rarely, you might have problems with SRAM latching

Dyer et al, “An experimental Study of Single-Event Effects Induced in Commercial SRAMs by Neutrons and Protons from Thermal Energies to 500 MeV, Transaction on Nuclear Science, 51(5), Oct 2004, 2817--2824

Overview

- **A Quick Introduction to Space Physics**
  - The Earth's Magnetosphere
  - Cosmic Rays
  - Solar Cycles, Solar Flares, and Coronal Mass Ejections
  - A Few of My Favorite Particles
  - Airplane Environments
- **Cosmic Rays and Electronics**
  - The History of Radiation-Induced Errors
  - Current reliability issues with Cosmic Rays and Modern Electronics
  - Radiation-Induced Failure Modes
- **Soft Error Estimates for FPGA-based systems**
  - FPGAs, memories, processors
  - What do these numbers mean?
- **Mitigation Methods**
  - Masking Errors
  - Repairing Errors
- **Testing Methods**
  - Modeling
  - Fault injection
  - Accelerated testing
- **Summary and Conclusions**

The History of Terrestrial-Based Radiation-Induced Faults

- **Historically, the problem was "an inside job"**
  - Alphas: radioactive bat guano, radioactive water
    - The fabrication process was altered to remove radioactive contaminates
  - Thermal neutrons: ^{10}B
    - Continue to be a problem with many CMOS devices to this day
- **Neutrons less of an issue, due to small size of memories and large clock frequencies**

J. F. Ziegler. "SER – History, Trends, and Challenges: A Guide for Designing with Memory ICS", Cypress Press, 2004

Modern Electronics

- **Physics: smaller is not better**
  - Smaller transistors are smaller targets, but easier to upset (Q_{crit})
  - Denser designs are easier to upset with multiple-bit upsets
- **System Design: increasing sensitivity (cross section)**
  - Microscopic: more complex components each generation
  - Macroscopic: larger systems each generation, memory-based devices (SRAM, DRAM, FPGAs) become larger, overall target size increase
- **System Location: peak neutron radiation levels**
  - Multiprocessor and multi-FPGA systems for airborne applications are more common

Soft Errors and System Reliability

- **Soft errors are often undetected, unmitigated**
- **For airplanes, 75% of all unrepeatable system errors are caused by soft errors**
- **For large-scale, reliable systems unmitigated soft errors are disastrous:**
  - Sun Microsystems received bad press for soft error failures in their high end servers in 2000
  - Intel processors verging on radiation-hardened electronics
- **Meanwhile, the last few generations of electrical engineers and computer scientists were educated on the basis of IEEE standards and von Neumann continues to be the standard in fault tolerant computing**

Terrestrial-Based Radiation-Induced Failure Modes

- **Predominantly single-event effects (SEE):**
  - Traditionally, most terrestrial systems are concerned with single-event latch-up, single-event transients, single-event upsets, and single-event functional interrupts
  - Most SEEs are not destructive, except single-event latch-up, single-event gate rupture, and single-event burn out
  - Rest of the SEEs are non-destructive, but can make fault tolerant computing difficult
  - Single-event functional interrupts can cause overall system failure until the device is reset
  - Experimentally measured at cyclotrons to determine sensitivity/volume
    - Either per-bit or per-device sizes

Single-event latch-up (SEL)

- **Traditional reliability issue with CMOS due to parasitic transistors caused by well/substrate contact**
  - Once turned on, current increases rapidly and destroys the part
  - Radiation is another avenue for turning on the parasitic transistor
- **Military/aerospace parts often have an epitaxial layer to prevent SEL, by localizing charge collection**
- **Measured by per-device size**

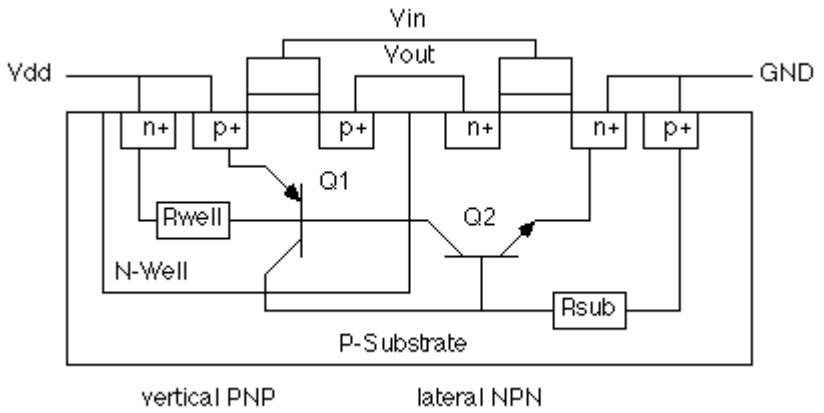

http://www.ece.drexel.edu/courses/ECE-E431/latch-up/latch-up.html

Single-event transients (Transients or SET)

- **Radiation-induced charge temporarily changes the value of gate**
  - No way to tell the difference from a real signal and a transient-affected signal
  - Transients in logic gates are a problem if latched, causes data corruption
  - Transients in the clock or reset trees can cause much more global issues
- **Decreasing clock frequencies make it easier to latch a transient: transient pulse and clock signal are roughly the same**
- **Measured per-bit**

Critical Pulse Width for Unattenuated Propagation

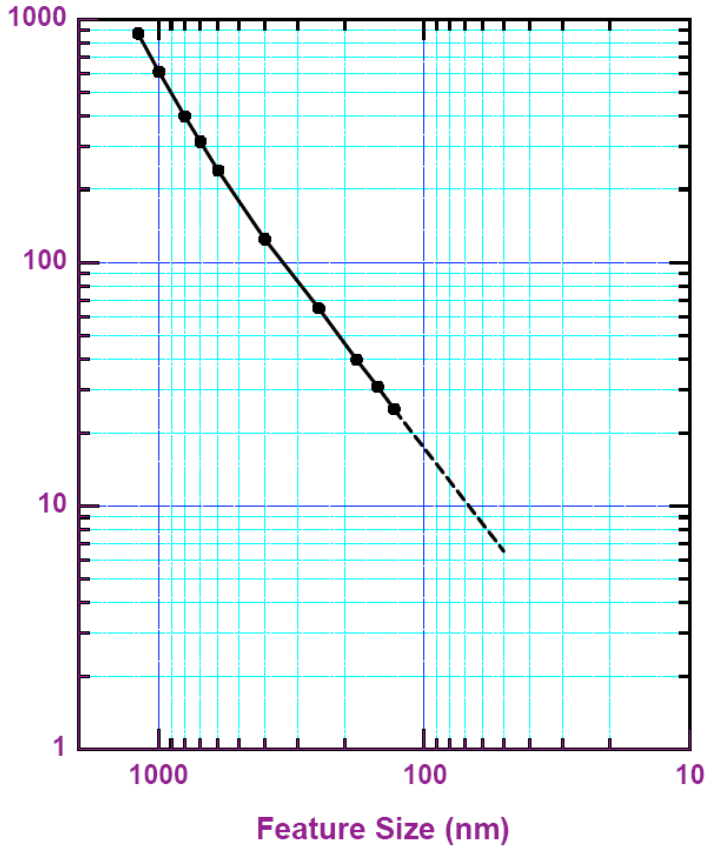

Mavis, "Single-Event Transient Phenomena: Challenges and Solutions." MRQW, 2002.

Single-Event Upsets (upsets or SEUs)

- **Cause bit flips in memory-based**
  - Data changes from 1→0 or 0→1
  - In some parts single-bit upsets (SBUs) are as common as multiple-bit upsets (MBUs)
- **Strongly affected by feature size:**
  - Smaller feature size means smaller targets, smaller Q_{crit}, more MBUs
  - Even with a decrease in per-bit cross-section, often see an increase in per-device cross-section increase
- **Measured per-bit or per-device**

Single-Event Functional Interrupts (SEFIs)

- Device will not operate functionally until reset
- Often caused by an SET or SEU in control logic for the device
- Measured per-device

Neutron-Induced SEEs in FPGA-Based Systems

- **FPGA systems are not exempt from SEE:**
  - The entire system (FPGAs, microprocessors, memory) is sensitive
  - SEE problems dependent on the subsystem
    - Which SEE manifest
    - How errors manifest
    - How to repair/mitigate errors

Neutron-Induced SEEs in Xilinx FPGAs

- **SEE mechanisms:**
  - Predominant mechanism is the SEU
  - SEU-induced SEFIs very, very rare
  - 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 (VI) or power network (VII) 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

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

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

Overview

- **A Quick Introduction to Space Physics**
  - The Earth's Magnetosphere
  - Cosmic Rays
  - A Few of My Favorite Particles
  - Solar Cycles, Solar Flares, and Coronal Mass Ejections
- **Cosmic Rays and Electronics**
  - The History of Radiation-Induced Errors
  - Current reliability issues with Cosmic Rays and Modern Electronics
  - Radiation-Induced Failure Modes
- **Soft Error Estimates for FPGA-based systems**
  - FPGAs, memories, processors
  - What do these numbers mean?
- **Mitigation Methods**
  - Masking Errors
  - Repairing Errors
- **Testing Methods**
  - Modeling
  - Fault injection
  - Accelerated testing
- **Summary and Conclusions**

Soft Error Rates (SER)

- **SER**

$$SER_{device} = flux * \sigma_{device} = flux * \sigma_{bit} * (bits/device)$$

- **Mean time to upset (MTTU)**

$$MTTU = 1/SER$$

- **Scaling soft error rates**
  - Scaling for locations

$$SER_2 = (flux_{loc2}/flux_{loc1})*SER_1$$

  - Scaling for system size

$$SER_2 = (\sigma_{device1} / \sigma_{device2})*SER_1$$

  - Scaling for both

$$SER_2 = (flux_{loc2}/flux_{loc1})*(\sigma_{device1} / \sigma_{device2})*SER_1$$

Caveats: Error in Soft Error Rates (SER)

