Outside of Normal Operating Conditions: Using Commercial Hardware in Space Computing Platforms for Ubiquitous Sensing

Heather Quinn

Los Alamos National Laboratory

Overview

- Ubiquitous Sensing for National Security Needs
- Flying Commercial
- Radiation Tests to Reduce Risk
- Conclusions

Sensing Applications for National Security

- **In recent years, much of LANL's mission has focused on persistence surveillance of targets and interests to provide an overall reduction in threats to the US**
- **This data plays an important role in national security and policy decisions**
- **Data are collected from a number of platforms: distributed sensor networks (DSNs), airplanes, unmanned aerial vehicles (UAVs), and satellites**
- **The collected data is from a number of sensor types: imagery, seismic, radiation, temperature, radio frequency**
- **Many of these sensors grew out of science programs**
  - Satellite-based detectors that could sense neutrons in the ground have been used to determine whether there is water on Mars and whether there is nuclear proliferation

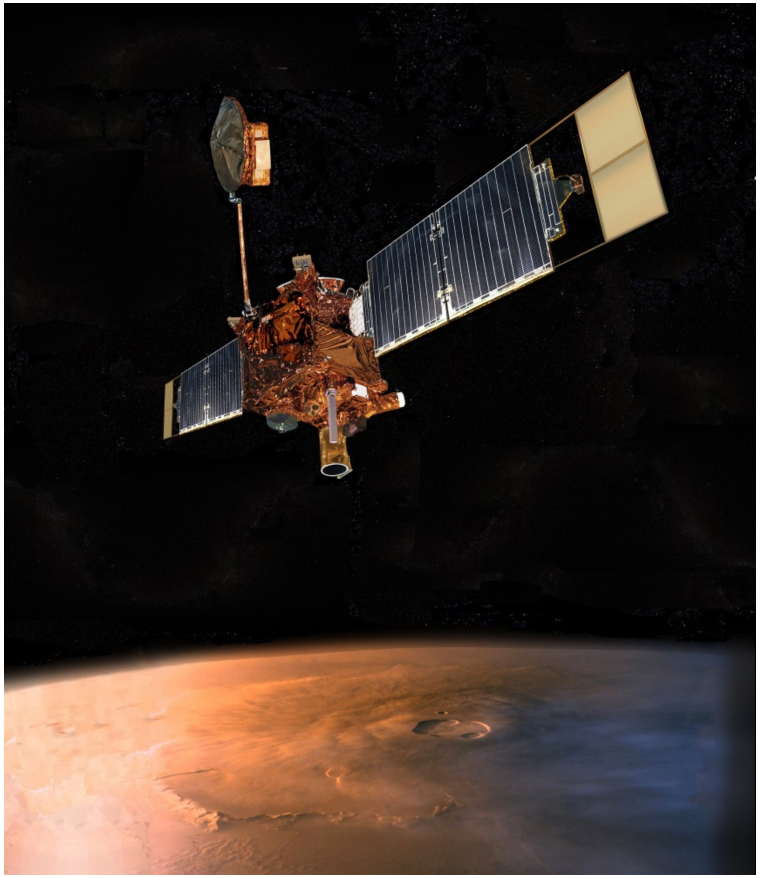

http://mars.jpl.nasa.gov/mgs/gallery/images/mgs-mons.jpg

Transitioning to Ubiquitous Surveillance

- **The lab is striving for a global reduction of threats**
- **The lab's mission is to grow our sensing capabilities so that we could provide constant, global – ubiquitous – surveillance**
  - Increasing the view of our sensing capabilities provides more information, giving us global coverage
  - Increasing the sensitivity of our sensing capabilities provides more accurate information
  - Increasing the number and types of surveilling platforms to provide options for collecting data
- **The better, the wider, the more proliferate our sensing capabilities are, the less likely we are to miss important events around the world**

Examples of Ubiquitous Sensing

- **DSNs:**
  - Smart paint that can monitor the integrity of physical infrastructure, such as buildings or bridges
  - Intelligent rocks that can monitor the movement of radioactive materials on highways
- **Airplanes/UAVs:**
  - Wide area persistence imagery that can track movement through cities
- **Satellites:**
  - Neutron detectors that can globally monitor the adherence to the Comprehensive Test Ban Treaty.
  - Imagery that can globally monitor whether nuclear plants are being built that could be later disguised

