



An Introduction to Mission Risk and Risk Mitigation for Xilinx SRAM FPGAs

Heather Quinn

Los Alamos National Laboratory

Acknowledgements

This work was funded by the DOE under the JAS, SOPRANO and DAPS projects, the Rhino project through a NASA ROSES grant, and the DOD under the FPGA Mission Assurance Center.

This presentation includes the work and help of LANL, BYU, Xilinx, JPL, SRI, SEAKR, and The Aerospace Corporation staff members.

Introduction

- **Many organizations have started using commercial SRAM-based FPGAs in space-based computing platforms**
 - Well-suited to DSP-oriented satellites
 - Custom hardware speedups without the cost of manufacturing an ASIC
 - Reconfigurability can extend the useful lifetime of the system by allowing the system to reconfigure to meet changing mission or science needs
- **Unfortunately, the devices are susceptible to SEUs which can make fault-tolerant computing challenging but not impossible**
- **In this talk we will discuss:**
 - What the mission risk is
 - What the advantages of using these devices are
 - How SEUs affect fault-tolerant computation
 - How to mitigate errors to increase reliability and availability

Caveats:

Field-Programmable Gate Arrays (1 of 2)

- **In this talk, we will focus on SRAM-based, reprogrammable FPGAs, where logic is implemented in lookup tables and the routing uses programmable switches**
 - Xilinx is the preferred vendor, because they have published several reports verifying latchup-immunity [1] [2]
- **For the rest of the talk, the term “FPGA” will be used to indicate only the Xilinx reprogrammable, SRAM-based FPGA**
- **We will discuss a comparison of SRAM-based FPGAs with Anti-Fuse-based or Flash-based FPGAs, but all other types of FPGAs will have their process technology defined**
- **While older technology will be occasionally discussed in this talk, much of the talk will focus on the Virtex-4**

[1] G. M. Swift, “Virtex-II static SEU characterization,” Xilinx Radiation Test Consortium, Tech. Rep. 1, 2004.

[2] G. Allen, G. Swift, and C. Carmichael, “Virtex-4VQ static SEU characterization summary,” Xilinx Radiation Test Consortium, Tech. Rep. 1, 2008.

Caveats:

Field-Programmable Gate Arrays (2 of 2)

- **While in this presentation we will attempt to make comparisons between SRAM-based, Flash-based, and Anti-Fuse-based FPGAs, these types of comparison are problematic:**
 - The logic on Flash-based and Anti-fused-based FPGAs are simpler than SRAM-based FPGAs, so the logic capabilities and capacities are hard to directly compare
 - The Anti-Fuse-based FPGAs are not reconfigurable
 - Each of the devices are subject to their own radiation-induced faults
 - **We suggest people be aware that politics in the radiation community and with the vendors have skewed some of the arguments about using one type of device over another type of device**
 - Get the facts, make your own decisions
- Or —
- Experiment with the hardware, make your own decisions

Mission Risk:

Caveats

- **At the end of the day, the ability to use this device on orbit will be based on your satellite's mission risk**
- **At LANL, we try to use Xilinx FPGAs whenever:**
 - Occasional data loss due to SEUs is permissible
 - Fast data processing is needed
- **In LANL's view, these devices are fast, power-efficient data processing devices**
 - Our plan is not to use them in mission critical areas of the satellite, such as the bus interface
 - When we do use them, we use the newest and largest device available
- **Scenarios do exist where no reasonable radiation-hardened alternative exists**
 - For these missions, mitigation and recovery from errors will be essential
 - For these missions, *testing* the mitigation and recovery plan will be essential

Mission Risk:

What are the Concerns?

- **While many organizations are already using these devices in space, many other organizations remain nervous**
- **The detractors of using these devices are:**
 - The “soft” SRAM cells used for the configuration memory
 - The emerging flight heritage
 - The perceived error rate of these devices
 - The perceived difficulty of mitigating errors in the system
- **Some organizations are worried that these devices are too unreliable to use for many applications due to the SEU susceptibility**
- **We will discuss today whether these concerns are valid or not so that engineers, system designers and program managers can make rational decisions about using Virtex devices in their spacecrafts**
- **Let’s move the decision making from fear to facts**

Mission Risk:

What Are the Facts Needed to Make the Decision

Mission Risk vs. Errors

■ Mission Risk

- System designers need to ask themselves and their teams what does error mean in your system?
 - The satellite becomes inoperable?
 - Bad data?
 - Actionable bad data?
- At LANL, we have found a number of programs can tolerate occasional bad data, as long as it is not actionable – Vela Incident
- Two views of the mission risk problem:
 - SEUs as noise
 - Errors as probabilities

■ Errors

- We're focusing on radiation only today:
 - The first order concerns are total ionizing dose and single-event latchup, as these effects can destroy the device
 - The second order effects are single-event upsets and single-event transients, as these effects can destroy data

Mission Risk:

Total Ionizing Dose

- **Total ionizing dose provides an upper bound on how much radiation a device can tolerate in space before the process is too degraded to use**
 - This provides a time limit for how long a device can survive in a particular orbit
 - Can be estimated using CREME96
- **The Xilinx Virtex-4 is confirmed to have a TID of 300 kRad (Si) [1]**
- **The Actel Flash-based RT ProASIC3 can withstand 15-25 kRad of dose**
 - At 15 kRad 10% propagation delay occurs in the circuit and the device is not programmable [2]
 - The TID problem is a problem with all Flash technology
- **The Actel Anti-Fuse-based RTAX-S/SL is confirmed to have a TID of 300 krad (functional) and 200 krad (parametric) [3]**
- **Unless your system is in a particularly vicious orbit both the Virtex-4 and the RTAX-S/SL are both reasonable to use in terms of TID performance**

[1] <http://www.epn-online.com/page/new57560/xilinx-introduces-virtex-4qv-fpgas.html>

[2] <http://www.actel.com/products/milaero/rtpa3/default.aspx#radiation>

[3] <http://www.actel.com/products/milaero/rtaxs/default.aspx>

Mission Risk:

Single-Event Latchup

- **Single-event latchup is a radiation-induced version of a destructive event caused by the vertical thyristors in CMOS technology**
 - As thyristors exist in all CMOS technology, it is up to the vendors to design devices so that latchup does not occur
 - Most satellite designers will not fly devices that latch below 80 MeV-cm²/mg
- **The Xilinx Virtex-4 is SEL-immune to 125 MeV-cm²/mg [1]**
- **The Actel Flash-based RT ProASIC3 is SEL-immune to 96 MeV-cm²/mg [2]**
- **The Actel Anti-Fuse-based RTAX-S/SL is SEL-immune “in excess of 117 MeV-cm²/mg” [3]**
- **All three devices are reasonable to use in terms of SEL immunity**

[1] <http://www.epn-online.com/page/new57560/xilinx-introduces-virtex-4qv-fpgas.html>
[2] <http://www.actel.com/products/milaero/rtpa3/default.aspx#radiation>
[3] <http://www.actel.com/products/milaero/rtaxs/default.aspx>

Mission Risk:

What is the Real Problem?