- **SER**

$$SER_{device} = flux * \sigma_{device}$$

- As discussed, hard to get accurate number
- Total error ~2-10x

- Experimentally measured: errror in dosimetry of the cyclotron, test fixture, low number of upsets counted.
- Total error ~2x

Caveats: Interarrival Times

- **Radiation-matter interactions are Poisson random processes**
  - MTTU tells you the time frame that the upset will occur in
  - Poisson tells you that the upset will arrive at random in that time frame
- **The closer you get to the end of the time frame, the more likely that the upset has occurred**
- **The actual arrival time could be any time within that time frame…or not all…or twice in that time frame**
  - P(0 upsets) = 37%
  - P(1 upset) = 37%
  - P(2 upsets) = 18%
  - P(3 upsets) = 6%
  - P(>=2 upsets) = 26%
- **The clock does not reset when you land the plane**
  - If the time frame is longer than the mission length then it only means that it might take several missions to manifest an upset not that you are SEU-free
  - On the other hand, if the time frame is short than the mission length, landing the plane and unbiasing the system is one way to clear out accumulated errors

Caveats: General Trends in Soft Error Rates

- **We are not here to beat up on vendors**
  - Increasing system size or flux unavoidably increases SER
- **With current trends in system design, soft errors will become more prevalent:**
  - Microscopic: design components to be less likely to upset
  - Macroscopic: design large systems to be error resistant
- **Research and development <u>now</u> while the problem is still manageable**
  - Determine the scope
  - Find low impact mitigation methods
  - Change our system design methods
  - Intel recently headed this way with it's "Cosmic Ray Detector"

FPGA Estimates

- **Neutron estimates were determined from Xilinx tests (earlier data) and joint Xilinx/LANL tests**
  - Rosetta test: Atmospheric and accelerator 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 VI-V6 Normalized to V1000 Bitstream