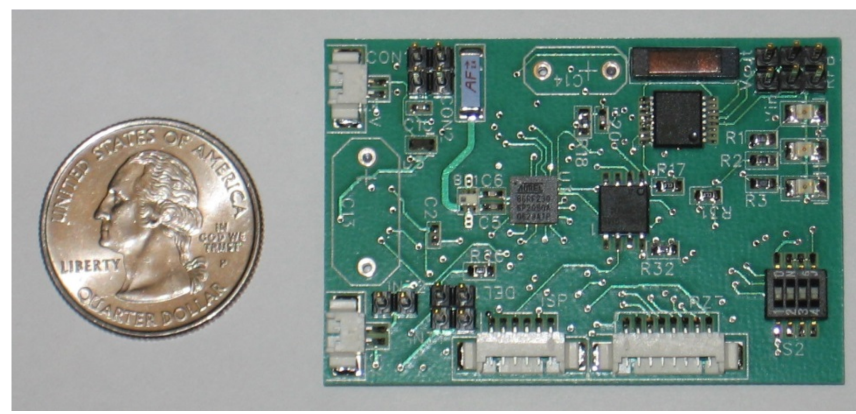

http://int.lanl.gov/news/index.php/fuseaction/home.story/story_id/11142

Challenges of Ubiquitous Sensing

- **Designing wide-area, extremely sensitive sensors is challenging**
  - Done with one, expensive and expansive sensor or tons of less expensive, less capable sensors?
  - How to blend different sensor types and capabilities?
- **Wide area, constant surveillance stresses computation and communication systems**
  - Do you need to trade off computation for communication?
  - How much can processing can be completed on the system?
- **The amount of data collected from these efforts presents many challenges**
  - We could reduce transmission of unusable or uninteresting data, transmit information instead of data, prioritize data for retrieval
  - We could not do that in the late 1990s using radiation-hardened electronics

LANL's Approach to High-Performance On-Orbit Processing: Using Commercial Technologies through Advanced Engineering

- **Use commercial-based technologies for high performance portions of the space systems**
  - Leverage billions of dollars of world-wide commercial investment in semiconductor technology
  - Employ well-tested technologies with large user bases rather than unique space solutions
  - Exploit inherent radiation tolerance (e.g., total ionizing dose) of these components
- **Use system-level, module-level, and application-level engineering to provide the robustness needed for the system (don't "over-engineer" systems)**
  - Employ an excellent understanding of both mission and technologies
  - Employ existing and new mitigation techniques to add robustness: e.g., redundancy, repair, and reconfiguration
- **Use more conventional radiation-hardened technologies in high-risk portions of the system or where performance and cost are not drivers**
  - Spacecraft interfaces
  - Critical non-volatile memory

COTS Electronics in Space

- **Fifteen years ago LANL partnered with Xilinx to determine if the commercially-available, radiation-tolerant Xilinx Virtex field-programmable gate arrays could be used in space**
  - Could these components provide the speed and agility we wanted without corrupting our data stream and affecting our national security mission?
- **To use this hardware in space a number of questions needed to be answered:**
  - Would radiation cause destroy the FPGA while in space?
  - Would radiation-induced errors make fault-tolerant computing impossible?
  - Could we mitigate the radiation problems?
  - Would the package survive the vibrations caused by the launch without breaking off the board?
  - Could the package handle the thermal cycles without breaking the FPGAs off the board or having temperature-related reliability problems?

Fault-Tolerant Computing with FPGAs in Space

- **Components did not exhibit destructive radiation effects but did exhibit single-event upsets (SEUs or upsets)**
  - Upsets cause memory cells to change values
- **Radiation testing showed that even a single SEU can cause the circuit to output bad data**
  - Accumulating SEUs increase the likelihood that output data is corrupted and increase device's current draw
- **The component is essentially "blank" and we could decide how to mitigate errors**
  - To date, best option for mitigation SEUs is to mask them through triple-modular redundancy (TMR)
- **The device is reprogrammable: the configuration ports could be used to fix the radiation-induced faults**
  - On-line reconfiguration, called scrubbing, used to remove SEUs
  - Off-line reconfiguration used to remove SEFIs