- **The difference between the RTAX-S/SL and the Virtex-4 are the single-event effects**
 - Single-event upsets: bit flips in SRAM and latches
 - Single-event transients: changes in the gate output
- **The Virtex-4 is susceptible to SEUs starting at 1 MeV-cm²/mg heavy ions and 63.3 MeV protons which can cause changes in the user circuit, changes to the user circuit's state and/or device functionality**
 - There is no published data on SETs in the configuration memory of the Xilinx devices. Even if there were evidence of SETs, the only difference would be a shift in the user FF cross-section, which is ~0.25% of the device by number of bits
- **The Actel Flash-based RT ProASIC3 has SEU and SET susceptibilities [1]**
 - Configuration Flash Cells: No errors observed
 - D-Type Flip-Flops : SEUs starting at 6 MeV-cm²/mg heavy ions and 63.5 MeV protons
 - SRAM Memory: SEUs starting at 1 MeV-cm²/mg heavy ions and 63.5 MeV protons
 - FlashROM Memory: No errors observed
 - Global Clock: SETs starting at 4 MeV-cm²/mg heavy ions
 - I/O Bank: SETs starting at 7 MeV-cm²/mg heavy ions
- **The Actel Anti-Fuse-based RTAX-S/SL has SEUs [2]**
 - "Logic" SEUs starting at >37 MeV-cm²/mg in heavy ions
 - "Memory" SEUs starting at >30 MeV-cm²/mg in heavy ions
- **All of the devices have some SEUs – the only difference is onset and quantity**

Mission Risk:

The Comparison So Far....

- **The Actel Flash-based RT ProASIC is probably not useful for multi-year missions due to low TID and will have many problems with SEUs and SETs**
 - Let's take that device off the table for now
- **The Xilinx Virtex-4 and the Actel Anti-Fuse-based RTAX have comparable TID and SEL-immunity**
 - The anti-fuse device will have fewer problems with SEUs than the Virtex-4, as the onset threshold is much higher and there are far fewer SRAM and latch cells in the anti-fuse device
- **We're done, right? Choose the anti-fuse device, right?**
- **Not yet – there still remains three more comparisons that should be included in this decision**
 - Mitigation
 - Device sizes
 - Reconfigurability

Mission Risk:

Mitigation Options

- The Actel RTAX-S/SL has many built-in mitigation options – EDAC on memory and triplication of user FFs
- The Virtex-4 has EDAC on the BRAM but no built-in mitigation on configuration memory. There are two available tools for the configuration memory:
 - The Xilinx TMRTool
 - The BYU BL-TMR Tool
- Since the Virtex-4 has a lower onset threshold for SEUs and more memory than the RTAX, what is the maximum availability for these devices?
 - There is a class of SEUs that cause single-event functional interrupts (SEFIs) that cause the device to not operate properly (either partially or fully) until reconfigured
 - Assuming that the rest of the SEUs can be mitigated, the SEFI rate is the upper bound on availability of the device

Mission Risk:

What Are SEFIs?

- **SEUs to the control logic and device state registers can cause the device to become temporarily interrupted**
- **Side Effects of SEFIs:**
 - Immediate loss of full device function
 - POR, GSIG, Scrub
 - Scrub SEFI could damage device
 - Reprogram by pulsing PROG as soon as possible
 - No impact to device function
 - SMAP/JTAG, FAR
 - Reprogram as soon as possible
 - Possible loss of full device function
 - Shutdown SEFI
 - Mitigate by scrubbing CFG_CLB column.

Mission Risk:

How Often do SEFIs occur on Orbit?

Mean Time to SEFI for Selected Orbits in YEARS, calculated by CREME96

Solar Minimum Quiet, AP8max, z=1 μm and 100 mils of Al

Orbit	Altitude (km)	Incl*	POR	GSIG	SMAP+	TOTAL	HI%
LEO	400	51.6°	1225	2161	1500	515	58
	800	22.0°	100	114	112	36	13
POLAR	833	98.7°	131	165	146	49	14
CONST	1200	65.0°	32	37	35	11	3
GPS	20200	55°	240	309	596	110	7
GEO	36000	0°	225	560	290	103	91

* Incl = Inclination HI% = fraction from heavy ions
SMAP+ = SMAP & FAR SEFIs combined

Mission Risk:

How do SEFIs affect Availability?

- **When SEFIs occur, the scrubber needs to detect and correct the SEFI**
- **Correction of the SEFI includes stopping the circuit and performing a complete off-line reconfiguration of the device**
 - SEFI detection \ll 1s
 - Complete reloading of the application from PROM/CRAM: ~ 6 minutes (almost entirely decompression time in SPARC software)
 - Complete reloading of the application from SDRAM: < 10 s
- **As a worst case scenario, this process will take 6.02 minutes and the device is inoperable the entire time**
- **As a best case scenario, this process will take between 1-11 seconds and the device is inoperable the entire time**
- **Poisson statistics also dictate that if an event will occur in a timeframe X, then there is a**
 - 37% chance no event occurs
 - 37% chance one event occurs
 - 18% chance two events occur
 - 6% chance three events occur
 - 2% chance four events occur
- **And the availability numbers are...**

Mission Risk:

Availability Rate from SEFIs Assuming 6 Minute Recovery

Orbit	One SEFI	Two SEFIs	Three SEFIs	Four SEFIS
LEO (400 KM)	0.99999998	0.99999996	0.99999993	0.99999991
LEO (800 KM)	0.99999997	0.99999994	0.99999991	0.9999999
Polar	0.99999998	0.99999995	0.99999993	0.99999991
CONST	0.9999999	0.9999998	0.9999997	0.9999996
GPS	0.99999999	0.9999999	0.9999999	0.9999999
GEO	0.99999999	0.99999998	0.99999997	0.99999996

- Even the worst, worst case scenario (CONST with 4 SEFIs in the timeframe) meets the minimum availability rate of most satellites – 5 “9s”
- In the best case scenario (LEO at 400KM) the maximum availability rate is 7 “9s”
- If that is not good enough for you...

Mission Risk:

Availability Rate from SEFIs Assuming 11 Second Recovery

Orbit	One SEFI	Two SEFIs	Three SEFIs	Four SEFIS
LEO (400 KM)	0.9999999993	0.999999998	0.999999998	0.999999997
LEO (800 KM)	0.9999999903	0.9999999806	0.9999999709	0.9999999612
Polar	0.999999992	0.99999999	0.99999998	0.99999997
CONST	0.99999997	0.99999999	0.99999999	0.99999999
GPS	0.999999997	0.999999994	0.999999991	0.99999999
GEO	0.999999996	0.999999993	0.99999999	0.99999999

- The worst case scenario (CONST) is 7 “9s”
- The best case scenario (LEO at 400KM) the maximum availability rate is 8-9 “9s”

Question

Is there anyone for which five to nine
“9s” is not good enough?

Next Question

Was our initial assumption that we could mitigate the rest of the SEUs reasonable?