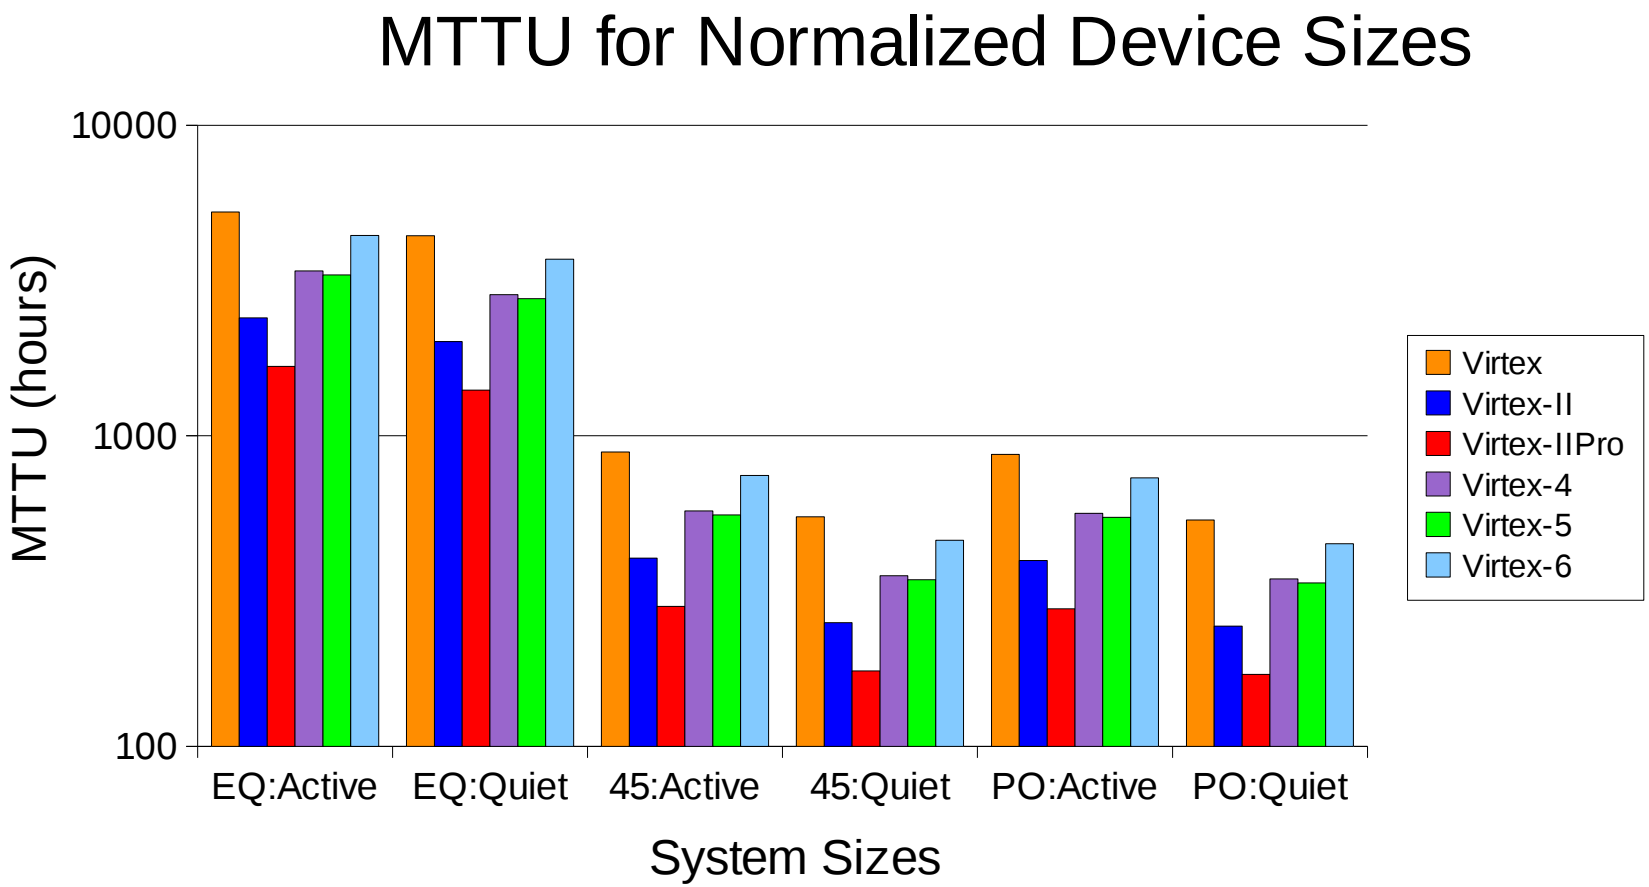

Virtex-I 1000 Bitstream:CFG: ~6Mb, BRAM: ~130Kb

Relative Bit Cross-Sections for V-I to Virtex-6

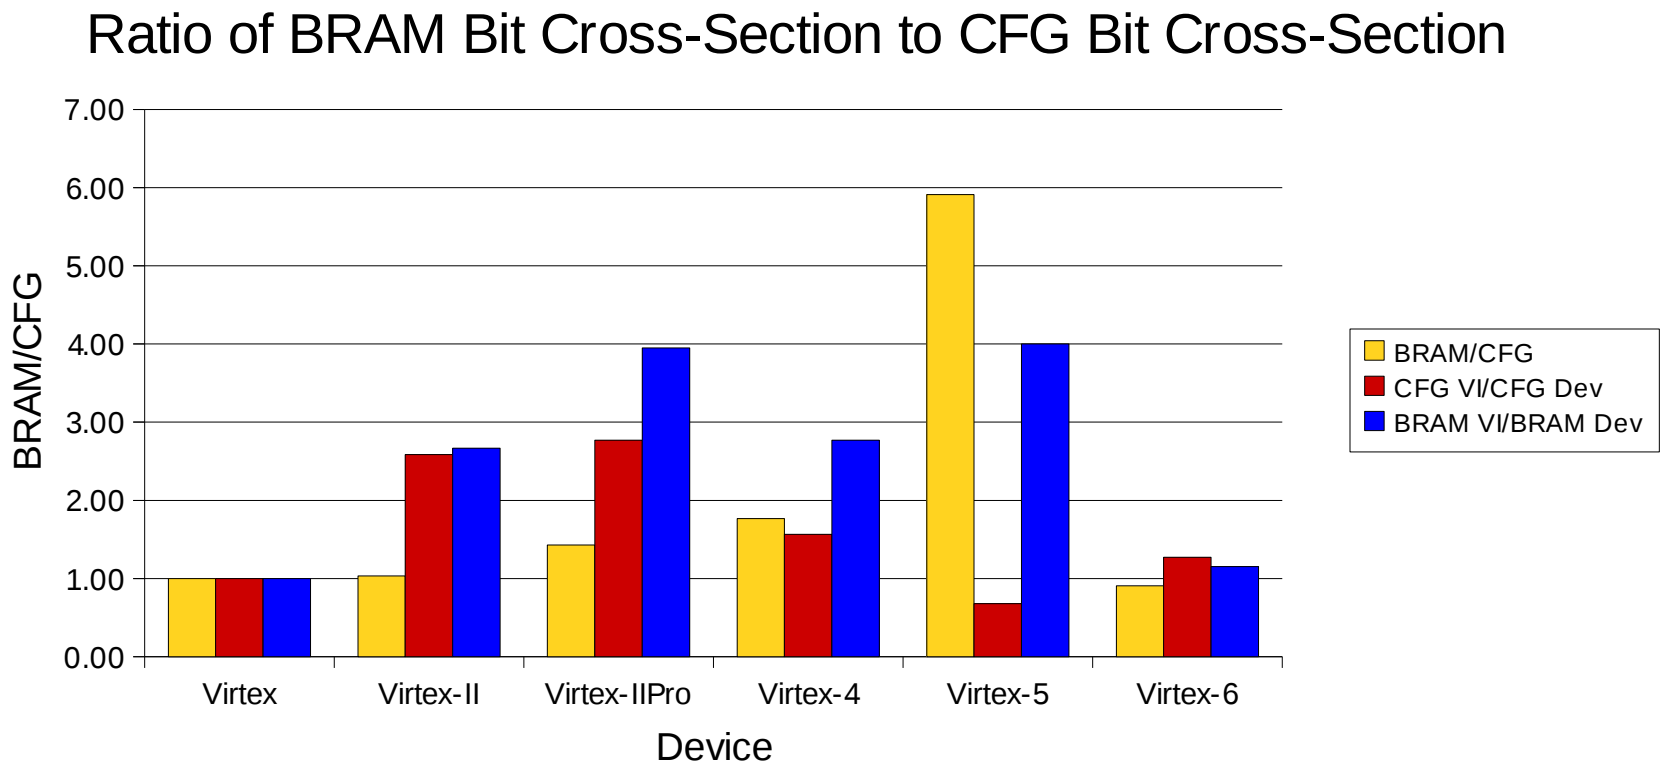

Comparison of Largest Device Sizes from V-I to V-6

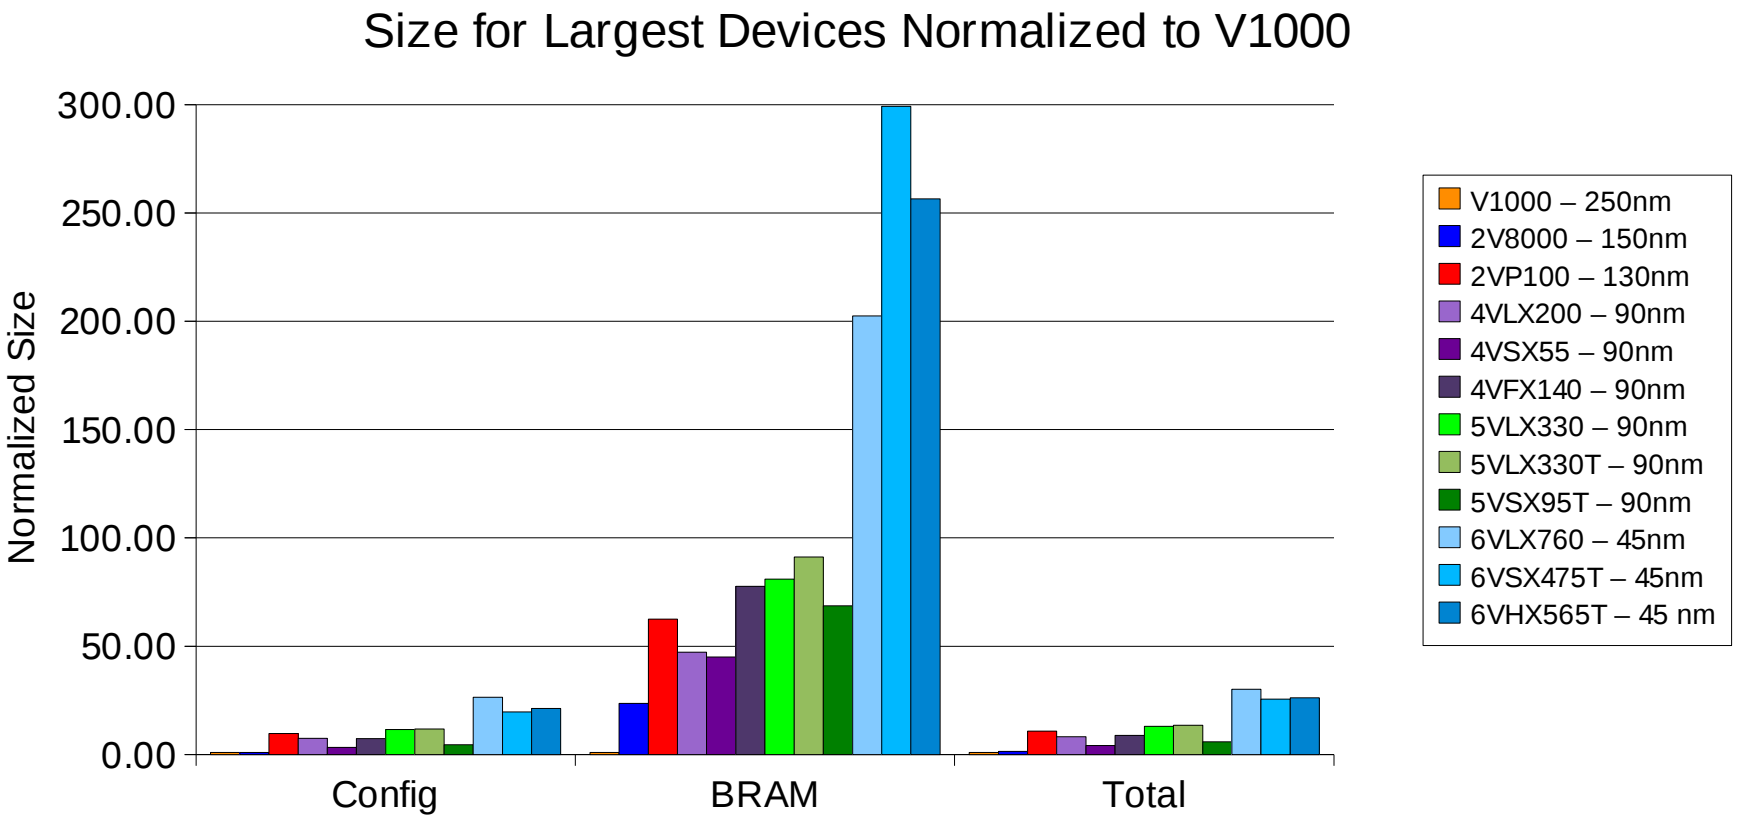

MTTU for Largest Devices for V-I to V-6

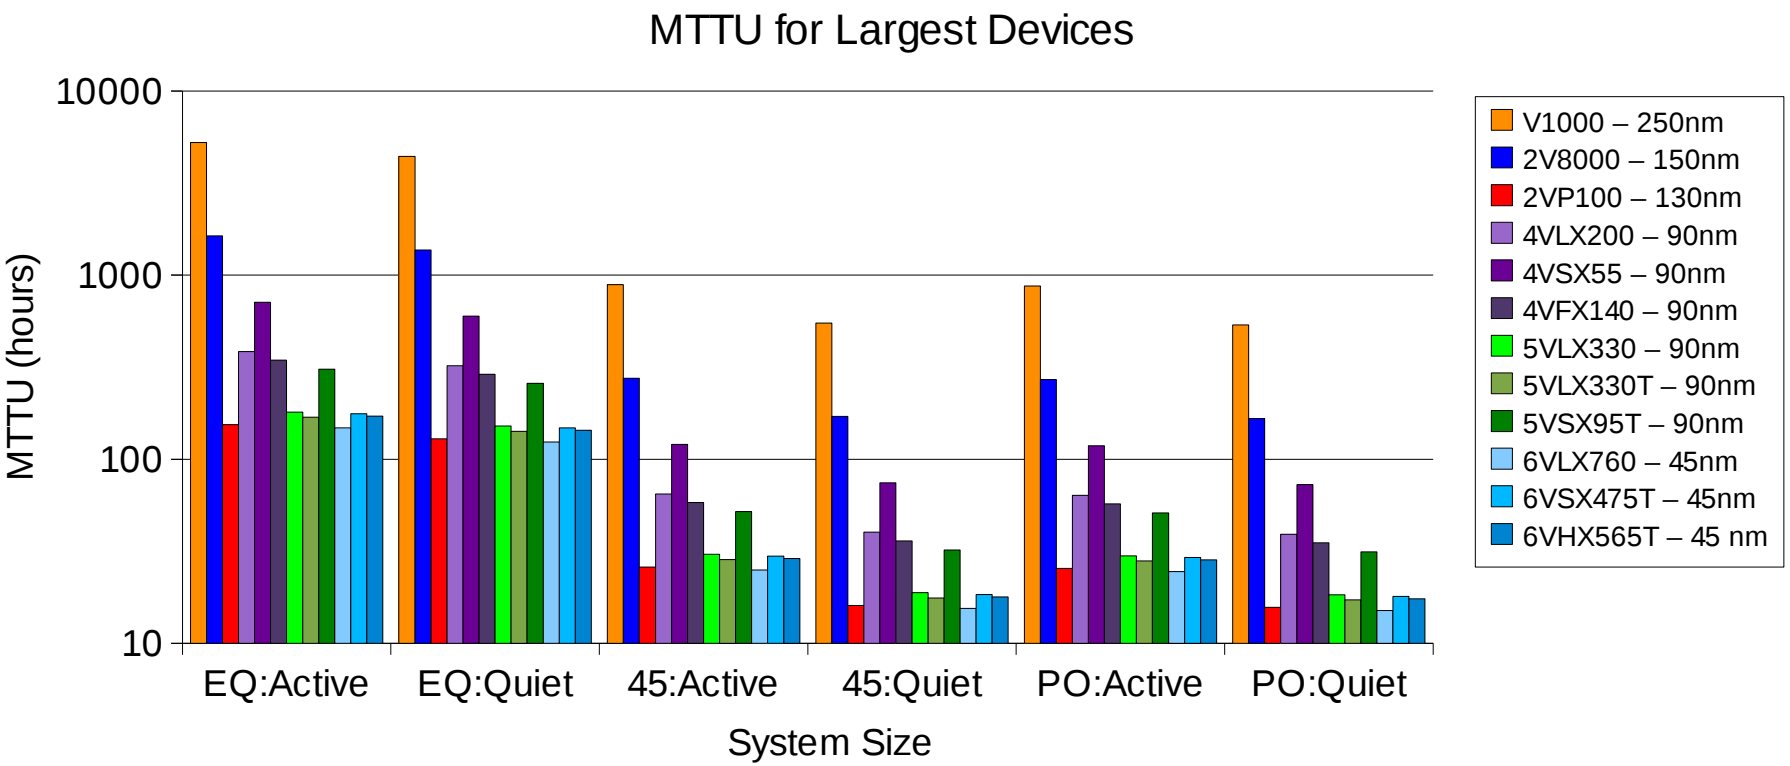

SRAM Memory Estimates

- **Neutron estimates were determined from several NSREC papers**
  - Dyer, TNS 2004
  - Granlund, TNS 2006
  - Armani, TNS 2004
- **No proton estimates, used 1.25 the flux for the calculation**
- **Some devices have a thermal bit cross-section**
  - In those cases the some times the thermal bit cross-section is larger than the fast neutron bit cross-section
- **One device latched**
  - Hitachi B has an SEL cross-section of 2.6E-9 cm^2 on a 174 MeV peak-flux neutron beam

Parts from Dyer Paper	Feature Size	Thermal Neutron Bit Cross-Section (cm^2/bit)	Fast Neutron Bit Cross-Section (cm^2/bit)	Ratio of Thermal/Fast Bit Cross-Section
Hitachi A	500nm	0	7.06E-14	-
Hitachi B	350 nm	0	4.24E-14	-
Toshiba	500 nm	8.7E-15	8.35E-15	1.04
Toshiba A	400 nm	2.7E-15	7.75E-15	0.35
Mitsubishi	400 nm	2.1E-13	7E-14	3.00
Samsung	400 nm	1.8E-13	1.41E-13	1.28

Dyer et al, "An experimental Study of Single-Event Effects Induced in Commercial SRAMs by Neutrons and Protons from Thermal Energies to 500 MeV, Transaction on Nuclear Science, 51(5), Oct 2004, 2817--2824

More Parts….