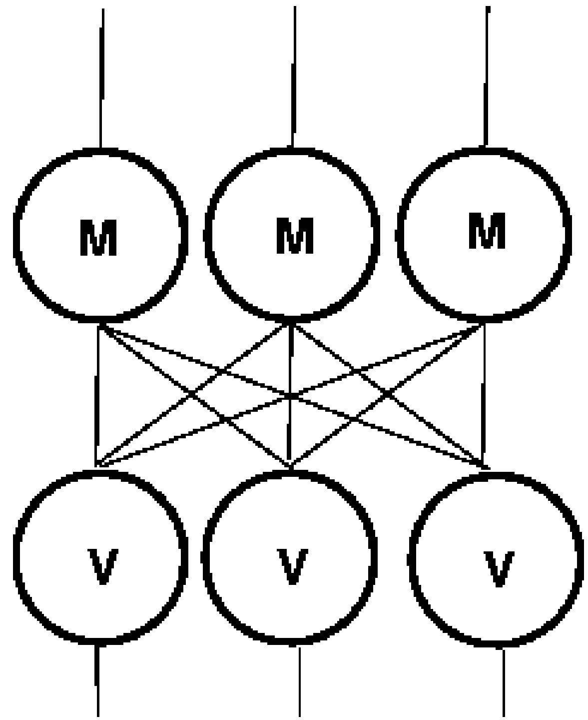

Cibola Flight Experiment: Demonstration of Fast On-Board Processor with COTS Parts

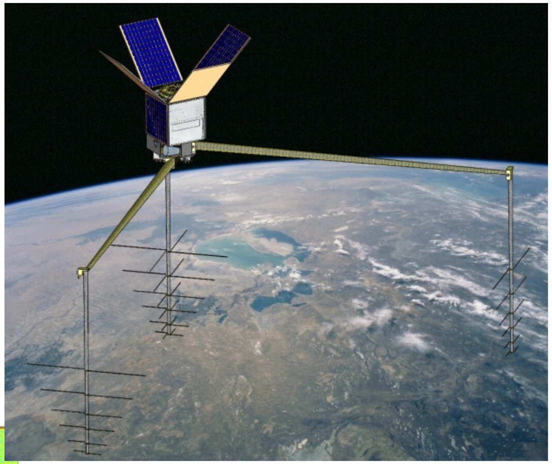

- **Launched March 2007**
- **Orbit: Circular 560 Km, 35.4 degree inclination**
- **Software Radio:**
  - Four channels, 20 MHz bandwidth each
  - Tunable from 100 to 500 MHz,
  - 3-board, 9 Xilinx Virtex FPGA 300-Gop/sec (peak) re-configurable computer (RCC)
  - 4-element antenna array

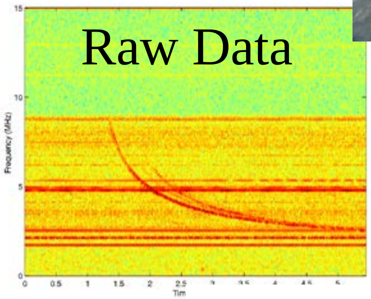

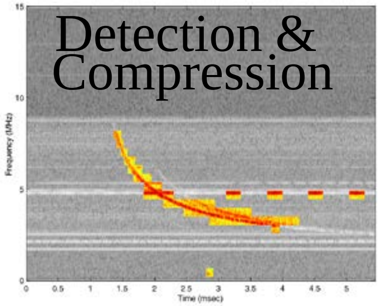

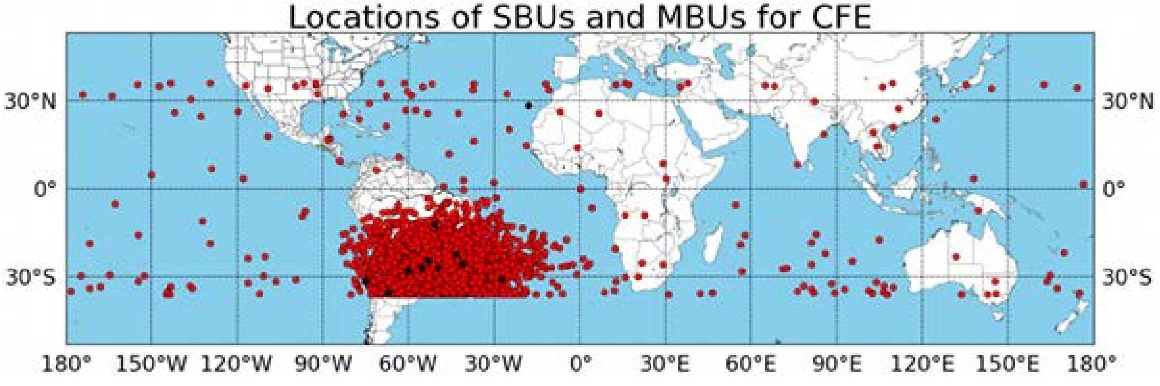

Mission Response Module: Second Demonstration of Fast On-Board Processor with COTS Parts

- **Launched into low Earth orbit in 2011 on a US Department of Defense satellite**
- **Software Radio:**
  - Four channels, 60 MHz bandwidth each
  - Two separate units with two Xilinx Virtex-4 FPGAs: each unit can tune to the same or different channel