Taking a Step Back

- **First we will discuss unmitigated circuits**
 - Let's get an idea of what is causing the errors and how these errors affect availability rates before moving onto mitigated circuits
- **Where Do Errors Occur?**
 - There are many different types of memory cells within the device each with their own sensitivity and consequence to radiation-induced faults, including configuration memory and user memory
- **Configuration memory controls much of the device:**
 - Defining the “equation” in LUTs
 - Defining the functionality of LUTs and user FFs
 - Defining the routing for the circuit
 - Defining the logical constants for the circuit
- **Designers have a few types of user memory available:**
 - FFs for pipelining data in the circuit
 - LUTs or BRAM configured as ROMs for data lookup – designers can approximate more complicated calculations by interpolating pre-calculated values stored in ROMs
 - LUTs or BRAM configured as RAMs for storing larger chunks of inflight data

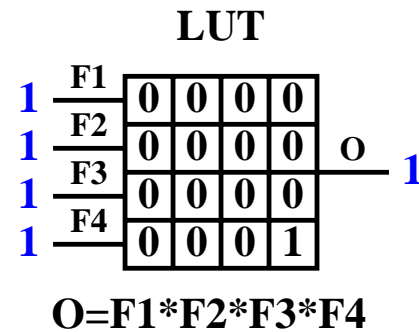
Configuration Memory:

Lookup Table Vulnerabilities

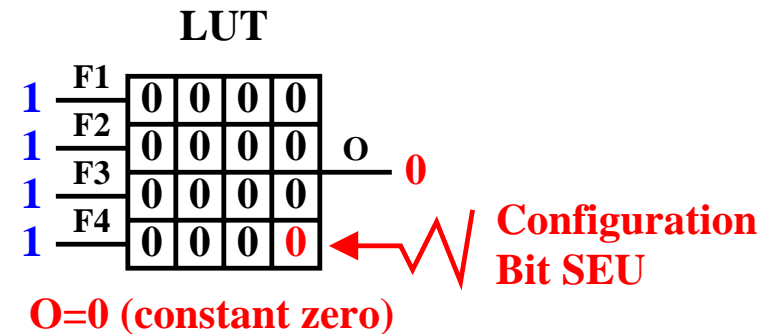
- **The logic in FPGAs is predominantly implemented in lookup tables (LUTs) which translate 2:1, 3:1, and 4:1 logic as a memory table with a decoder**
- **There are some embedded user cores in the device:**
 - Inverters
 - Multipliers/DSP units
 - DCMs
 - Processors
- **Under some circumstances the CAD tools will use LUTs instead of embedded cores:**
 - Not enough embedded cores were available
 - Inversion simplified into a LUT with another equation
- **Two types of LUT vulnerabilities:**
 - LUT equation changes
 - LUT “control” or “functionality” changes

Configuration Memory: LUT Equation Vulnerabilities

- **Configuration memory bits are used to store the LUT's values**
 - The LUT takes on a slightly different equation due to the changes
 - In the figures to the right, the 4 input AND gate equation is changed into a constant 0 equation
- **Bad news:**
 - Changes in the LUT equation can cause bad output data in unmitigated circuits
 - Errors from LUTs can cause bad circuit state
- **Good news:**
 - Except in cases of multiple-bit upsets, a LUT with an SEU in it is still correct for 15 out of 16 input combinations
 - Logic masking can cause output errors from one LUT to not become an output error for the circuit
 - Only 2-5% of upsets that occur in the V-4 device occur in the LUTs
 - At most 1.5% of the MBUs that occur in the V-4 device occur in the LUTs and mostly at very high LET heavy ions (very improbable)



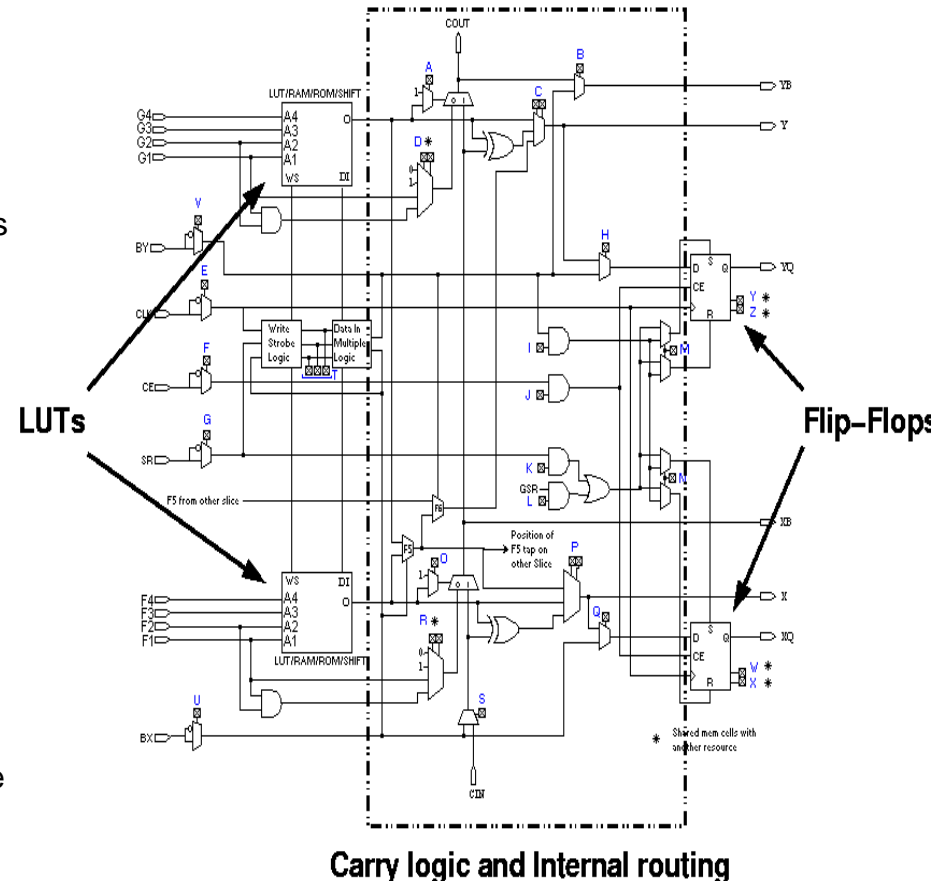
Original



After Upset

Configuration Memory: “Control” Bit Vulnerabilities

- **There are a number of configuration bits that are devoted to configuring the functionality of the slices:**
 - Inverting inputs
 - Using LUTs as LUTs, SRL16s, RAMs, or ROMs
 - Using user FFs as FDs, FDRs, FRSE, etc.
 - Using the fast carry chain
 - Etc.
- **Bad news:**
 - It is possible that the configuration memory bits devoted to state can be changed with an SEU
 - LUT turns into a shift register
 - FF loses the ability to be reset
 - Inverters not used
- **Good news:**
 - Only ~2% of upsets that occur in the V-4 device affect LUT functionality
 - At most 0.35% of the MBUs that occur in the V-4 device change the state in the slice and only with very high LET heavy ions