Parts from Granlund Paper	Feature Size	Thermal Neutron Bit Cross-Section (cm^2/bit)	Fast Neutron Bit Cross-Section (cm^2/bit)	Ratio of Thermal/Fast Bit Cross-Section
Cypress 1	250 nm	0	9.13E-15	-
Cypress 2	250 nm	1.85E-13	3.13E-14	5.91
Samsung 3	250 nm	1.82E-14	5.27E-15	3.45
Mitsubishi 4	180 nm	2.38E-13	3.89E-15	61.18
Cypress 5	160 nm	0	4.07E-14	-
Cypress 6	160 nm	0	3.64E-14	-
Cypress 7	160 nm	0	3.21E-14	-
Cypress 8	160 nm	0	4.53E-14	-

More Parts….

Parts from Granlund Paper	Feature Size	Thermal Neutron Bit Cross-Section (cm^2/bit)	Fast Neutron Bit Cross-Section (cm^2/bit)	Ratio of Thermal/Fast Bit Cross-Section
Cypress 9	160 nm	0	3.96E-14	-
Cypress 10	130 nm	0	3.80E-14	-
Cypress 11	130 nm	3.07E-14	3.95E-14	0.78
Hitachi 12	130 nm	0	2.44E-14	-
Hitachi 13	130 nm	0	3.00E-14	-
Samsung 14	130 nm	1.49E-14	4.54E-15	3.28

More Parts….

Parts from Armani Paper	Feature Size	Thermal Neutron Bit Cross-Section (cm^2/bit)	Fast Neutron Bit Cross-Section (cm^2/bit)	Ratio of Thermal/Fast Bit Cross-Section
Samsung	130 nm	1.9E-16	4.8E-16	0.40
Toshiba	220 nm	7.1E-14	Not measured	-
Hitachi	180 nm	2.8E-16	1.6E-15	0.18
Nanoamp	180 nm	6.9E-17	1.1E-16	0.63