Map of SBUs and MBUs

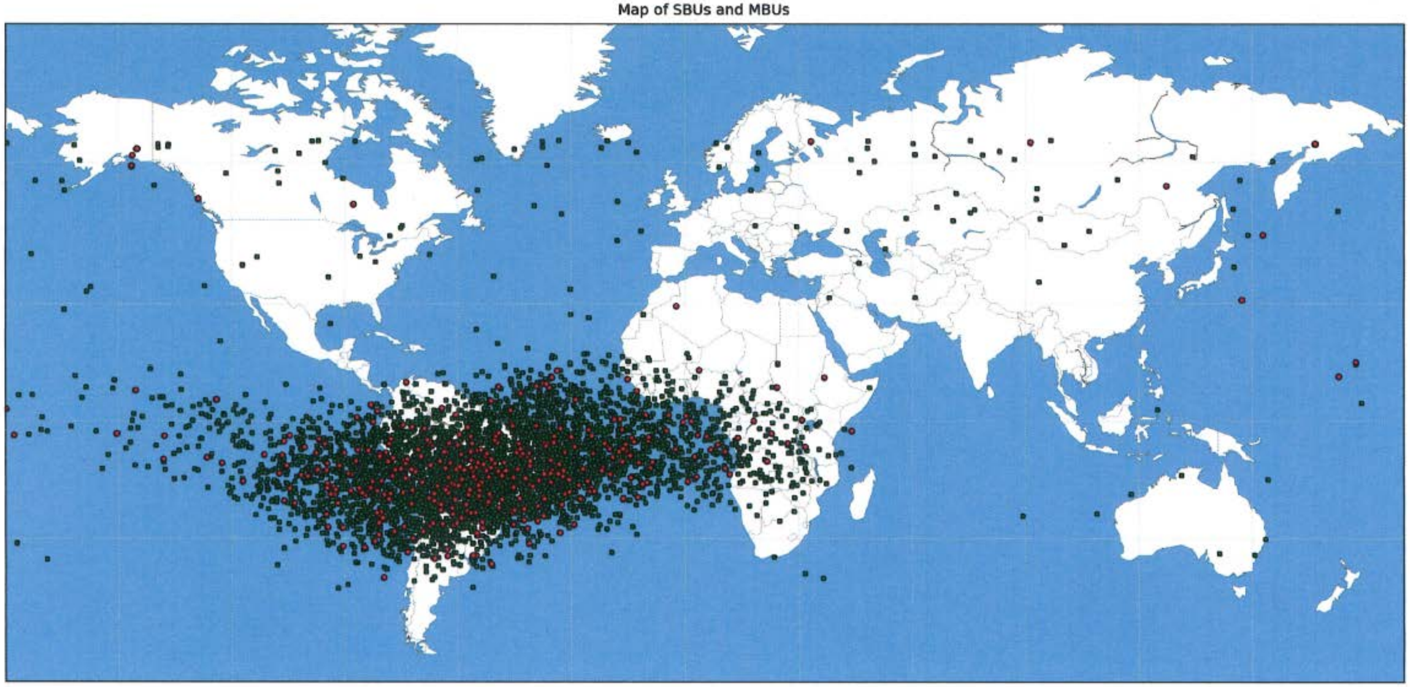

Life after CFE and MRM

- **In 2008, the DOE gave us the chance to transfer our knowledge from CFE and MRM (still in integration) to the operational DOE space mission:**
  - Space-based Nuclear Detonation Detection (SNDD) is a suite of payloads integrated into GPS satellites
  - Provide 24x7 converge of the Earth for Comprehensive Test Ban Treaty monitoring
- **One of the hardest/worst space missions:**
  - Long duration
  - Heightened radiation environment
  - Nuclear survivable
  - … and we want to build the payloads out of $2 commercial components that are designed to work in cars and toasters

Reducing Risk through Environmental Testing

- **Most commercial components have not been tested for radiation effects**
  - The only way to put them into GPS is to prove that it will not hurt the DOE or GPS mission
  - The parts need to be qualified for space usage, which means that we need extensive test radiation data
- **But first we needed to find candidate parts….**
  - FPGAs were an accidentally perfect first demonstration vehicle:
    - The Xilinx FPGAs did not have a sensitivity to destructive single-event effects
    - The FPGAs had a good, natural tolerance to total ionizing dose
    - The lab was filled with expert FPGA designers that could work with or around the design tools as necessary
  - We suddenly need to cope with an onslaught of really bad electronic components
    - Many of the components are highly sensitive to destructive single-event effects
    - Some of the components are so complex that there is an entire zoo of failure modes