Configuration Memory: Routing Vulnerabilities

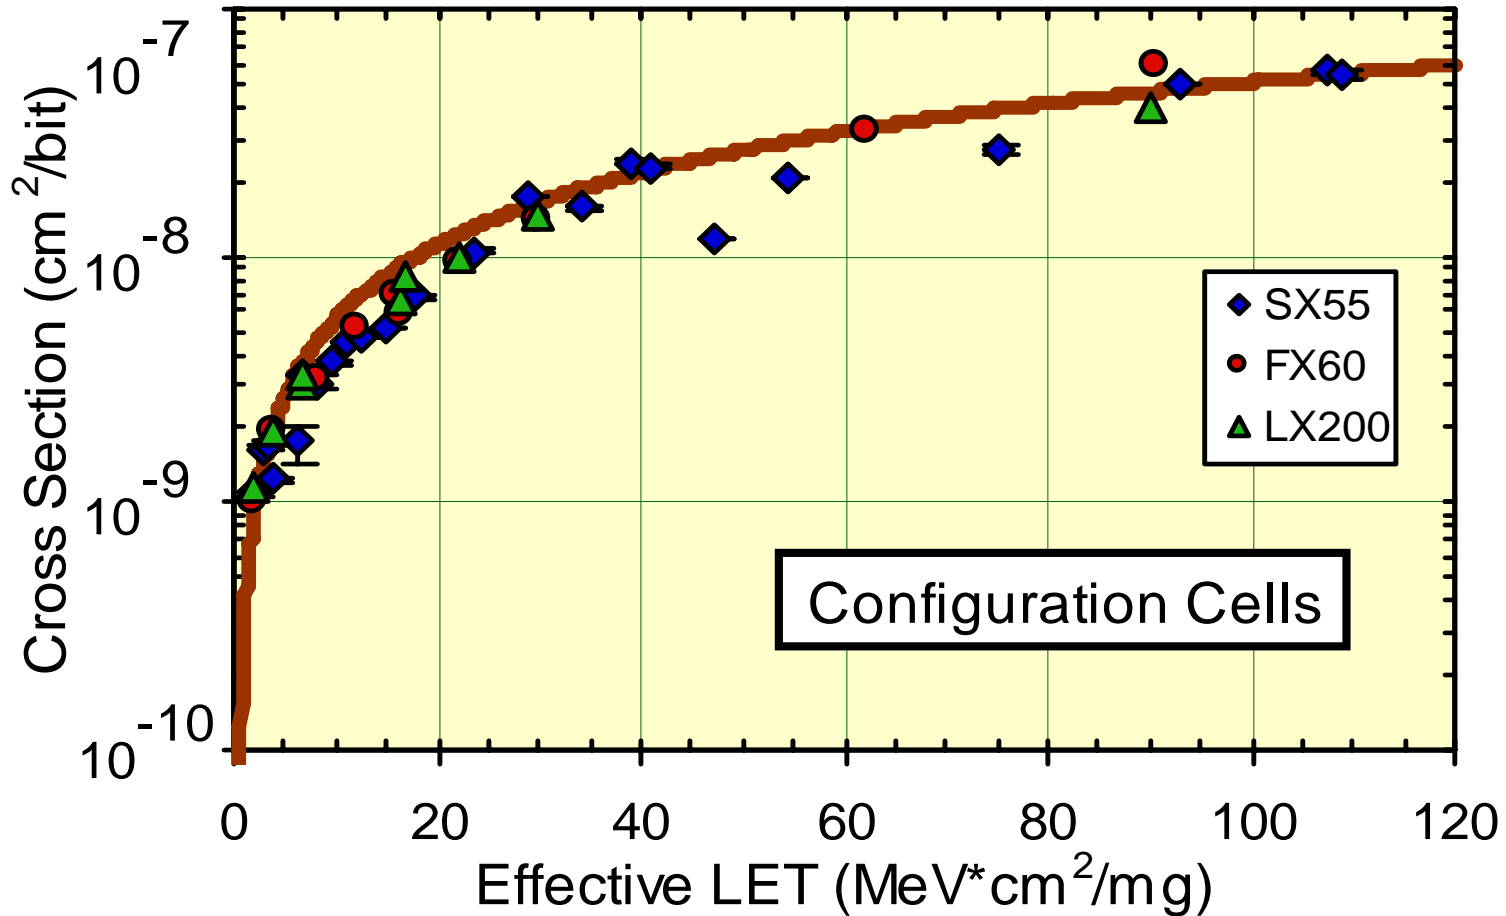
- **There are three main components in routing:**
 - Muxes, PIPs, and buffers
- **Routing comprises ~80% of the device**
 - Errors in routing affects routing data to LUTs, DSPs, Microprocessors, BRAM
 - Errors in routing affects routing global signal, such as clocks and resets
- **In analysis of failures in unmitigated designs 2/3rds of the “sensitive” cross-section (i.e., bits that when flipped cause noticeable errors in the output stream) is in the routing configuration memory**
 - Routing errors are not sensitive to input data
 - Corrupt routing is wrong no matter what the data is
- **Global signals are particularly vulnerable to SEUs**
 - Clock and reset trees can route to the entire device and SEUs can open or short the trees
 - Corrupting a global signal close to the input pin can affect the entire circuit
 - Corrupting a global signal near the leaves will have a more limited impact
 - Follow up research into domain crossing errors is showing that one of the vulnerabilities is that MBUs will switch the global signal routing – clocks from two domains switching, clocks and resets switching, etc.

Configuration Memory: Mux Vulnerabilities

- **Much of the routing in the Virtex-4 is mux-based.**
 - Muxes in the routing switches and the slices determine how the data is moved from point A to point B.
 - Routes are “defined” by moving data from one mux to the next mux until it reaches it’s destination
 - Muxes have a specific select line values stored in configuration memory that determines the input line on a route
- **Bad news:**
 - An SEU can change the configuration memory storing the select line values, causing the route to be driven by the wrong signal
 - Using a wire that is actively used by different logic (OMUX)
 - Using a wire that is being driven by a half latch, which imitates a stuck-at value
 - Opening or shorting a route
 - There are numerous muxes internal to each slice, muxes on every input and output of the LUTs and user FFs
 - SEUs in routing are 32-50% of all upsets in heavy ion on the V-4, based on energy
- **Good news:**
 - Relatively easy to protect by protecting the logic
 - Most of the routing is unused




Configuration Memory: Configuration Cell Upsets from Heavy Ions



From the Virtex-4QV Static SEU Characterization Summary: <http://www.xilinx.com/products/v4q/index.htm>

Configuration Memory: Configuration Cell Upsets from Protons



From the Virtex-4QV Static SEU Characterization Summary: <http://www.xilinx.com/products/v4qv/index.htm>

Configuration Memory: Configuration Cell Upset Rates

Configuration Upset Rates for Selected Orbits
per device-day, calculated by CREME96

Solar Minimum Quiet, AP8max, z=1 μm and 100 mils of Al

Orbit	Altitude (km)	Incl*	SX55	FX60	FX10	LX200	HI%
LEO	400	51.6°	0.73	0.69	1.61	2.03	69
	800	22.0°	7.56	7.12	16.7	21.1	2
POLAR	833	98.7°	6.02	5.67	13.3	16.8	22
CONST	1200	65.0°	23.3	21.9	51.6	65.1	5
GPS	20200	55°	4.08	3.79	8.92	11.2	6
GEO	36000	0°	4..28	4.03	9.5	11.9	94