SRAM Bit Cross-Section by Feature Size

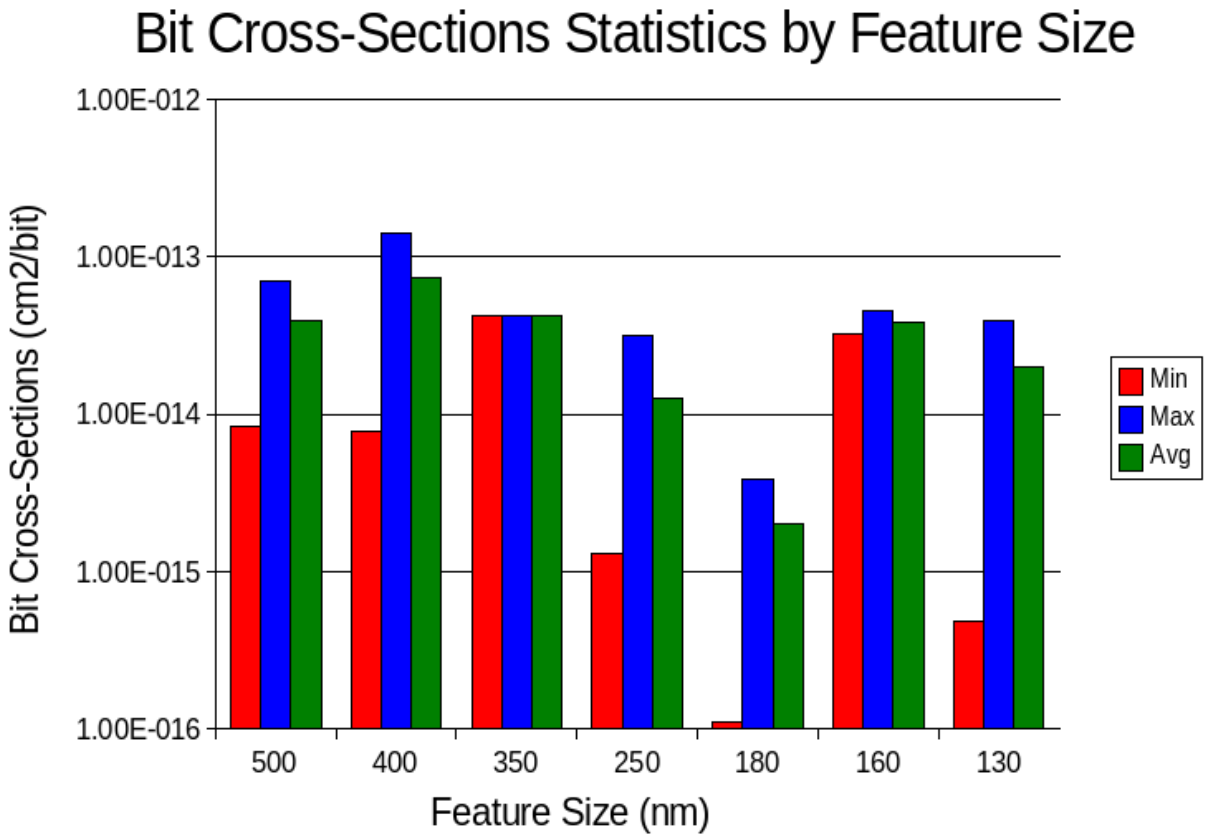

SRAM MTTU for Fast Neutrons: 512 Mb

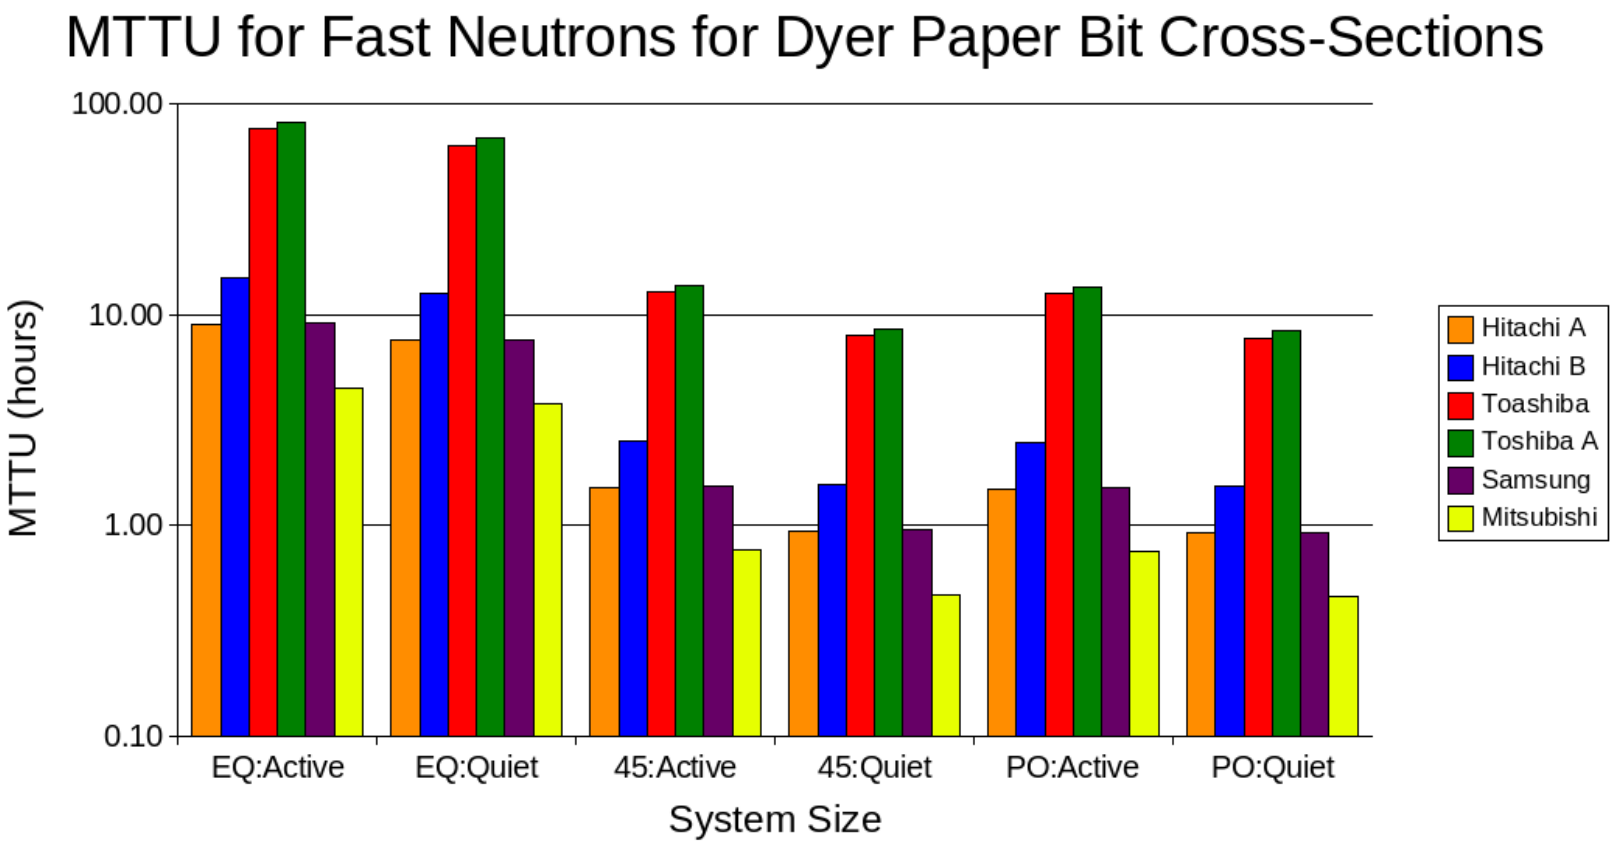

SRAM MTTU for Thermal Neutrons: 512 Mb

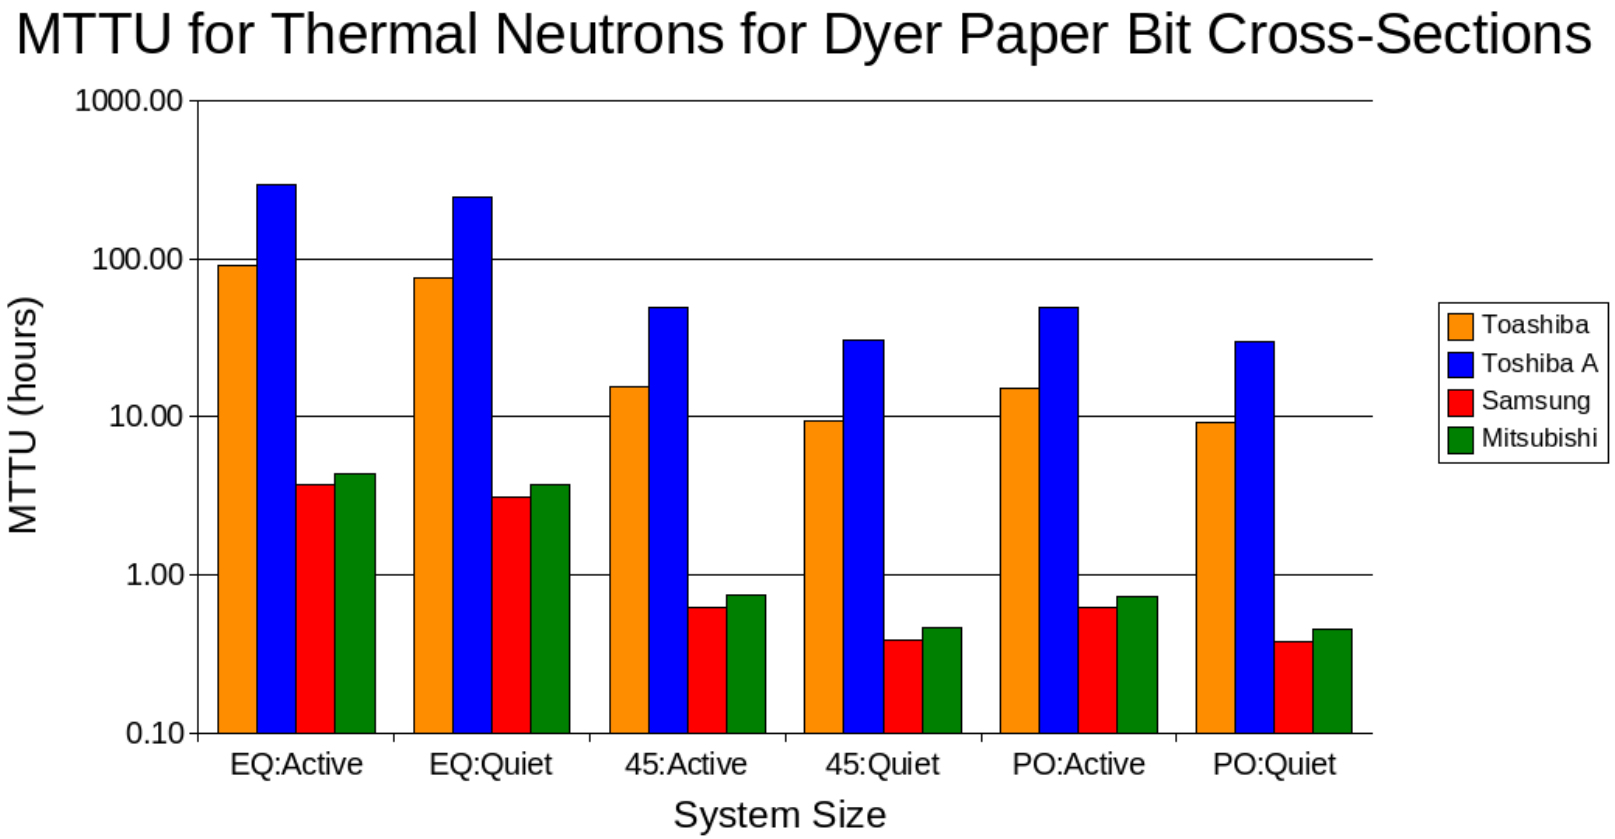

SRAM MTTU for All Neutrons: 512 Mb

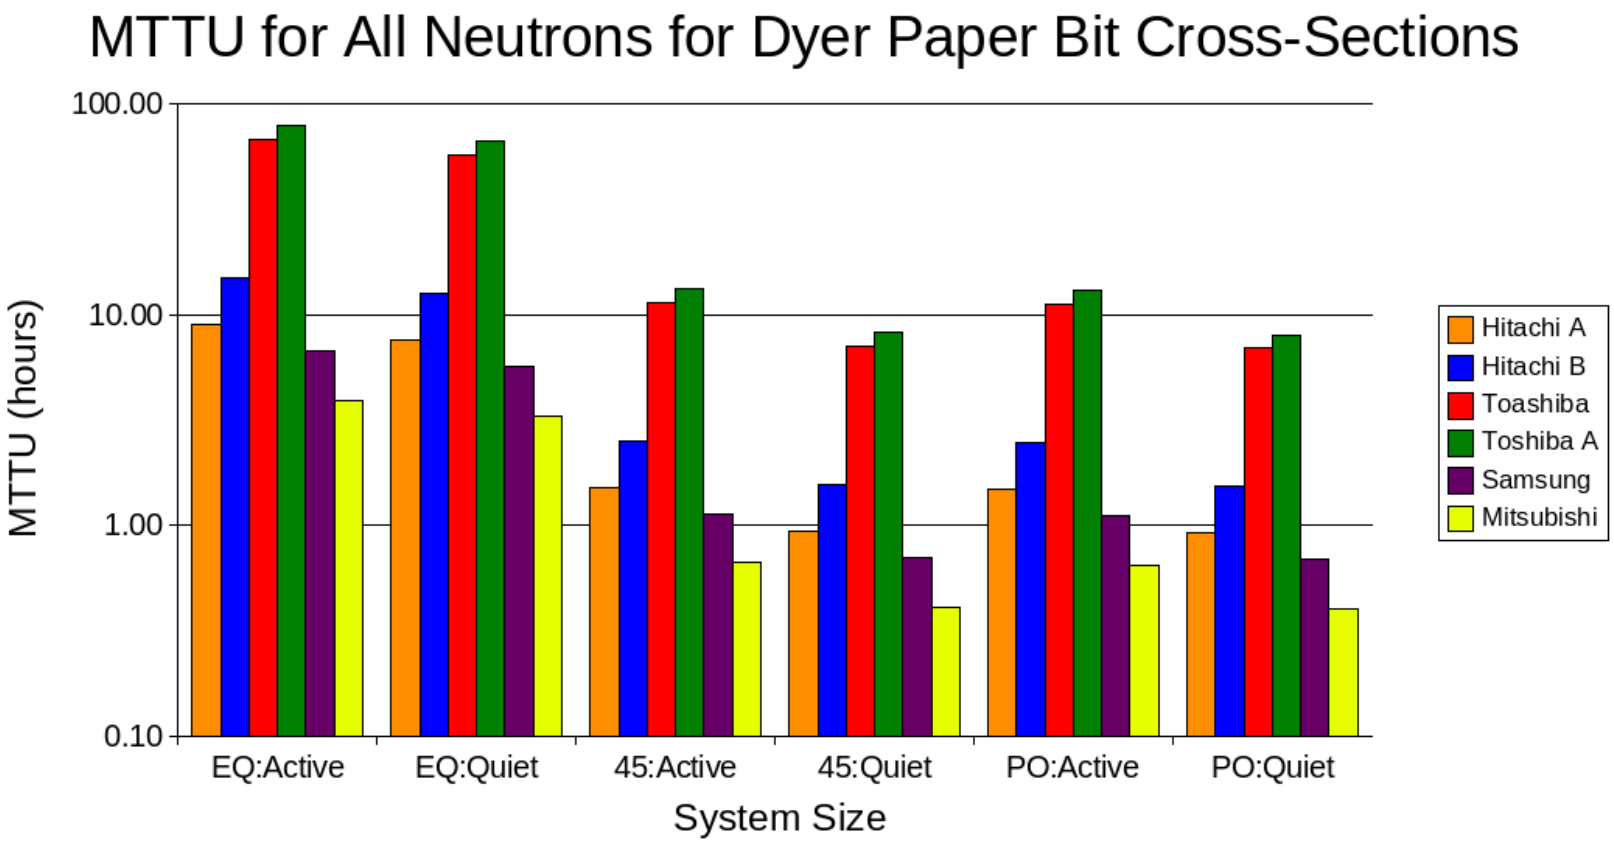

SRAM Bit Cross-Section and Voltage

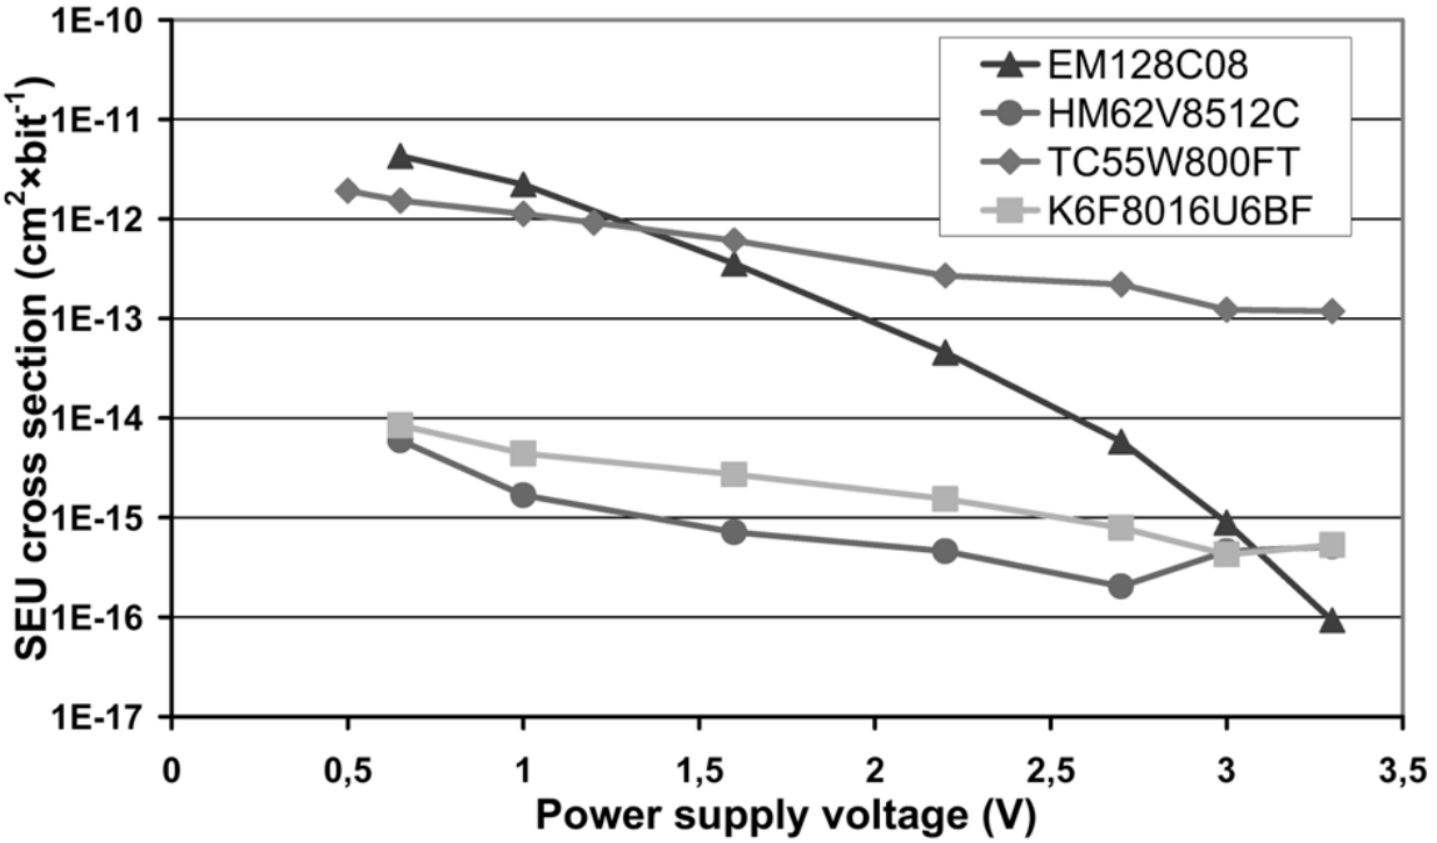

Armani et al, "Low-Energy Neutron Sensitivity of Recent Generation SRAMs," Transactions on Nuclear Science, 51(5), October 2004, 2811--2816

Other Factors to Consider with SRAM….

- **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**
  - FPGAs read the memory as “raw” signals, will need to decode the data on the input side and encode on the output side
  - Typical ECC protect is single error correct and double error detect
    - Protection from all but the MBUs

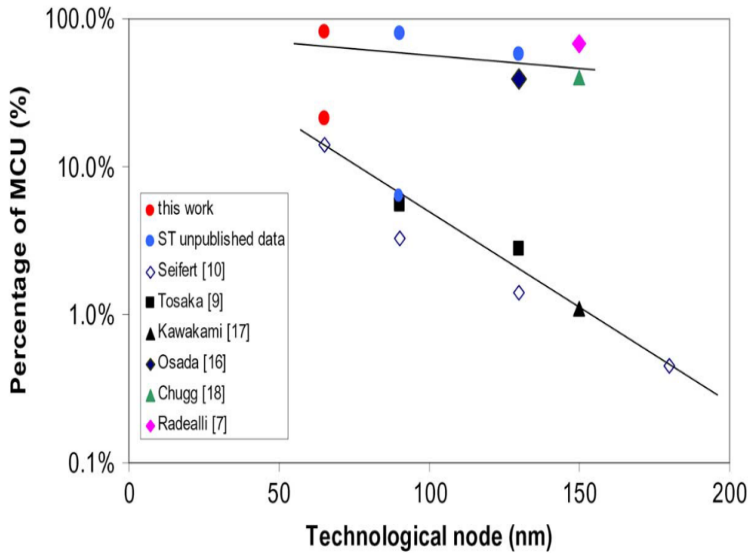