Screen All of the Components at LANSCE First

- **It was clear that we needed to start using LANSCE as a testing partner**
  - The fast neutrons are a reasonable analog to high-energy protons
  - We needed to “slow down” our tests so that we could observe the errors in the components one at a time
    - The single-event effects are from a indirect ionization reaction is 5-7 orders of magnitude smaller than the direct ionization effects we were getting at heavy ion accelerators
    - The flux was not as high as proton accelerators
  - The neutrons are non-ionizing, so we do not churn through parts due to dose-related problems
  - The location is also extremely convenient for us
- **We developed a policy to screen parts at LANSCE before moving onto heavy ion testing**
  - If the component could not survive a LANSCE test, then it would not survive the rest of the qualification process or a long space mission
  - It is still possible to have failures at heavy ion facilities, but not as many

Advantages of Testing at LANSCE

- **We tested both parts at LANSCE and LBL**
- **One part had no destructive failures at LANSCE, but had some destructive failures at a high threshold at LBL**
  - Could still be a reasonable part to deploy
- **The other part had destructive failures at LANSCE and many destructive failures at LBL**
  - We did not need to do more testing at LBL after the failures at LANSCE….

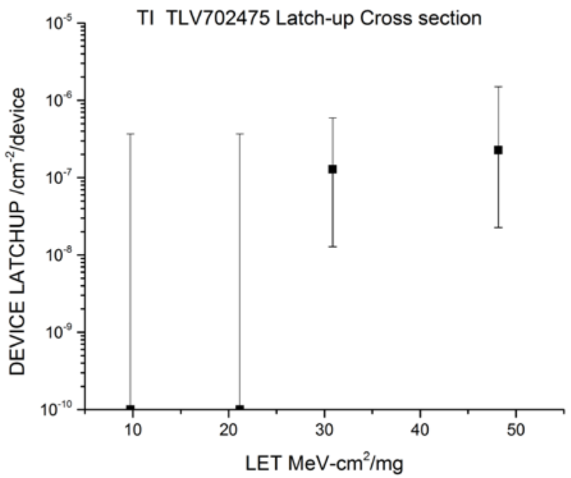

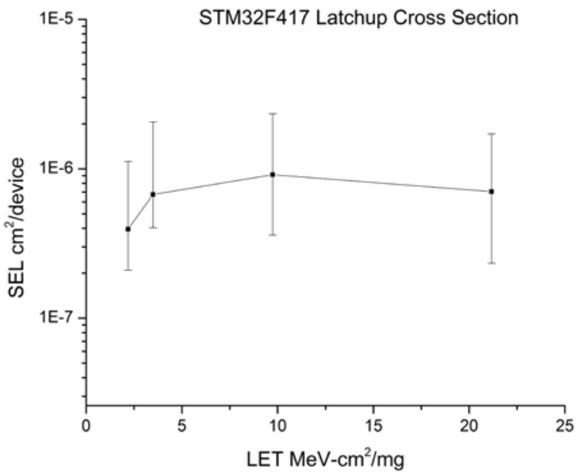

Memory

- **Memory is an essential part of computational systems**
  - For many systems the source of where many radiation-induced errors comes from
  - For deployed systems need to find reasonable memory components
  - Particularly difficult to find dynamic RAM without destructive failure modes and low SEFI sensitivities
- **Over the years we have tested many different samples of SDRAM from many of the SDRAM manufacturers**
  - Memory array has very low sensitivity to SEUs, but the memory array is very large
  - The SEFI sensitivity is very high, but is on the order of a single SEU across the entire memory array

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

SDRAM SEFIs

- **The SDRAM SEFI failure mode is particularly destructive to data:**
  - The radiation strike causes one to many bits to be overwritten for entire column of the device
  - Many of these errors cannot be corrected with standard “correct one, detect two” encoding schemes
- **To use many of these parts, would need to use block encoding schemes, which might not work with the computational model for how the memory is accessed**