* Incl = Inclination HI% = fraction from heavy ions

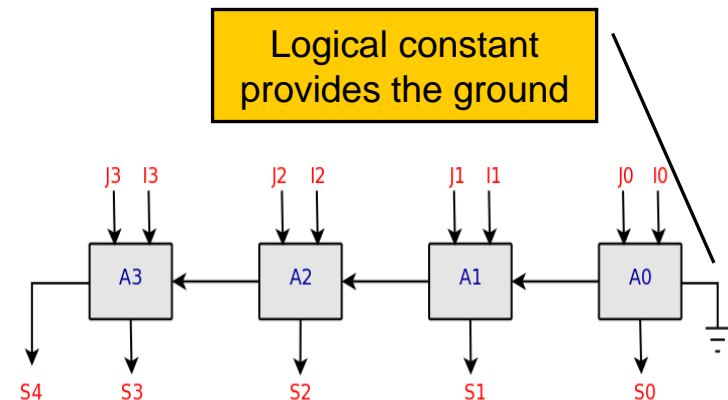

SMAP+ = SMAP & FAR SEFIs combined

From the Virtex-4QV Static SEU Characterization Summary: <http://www.xilinx.com/products/v4qv/index.htm>

UNCLASSIFIED

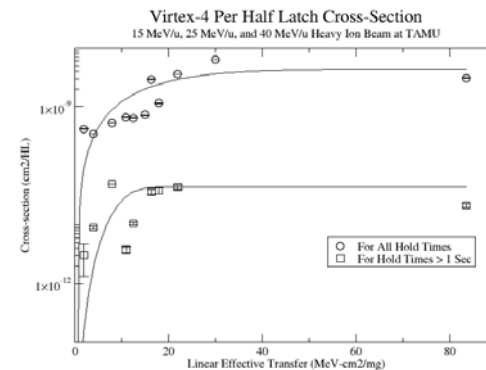
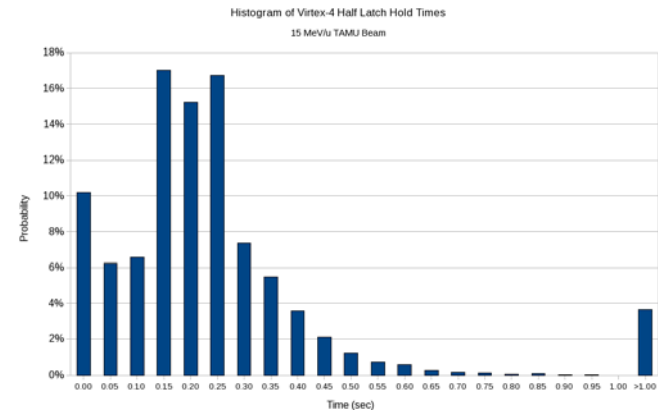
Configuration Memory: Logical Constants

- **Logical constants are needed to generate constant zero and one logic values used internally by FPGA designs**
 - Artifact of mapping VHDL designs to the specific FPGA architecture
 - Not under design-control, unless the designer is going to extraordinary measures to avoid them in VHDL/Verilog
 - Easy to mitigate at either the EDIF or XDL level
- **“Implicit” logical constants**
 - Inputs to I/O, logic, RAM, clocking, and other resources
 - Implemented in half latches (weak keepers)
- **“Explicit” logical constants:**
 - Tie-offs to the zeroth bit of the carry chain for adders and unused multiplier/DSP inputs
 - Implemented as constant LUTs in the Virtex-I and Virtex-II, implemented as architectural posts in the Virtex-4



Configuration Memory: Half Latch Data for Virtex-4

- **Multi-modal data**
 - Peaks at 0.04, 0.16, 0.18, 0.19, and 0.26 secs
- **The average time that a HL holds is 0.23 secs for the entire data set with a standard deviation of 0.13 secs such that 68% of all half latches leak off within 0.10-0.36 secs**
- **On average 96.4% of all half latches leak off within 1 sec**
- **Good news:**
 - Per-HL cross-section for all HLs is 2-3 orders of magnitude smaller than per-bit cross-section of configuration bits
 - Per-HL cross-section for HLs that hold for longer than one second is 5-7 orders of magnitude smaller than the per-bit cross-section of configuration bits
 - Can extract to posts using MAP/PAR




Configuration Memory: Architectural Posts Vulnerability – None!

■ Architectural Posts

- In the Virtex-II, these posts are an abstraction of a half latch and have the reliability concerns of a half latch
- In the Virtex-4, post ties to ground rail

■ Very good news:

- So far, no known single-bit failure modes associated with post
- No known configuration bit associated with post
- Since post only supplies constants for slices within the CLB, routing between post and slices is within the local switch box. Route sharing within local switch box not observed.
- Using posts for extracting half latches can keep the logical constant network local and not global

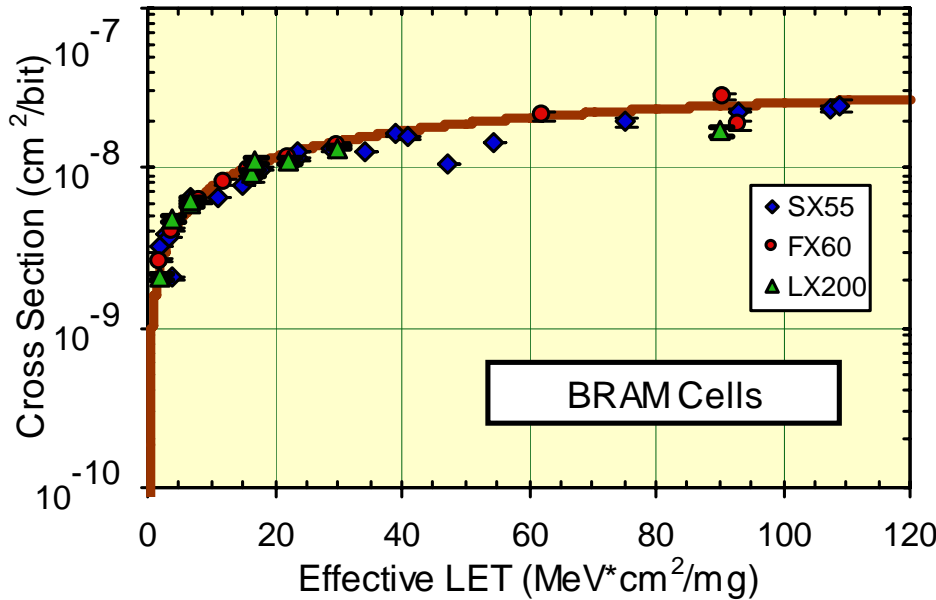


User Memory

- **SEUs in these memory cells are a concern and lead to a corruption of circuit state**
 - SETs in the device would be visible through a change in the user FF cross-section
 - No published data on SETs in the user FFs
- **SEUs in user memory are difficult to mitigate**
 - Mitigate the logic attached to user FFs
 - Triplicate the BRAM
- **User memory that can be written to cannot be scrubbed traditionally without corrupting the contents of the memory**
 - Use Xilinx's BRAM scrubber to scrub user memory in BRAM
 - Mitigate the logic around the other user memory so there is no need to scrub