Gasiot et al, “Multiple Cell Upsets as the Key Contribution to the Total SER of 65 nm CMOS SRAMs and Its Dependence on Well Engineering, Transactions on Nuclear Science, 54(6), December 2007, 2468--2473

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

SRAM MTTU for All Neutrons: 512 Mb with ECC

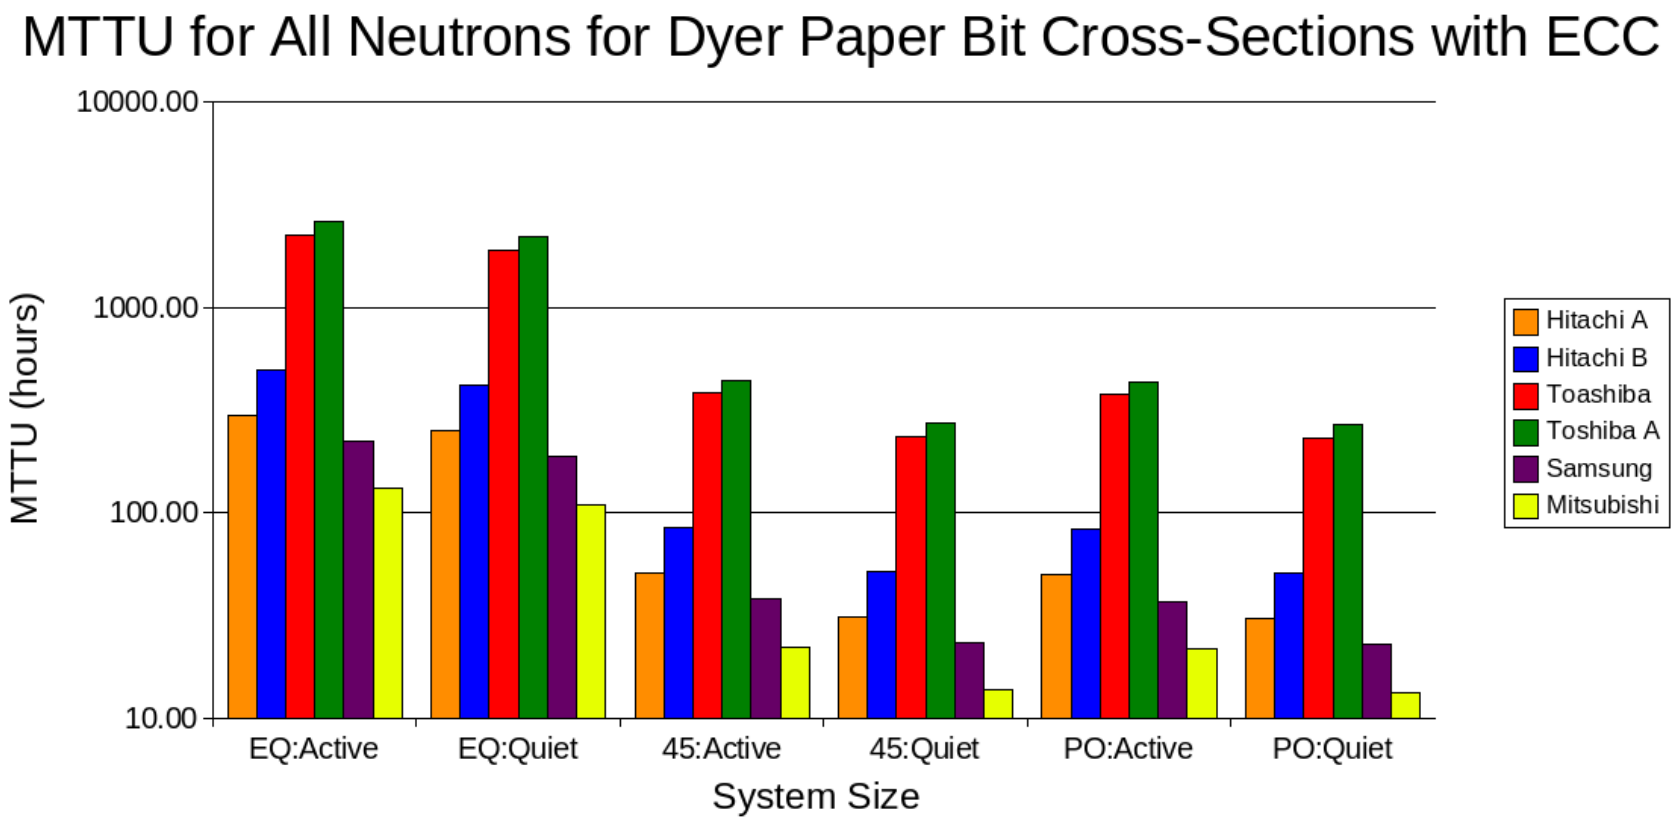

Assuming 0.03% MBUs

SEUs in DRAM

- **SEE mechanisms:**
  - Predominant mechanisms are SEUs in the memory array and SEU/SETs in the control circuitry that cause a SEFI mode of burst errors
  - SEL or high-current events an ongoing problem
  - “Stuck bits” are not uncommon – a single bit or a page will become stuck at some value for some number of seconds, which will eventually relax back to a writeable state
- **Failure modes:**
  - Data corruption
- **Mitigation Methods**
  - Mask errors through redundancy, bit interleaving, Hamming codes for bitflips
  - RAID striping might be necessary to mitigate SEFI modes
  - Repair errors through scrubbing