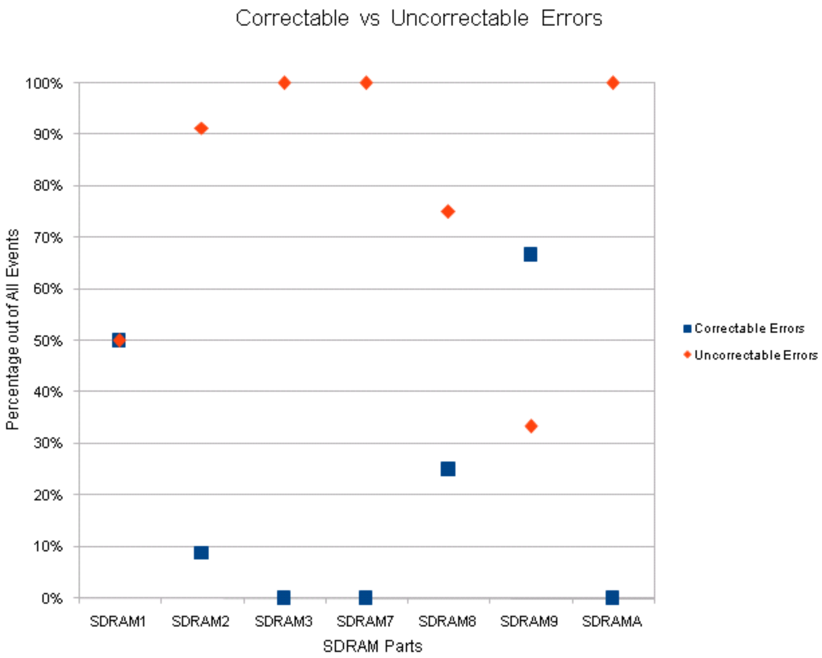

ARMs and Microcontrollers

- **Finding a reasonable microcontroller for background and configuration tasks will allow us to reserve the radiation-hardened microprocessor for mission critical processing**
- **We have tested:**
  - ST Micro ARMs
  - Texas Instruments MSP430, DSPs and ARMs
  - Xilinx ARM
  - NXP ARM

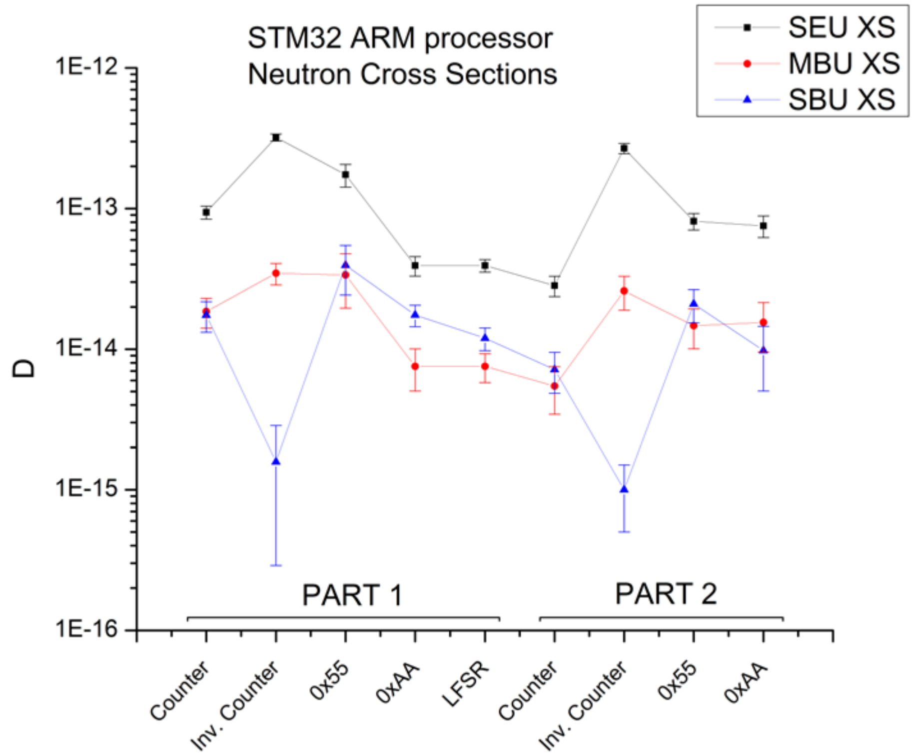

Los Alamos NATIONAL LABORATORY EST.1943

Operated by Los Alamos National Security, LLC for NNSA

Texas Instruments C6474 Tri-core DSP

- **We were particularly interested in this DSP for several trips:**
  - Large amount of memory
  - Fast computation of signal processing data sets
- **The error rates were particularly high, but we found that it was possible to mask the errors in hardware using software mitigation**

	Cross-Sections
SEU bit-cross-section	7.30×10^{-16} cm^2/bit (4.45×10^{-16} cm^2/bit, 1.12×10^{-15} cm^2/bit)
SEU device cross-section	1.65×10^{-7} cm^2/device (1.01×10^{-7} cm^2/device, 2.54×10^{-7} cm^2/device)
SEFI cross-section	4.13×10^{-10} cm^2/device (1.76×10^{-10} cm^2/device , 8.16×10^{-10} cm^2/device)