User Memory:

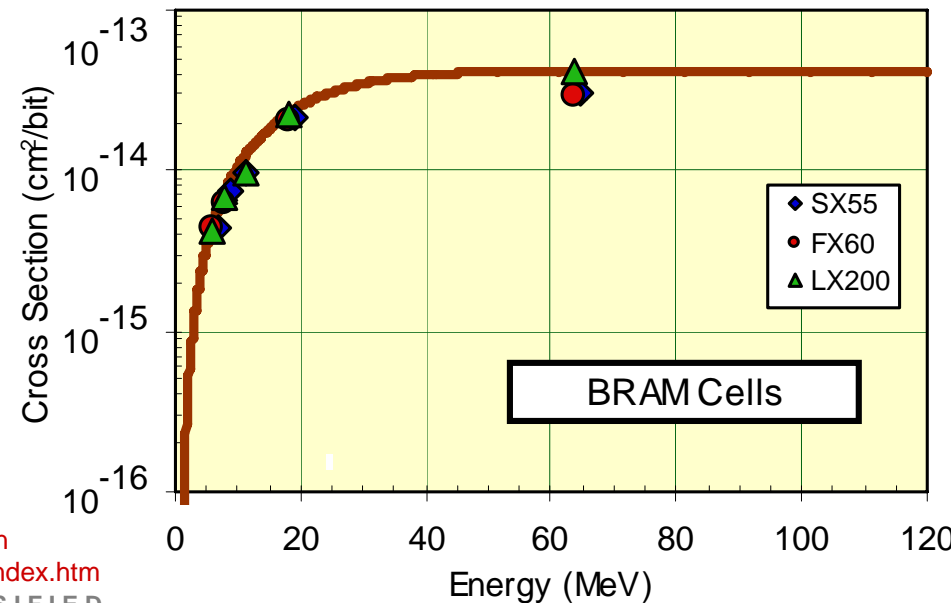
BRAM Cell Upsets from Heavy Ion and Proton



Heavy Ion

Both heavy ion and proton saturated cross-sections are comparable to the Actel SRAM results

Proton



User Memory: BRAM Cell SEU Upset Rates

BRAM Upset Rates for Selected Orbits
per device-day, calculated by CREME96

All Bits Used, Solar Minimum Quiet, AP8max, z=1 μm and 100 mils of Al

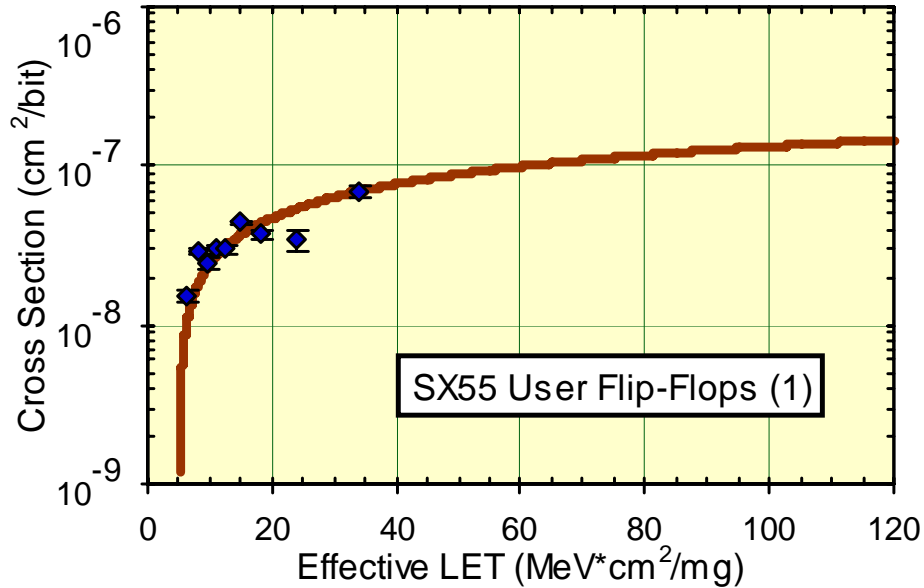
Orbit	Altitude (km)	Incl*	SX55	FX60	FX10	LX200	HI%
LEO	400	51.6°	0.72	0.52	1.24	0.75	84
	800	22.0°	4.05	2.94	6.99	4.25	5
POLAR	833	98.7°	4.00	2.90	6.90	4.20	37
CONST	1200	65.0°	13.3	9.63	22.9	13.9	10
GPS	20200	55°	4.02	2.92	7.43	4.22	2
GEO	36000	0°	4.49	3.26	7.75	4.71	98

* Incl = Inclination HI% = fraction from heavy ions
SMAP+ = SMAP & FAR SEFIs combined

From the Virtex-4QV Static SEU Characterization Summary: <http://www.xilinx.com/products/v4qv/index.htm>

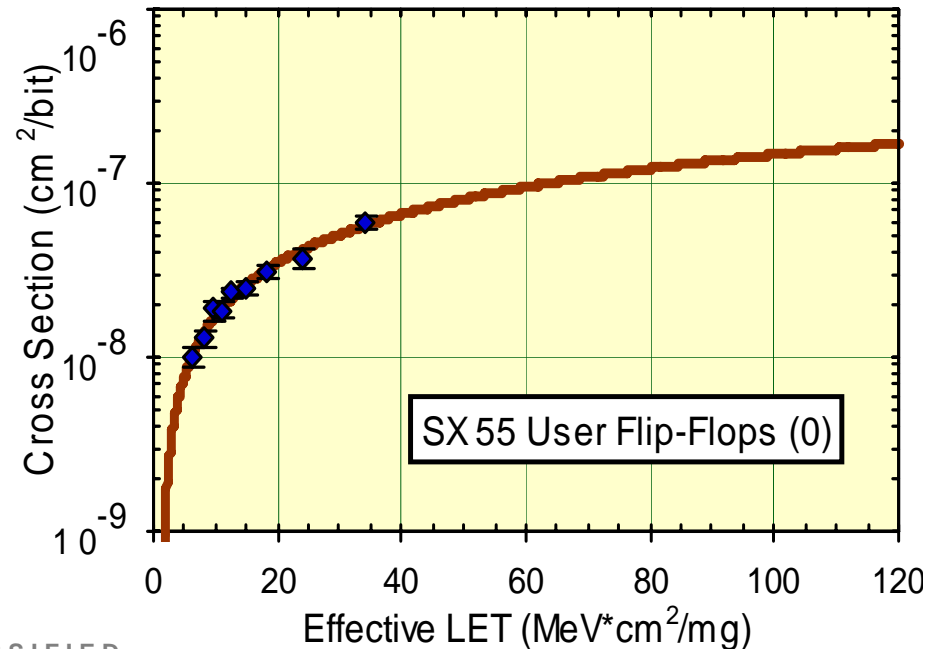
UNCLASSIFIED

User Memory: User Flip-Flop Test Results



Results are comparable to
the Actel D-Latch saturated
cross-section results

From the Virtex-4QV Static SEU
Characterization Summary:
<http://www.xilinx.com/products/v4qv/index.htm>