SDRAM SEU Bit Cross-Sections

Sample	SEU Bit Cross-Section (cm^2/bit)	SEFI Bit Cross-Section (cm^2/device)
SDRAM1	$2.14x10^{-20}$	$4.76x10^{-12}$
SDRAM2	$2.15x10^{-20}$	$1.62x10^{-10}$
SDRAM3	$7.54x10^{-20}$	$7.71x10^{-12}$
SDRAM7	$7.23x10^{-20}$	$1.79x10^{-11}$
SDRAM8	$1.72x10^{-20}$	$6.94x10^{-11}$
SDRAM9	(0, $2.32x10^{-19}$)	$1.26x10^{-11}$
SDRAMA	$4.43x10^{-20}$	(0, $2.20x10^{-11}$)

SDRAM EDAC Failures

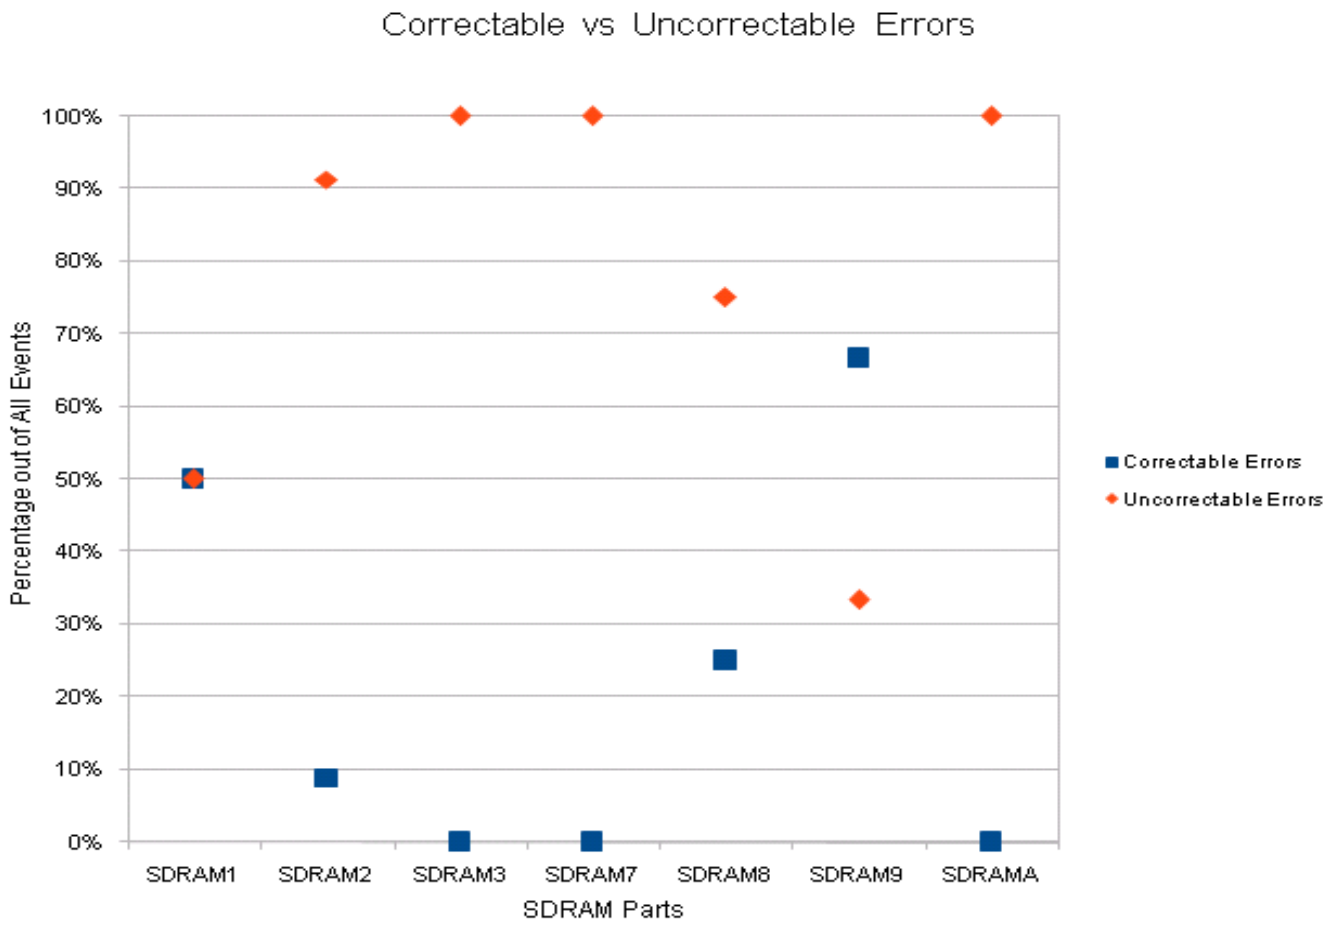

Microprocessor Estimates

- **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 rad-hard 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 fail ever 123-1210 hours at 60,000'.
- **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 RAM, 1 microprocessor**
- **3 5VLX330T FPGA, 512 Mb RAM, 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

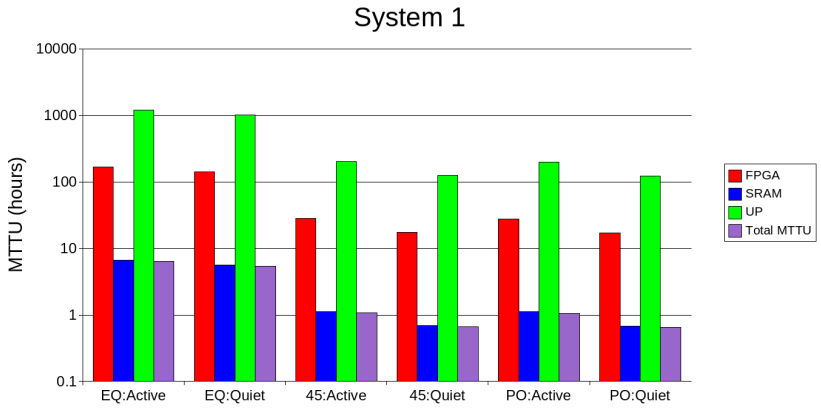

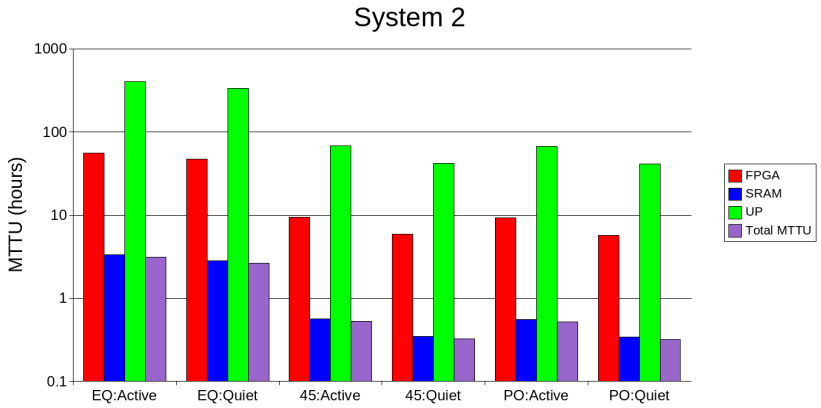

Case 2: ECC-Protected SRAM

- **ECC protection increases the MTTU so that it is ~2x of the FPGA**
  - MTTU of the system is dependent on a combination of the SRAM and FPGA