Use LANSCE to Experiment with Mitigation

- **After eliminating parts with destructive effects, we might still have components that will destroy the data**
- **Testing mitigation methods requires ensuring only one error is in the system at a time**
  - Impossible to do at heavy ion facilities for many components
  - Many accelerators cannot be tuned to a flux that low without problems with dosimetry
  - UC-Davis and LANSCE can both be tuned low enough to allow for mitigation tests
  - LANSCE is still in a great location for us
- **The last several years we have been doing extensive testing to show that our mitigation method for FPGAs can be ported to microprocessors**

Unmitigated Software Test Results on the C6474 DSP

- **Not all SEUs will create SDC, crashes, or other types of errors**
  - Device utilization, logical masking, and compensating failures lower the error rate
  - SEUs can be categorized into ones that create observable errors by affecting calculations and ones that do not
- **The length of time the data is in the cache is important**
  - For data that is read once, the SEU would need to occur in between writing and reading – any SEUs after reading would not be observed and likely overwritten
  - Global values or constants are more likely to have observable errors because the values are read repeatedly without refreshing
- **The amount of data needed for a calculation is important**
  - The more data that a calculation uses, the more likely SDC will affect the calculation

Unmitigated Software Test Results on the C6474 DSP

- By studying the amount of time data remains resident in the L2 cache, we can understand the difference in the reliability of long-term and short-term resident data variables
- Some data will be read many times and some data will be read only once
- These results show that there is nearly 15 times decrease in noticeable errors from data that is read frequently to data that is read once
- This result indicates that selective TMR approaches will be more useful for data that is written once and read many times, such as global constants.

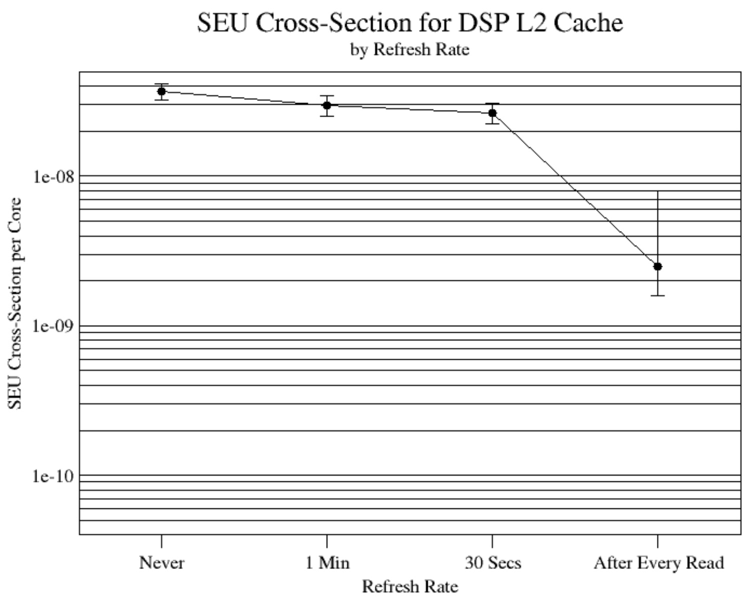

Mitigated Software Test Results

- **While the SEU bit cross-sections are quite small, the SEFI cross-sections are 400 times larger**
  - For many calculations dual module redundancy (DMR) would not be strong enough
  - Triple-modular redundancy (TMR) would provide masking, which can be useful for higher error rates
    - DMR fails at 2x the rate of the unmitigated code and must be reset after each error
    - TMR fails at 3x the rate of the unmitigated code and can mask at least 1 error
- **The TMR granularity is important**
  - The more data that are used, the more likely the calculation fails
  - Fine-grained granularity can tolerate more errors
- **The software structure is important**
  - The reliability of recursive codes will be dependent on the iteration – the more iterations, the more likely a failure could be accumulated

Benchmarking at LANSCE

- **Recently have been part of a collaboration for developing standard benchmark codes/circuits for radiation tests of mitigated software/hardware**
  - Need to be able to determine whether the mitigation process is masking errors in the system
  - Need to be able to determine which mitigation technique to use for the (power, speed, effectiveness) tradespace
- **Ten organizations have been collaborating for over a year to design a benchmark**