User Memory:

User Flip-Flop Upset Rates

Flip-Flop Upset Rates for Selected Orbits
 # per device-day, calculated by CREME96

Full, 50% ones, Solar Minimum Quiet, AP8max, z=1 μm and 100 mils of Al

Orbit	Altitude (km)	Incl*	SX55	FX60	FX10	LX200	HI%
LEO	400	51.6°	0.007	0.007	0.017	0.024	70
	800	22.0°	0.070	0.072	0.18	0.25	2
POLAR	833	98.7°	0.057	0.059	0.15	0.21	23
CONST	1200	65.0°	0.22	0.23	0.57	0.81	5
GPS	20200	55°	0.0364	0.0374	0.0935	0.1319	5
GEO	36000	0°	0.039	0.048	0.099	0.14	96

* Incl = Inclination HI% = fraction from heavy ions
 SMAP+ = SMAP & FAR SEFIs combined

Mission Risk:

Errors in the Output Data Stream

- **In unmitigated circuits approximately only 1-20% of the device will cause a noticeable output error from the user circuit**
 - These numbers are based on fault injection and beam testing using random input data
 - These numbers are based on designs that do not have mitigation applied to user circuit
 - Approximately 1/3rd of errors are in the logic and 2/3rd are in routing
 - These numbers are also design-dependent, which means that testing will be necessary to determine the sensitivity of your design to output errors
- **Each circuits has its own inherent sensitivity to errors**
 - Many digital signal processing applications are very insensitive to errors
 - Circuits with a lot of feedback loops, state and where the output data is “well-tied” to the input data are more sensitive to errors
- **It takes 5 seconds for the scrubber to detect the error in the bitstream and fix it**
 - Resynchronizing the circuit could be very quick or very slow depending on the circuit

Error Rates for Unmitigated Circuits

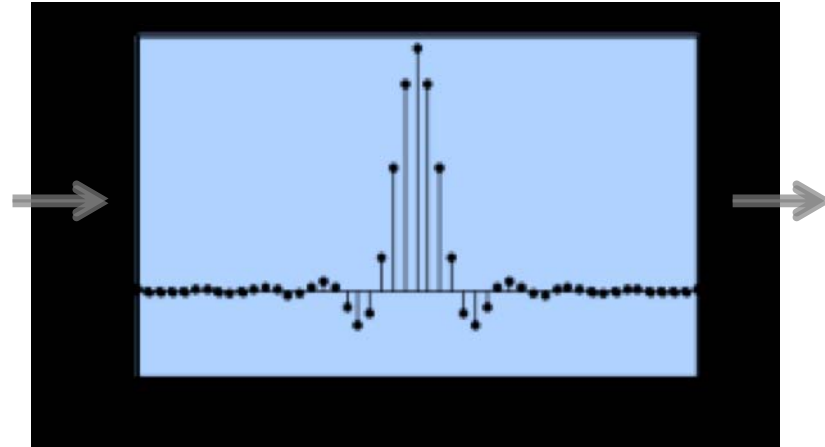
- The LX200 is a 51Mb device
- When there is more 60 errors per year, the availability rate will fall below the 5 “9”s mark
 - Small circuits in High LEO, Polar and CONST will need to be mitigated
 - Nearly all large circuits in all orbits will need to be mitigated

Orbit	1% -- upsets per year	1% -- lost seconds per year	20% -- upsets per year	20% -- lost seconds per year
LEO (400KM)	7.41	37.05	148.19	740.95
LEO (800 KM)	77.02	385.08	1540.30	7701.5
Polar	61.32	306.60	1226.40	6132.00
CONST	237.62	1188.08	4752.30	23761.5
GPS	40.88	204.4	817.60	4088
GEO	43.44	217.18	868.70	4343.5



Example of a Digital Signal Processing Circuit

- Brian Pratt from BYU studied the effect of SEUs and input data noise on a FIR filter
- 49 taps
- 24 multipliers (symmetric coefficients)
- Square-root raised cosine (SRRC) pulse shape with 50% rolloff
- 16-bit fixed-point input (Q2.14 format)
- 18-bit fixed-point output (Q4.14 format)
- 15% of Slices occupied on Virtex 1000 FPGA
- Total sensitive configuration bits: 149,696/5,810,024







Example of an Unmitigated DSP Circuit: Results Table

Input SNR	Less than 0.1dB loss in SNR	Less than 1dB loss in SNR	Less than 3dB loss in SNR	Less than 6dB loss in SNR
No noise	69,160 trials (46.2%)	81,419 trials (54.4%)	89,619 trials (59.9%)	95,134 trials (63.6%)
20 dB	121,370 trials (81.1%)	129,223 trials (86.3%)	133,441 trials (89.1%)	136,230 trials (91.0%)
10 dB	128,741 trials (86.0%)	135,997 trials (90.8%)	139,586 trials (93.3%)	142,135 trials (94.9%)
5 dB	132,484 trials (88.5%)	139,126 trials (92.9%)	142,230 trials (95.0%)	143,825 trials (96.1%)

- Total trials: 149,696
- Number of sensitive configuration bits in the design

Example of an Unmitigated DSP Circuit: Application-specific Cross-section

FPGA
Virtex 1000



**Full FPGA
(static cross section):**
12,288 Slices
**5.8 Million config bits
(100%)**

**FIR Filter
(dynamic cross
section):**
1,869 Slices
**149,696 config bits
(2.5%)**

**FIR Filter in a 20dB
SNR environment
tolerating 1dB
additional SNR loss:**
**20,473 config bits
(0.35%)**

Question

What does this mean in terms of error rates?

Example of an Unmitigated DSP Circuit: Error Rates

- 20,473 configuration bits caused errors in the system
- The LX200 is a 51Mb device
- Depending on the orbit, the circuit could go between 39 and 1235 days between noticeable output errors

Orbit	Per Device-Day	Per FIR-day	Days Between Noticeable Upsets
LEO (400KM)	2.03	0.0008	1235
LEO (800 KM)	21.1	0.0084	119
Polar	16.8	0.0067	149
CONST	65.1	0.0260	39
GPS	11.2	0.0045	224
GEO	11.9	0.0047	211


Example of an Unmitigated DSP Circuit:

Availability

- **SEU Detection:**
 - Done in hardware with a scrubber running continuously at full system rate (32-b SelectMAP at 33 MHz)
 - No software overhead except to initialize and start
 - < 5s – probably much faster
- **For this unmitigated design between 1.48 to 47.39 seconds will be lost each year to SEUs**
 - Circuit easily meets the 5 “9s” criteria

Orbit	Lost Seconds/year	Availability/year
LEO (400KM)	1.48	0.9999999531
LEO (800 KM)	15.36	0.9999995129
Polar	12.23	0.9999996122
CONST	47.39	0.9999984973
GPS	8.15	0.9999991759
GEO	8.66	0.9999997253