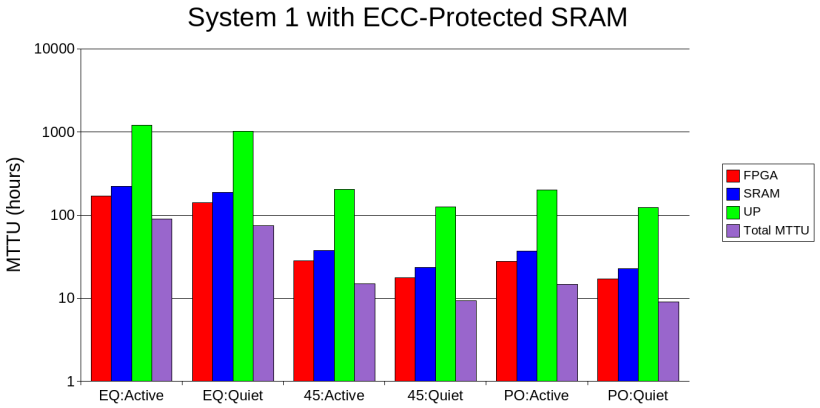

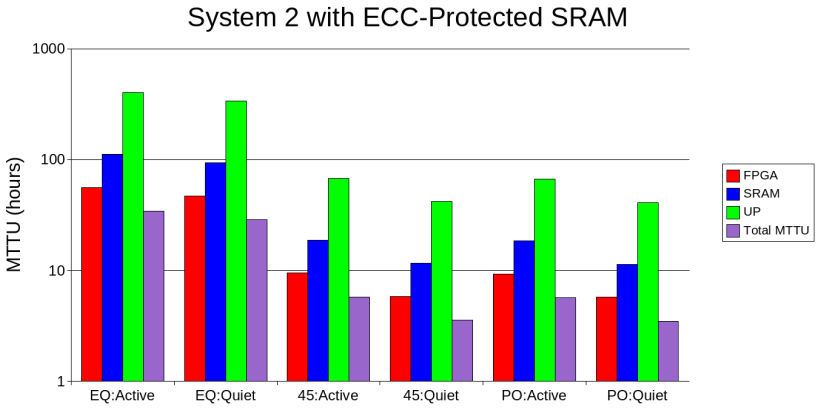

Case 3: ECC-Protected SRAM and 20% FPGA Utilization

- **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**

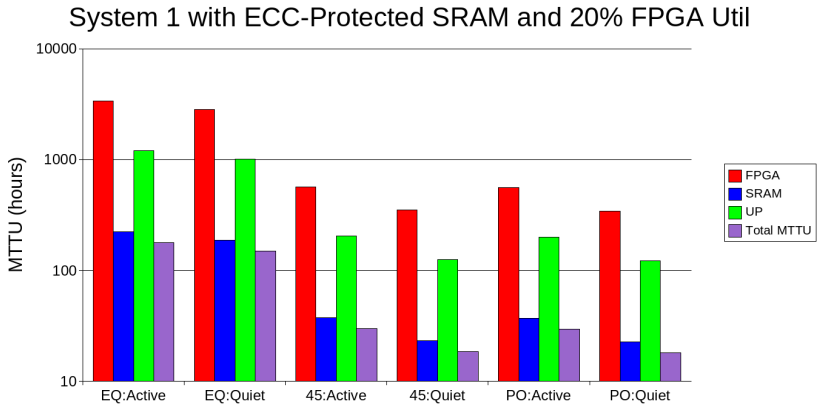

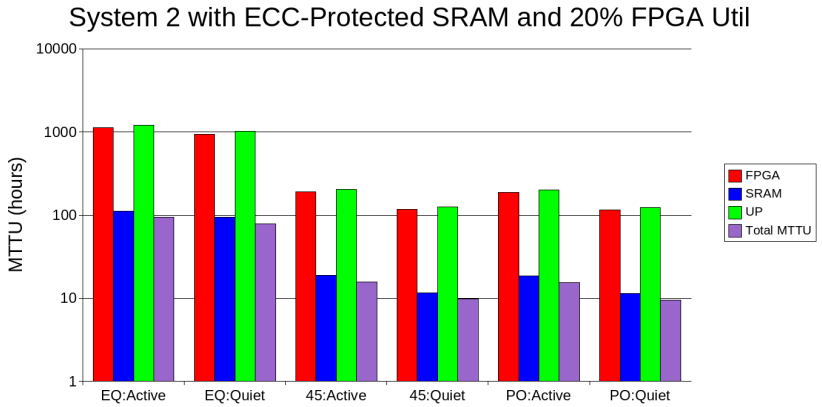

What Do These Numbers Mean?

- **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 the nuclear explosion?
    - Do you miss yet another chance to take the same picture?
  - 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
  - The systems I work with in LEO have a fault every 86 minutes
    - The device will have 6112 faults a year
    - Assuming an unmitigated design with 1 minute repair the availability is 0.98837
    - Our system requirements are 0.99999 or 5 minutes of downtime per year which possible but faults have to be mitigated

Overview

- **A Quick Introduction to Space Physics**
  - The Earth's Magnetosphere
  - Cosmic Rays
  - Solar Cycles, Solar Flares, and Coronal Mass Ejections
  - A Few of My Favorite Particles
  - Airplane Environments
- **Cosmic Rays and Electronics**
  - The History of Radiation-Induced Errors
  - Current reliability issues with Cosmic Rays and Modern Electronics
  - Radiation-Induced Failure Modes
- **Soft Error Estimates for FPGA-based systems**
  - FPGAs, memories, processors
  - What do these numbers mean?
- **Mitigation Methods**
  - Masking Errors
  - Repairing Errors
- **Testing Methods**
  - Modeling
  - Fault injection
  - Accelerated testing
- **Summary and Conclusions**

Mitigation Methods

- **Usually need to mask errors until you fix them**
- **Masking errors**
  - Redundancy-based systems
  - Hamming codes with or without bit interleaving
- **Repairing errors**
  - Scrub for memory-based devices
  - Reset and restart computation in extreme cases
- **The difficult aspect of airborne systems is that almost all mitigation methods were developed for systems with high error rates**
  - Harder to balance reliability with size, weight and power (SWaP) requirements
  - Some method will seem overzealous

Redundancy-Based Systems

- **Multiple copies of the system or design, output values are voted**
  - Triple-Modular Redundancy (TMR) being the most popular
  - Dual redundant lighter weight, but can only detect errors
- **Redundancy-based systems good for FPGAs**
  - Can TMR design inside one chip
  - Can TMR across three chips (spaceshuttle method)
- **TMR is impossible to get right on your own**
  - Incorrect implementations of TMR can be worse than unmitigated designs
  - Even correct implementations of TMR often have an exposed MBU cross-section
  - Xilinx and BYU both have tools
  - BYU tool does partial TMR, which might be best for airborne systems

Unprotected Cross-Section for Four Implementations	Edge Detection	Noise Filter
No TMR	25,938	23,554
Partial TMR: logic TMR'd, no input or output signals TMR'd	28,938	19,005
Partial TMR, logic, clock and reset TMR'd, no input or output data signals TMR'd	211	211
Full TMR	0	0

Hamming Codes

- **Can use error correcting codes for memories**
  - Can detect and/or correct errors in how they are stored
  - Bit interleaving can be used to reduce the chance of MBUs breaking the error correction
  - FPGAs access the memory as raw signals, so even non-ECC memory can be used with the right FPGA circuitry
  - Overhead associated with this method: lose some bits in the memory, lose some space on the FPGA
    - Overhead will depend on the implementation
    - Even a SECDED encoding and bit interleaving with two words would be useful
- **Some of the FPGA BRAMs have ECC protection and minimal bit interleaving**

Scrubbing Memories and FPGAs

- **Memories that have coding schemes on them can correct many errors**
  - Can scrub every time a word is read that has a detectable error
  - Can also scrub whenever possible to keep errors from accumulating
  - Depends on what the duty cycle of the SRAM is and how often errors manifest
- **FPGAs can be scrubbed through external devices or using the Xilinx internal scrubber**
  - Accumulating errors can defeat TMR, if you are using it
  - In space, we use an external scrubber in Actel devices to scrub the Xilinx FPGAs
  - External scrubbing also impossible to get right on your own
    - BRAM test frames need to be scrubbed, which is not published in open literature
    - Bad scrubbing creates SEL-like conditions on the device
    - Xilinx is prepared to help, so let them help you
  - Xilinx's internal scrubber is not sufficient for space (too many single points of failure), but probably best option for airborne systems that need scrubbers

Clever Low-Weight Mitigation Schemes

- **The field is wide open**
  - Mitigation methods for microprocessors have not been as explored as other fields
  - Could find ways to use the multi-core systems in a redundant method
  - ???

Overview

- **A Quick Introduction to Space Physics**
  - The Earth's Magnetosphere
  - Cosmic Rays
  - Solar Cycles, Solar Flares, and Coronal Mass Ejections
  - A Few of My Favorite Particles
  - Airplane Environments
- **Cosmic Rays and Electronics**
  - The History of Radiation-Induced Errors
  - Current reliability issues with Cosmic Rays and Modern Electronics
  - Radiation-Induced Failure Modes
- **Soft Error Estimates for FPGA-based systems**
  - FPGAs, memories, processors
  - What do these numbers mean?
- **Mitigation Methods**
  - Masking Errors
  - Repairing Errors
- **Testing Methods**
  - Modeling
  - Fault injection
  - Accelerated testing
- **Summary and Conclusions**

Testing Methods

- **Testing will save a lot of headaches**
  - Errors can often be hard to trace, repeat, and fix
  - Considering the difference between sea level and 60,000' neutron fluxes, you might not find errors until the system is airborne
  - Once airborne, you are too far in the process
- **Accelerator testing is the gold standard for determining faults**
  - MGH?
  - LANSCE at LANL is the "premier" facility
- **We have been trying to bring in the fault analysis earlier in the process**
  - STARC Tool: determines unprotected cross-section and problems with TMR in EDIF circuit representations
  - SEU Emulation: determines unprotected cross-section in FPGAs by using "on-line" reconfiguration to inject faults, makes finding errors much easier since there are no multiple independent errors, have an idea which bit is causing the error
- **Testing, modeling, and fault injection of full systems are very hard to do**
  - System is tested piece-meal and then pray it works as a full system

Learning to Test

- **Even though it is only counting, it's easy to introduce error into the count**
- **Good testing depends on the ability to accurate account for errors in the system**
- **JEDEC89 standard provides some guidance, but there is a lot to testing**
- **Easier, quicker, more fun to just tag along with experienced testers**
  - My coworkers taught me how to use our test fixtures and modify them to do tests I was interested in
  - I was introduced to Gary Swift (JPL, now Xilinx) and Jeff George (Aerospace) through the XRTC who taught me about statistics, systematic error, correcting fluence, accounting for energy loss through the lid and Silicon
  - Gary introduced me to Larry Edmonds, who taught me about statistic models for TMR and determining whether MBUs were coincident SBUs
  - I spent five months at NASA Goddard where I learned more about system level requirements, statistically detecting and splicing data sets for anomalous data, physics of SEEs

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 encoding 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**