First Benchmark Test at LANSCE in December 2014

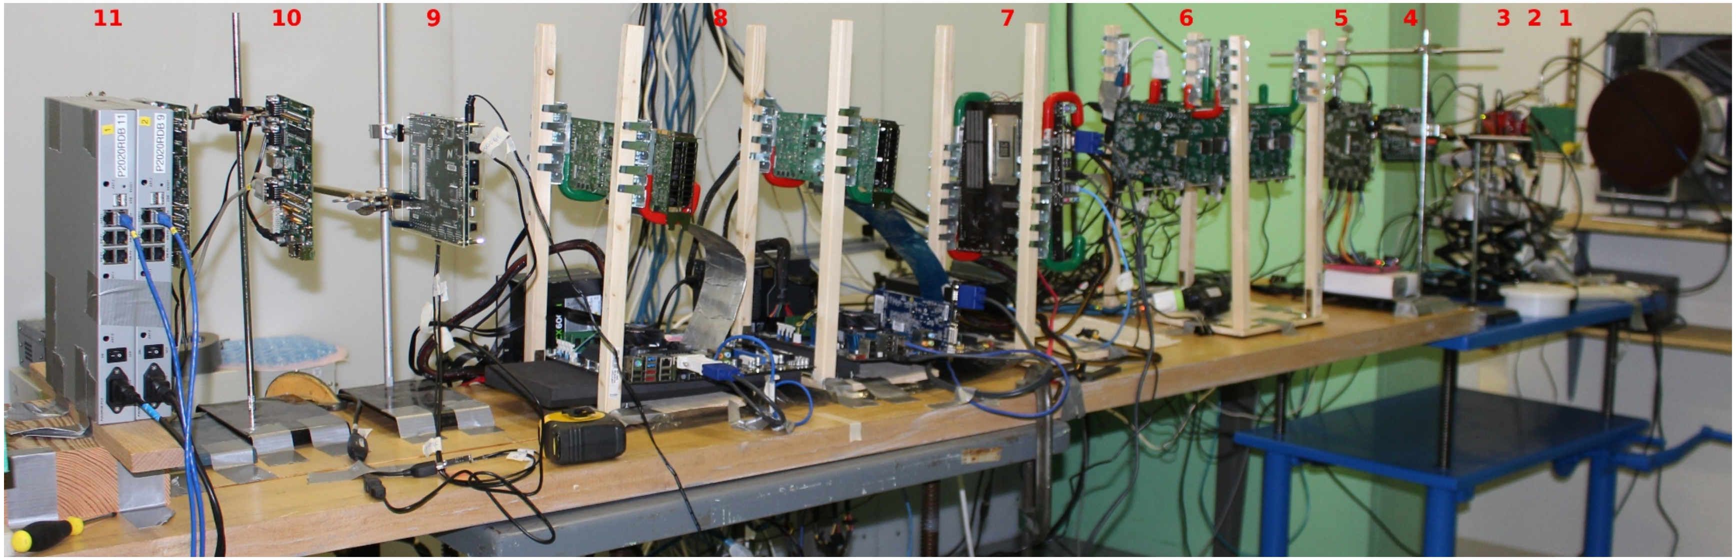

Microcontroller Results

Code	Tiva	MSP430F2619	MSP430FR5739
AES	0.30 (0, 1.1)	0.38 (0.04, 1.37)	0.85 (0, 3.1)
AES TMR	0.31 (0, 1.1)	3 (1, 5)	2 (0, 7)
Cache	75 ± 10	8 ± 2	10 (6, 15)
Cache TMR	0.27 (0, 1.0)	0.21 (0, 0.76)	2 (0, 8)
Coremark	0.75 (0.15, 2.20)	1.27 (0.51, 2.61)	N/A
M x M	59 ± 13	4 (2, 6)	1 (0, 4)
M x M TMR	10 (7, 14)	0.27 (0, 1.0)	2 (0, 8)
Qsort	59 ± 13	3 (2, 5)	25 (16, 38)
Qsort TMR	--	7 (4, 10)	--

Software was mitigated using Trikaya software technique for s/w mitigation

- **All of these components are very small, which is why the FIT rate is small**
- **These results show that AES-128 is naturally resistant to errors: very small amount of memory and processing**
- **Many similarities in results due to forcing similar amount of memory**
- **These values are not normalized to amount of work performed:**
  - Cache test makes the MSP430F2619 look like the most robust operation
  - In reality, it is doing far less processing than the Tiva
  - The slower processing in Coremark shows how the slower processing decreases resilience to errors

Conclusions

- **Ubiquitous sensing is an important aspect of national security and reducing global threats**
  - The amount of data collected drives the need for more efficient computational and communication systems
  - FPGAs have been useful in both ground-based and satellite-based systems
  - Expanding the program to look at more commercially available electronics alternatives to radiation-hardened electronics
- **Radiation testing showed that fault-tolerant computing could be difficult**
  - Many components are sensitive to SEL that could damage the component or SEUs/SETs that could damage the data
  - Mitigation is possible, but requires extensive testing
- **Testing partners, like LANSCE, are a valuable asset for next-generation computational design work**