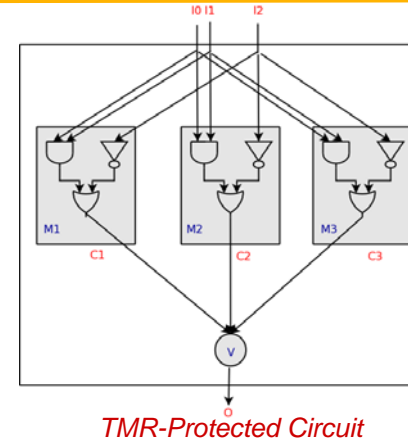

The First Thing to Keep in Mind: SRAM FPGA Radiation Effects Characterization Error Rates




Moving Forward: Mitigating Circuits

- **“Gold standard” for “perfect” triple-modular redundancy is**
 - Triplicated input and output data streams
 - Triplicated clock, resets, and other global signals
 - Triplicated logic
- **A fully TMR-protected FPGA design should mask all single-bit SEUs as long as there is only one in the system at a time.**
 - Scrubbing must ensure that all upsets are removed and the circuit resynchronized before the next upset occurs.
 - TMR is not as successful with either multiple-bit upsets or multiple independent upsets.
- **While the concept of TMR is simple, the implementation of TMR in FPGA designs is often not simple.**
 - The circuit description could vary widely from the circuit implementation.
 - A number of scenarios exist that can affect the reliability.

Circuit Design Influences in TMR-Protected FPGA Designs: Redundant Modules Removed

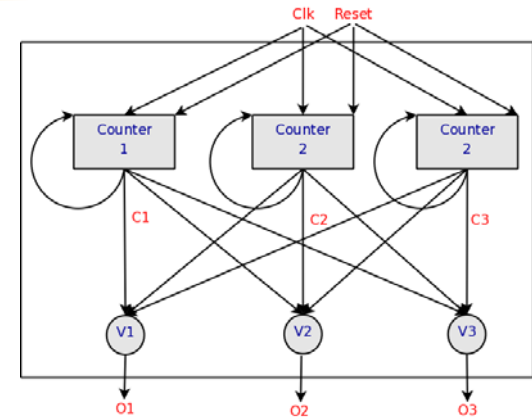


- **Synthesis tools remove redundant modules to optimize the circuit for speed and area.**
 - Common when inputs and outputs are single sourced.
- **Synthesized circuit is no long protected by TMR.**
 - Remaining voters increase the sensitive cross-section.




Circuit Design Influences in TMR-Protected FPGA Designs: Feedback Loops Improperly Protected

- **Technically, both counters are TMR-protected.**
- **The top counter has a persistent cross-section.**
 - If one SEU causes one of the counters to output bad data, the TMR-protected circuit will not output bad data.
 - Because the feedback loop is not cut with a voter, the broken counter will not resynchronize after the SEU is removed.
 - If another SEU causes another counter to output bad data, the TMR-protected circuit will output bad data.
- **The bottom counter does not have a persistent cross-section**
 - In this case, the counters feedback through voters.
 - Even if one of the counters is outputting bad data, all of the counters will feedback the correct data.



Counter with Persistent Cross-Section



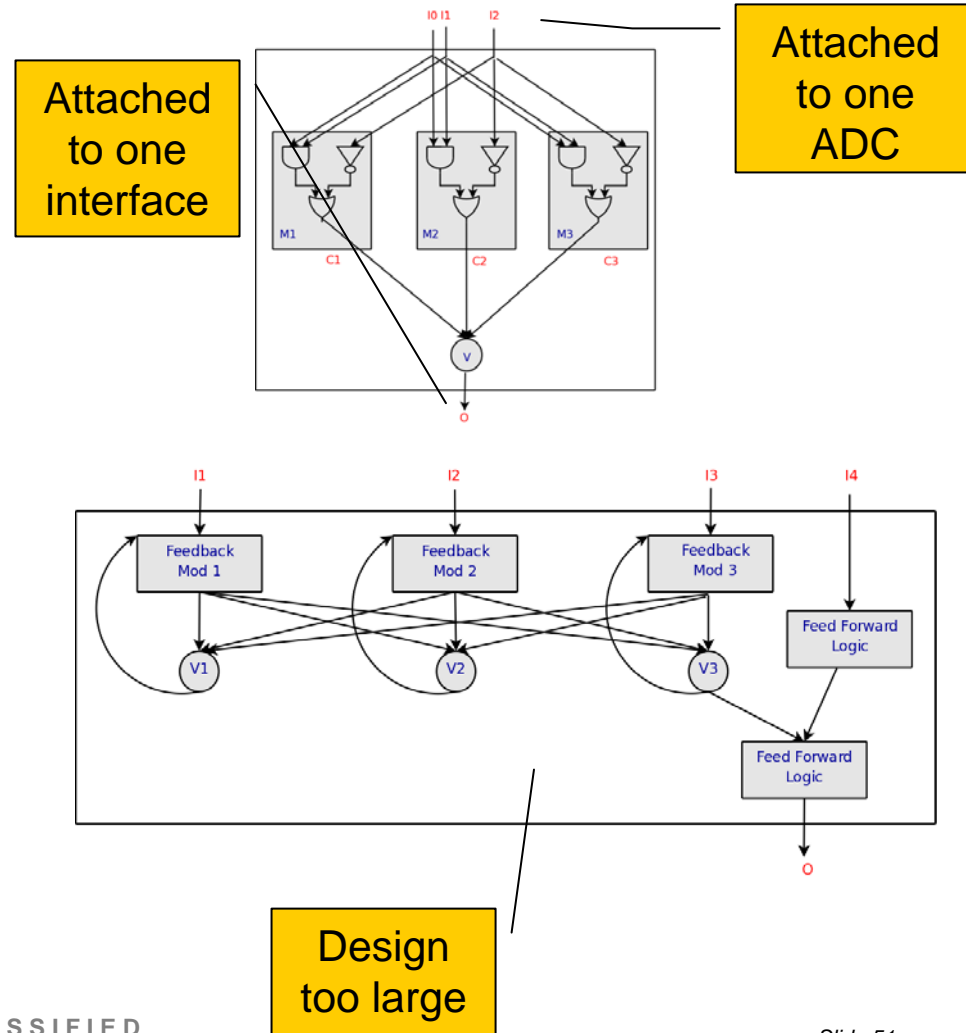
Counter without Persistent Cross-Section

Correct Circuit Design for TMR-Protected FPGA Designs

- **Do not manually apply TMR in VHDL.**
- **Use either the Xilinx TMRTool and BL-TMR (from Brigham Young University) to apply TMR.**
 - Both automatically apply TMR to post-synthesis circuit representations, called EDIF.
 - After synthesis, major circuit optimizations will not occur.
 - Less likely to lose TMR-based redundant modules
 - Feedback loops properly cut
- **Do not expect TMR to solve any existing problems in the design.**
 - If your design cannot meet timing, has design flaws or is really large without TMR, applying TMR will not fix these problems.

Circuit Design Influences in TMR-Protected FPGA Designs: Design Constrained Scenarios that Affect TMR

- Sometimes designers are unable to fully triplicate the design due to either area or pin issues.
- Triplicating input/output signals can be impossible due to pin constraints and can be difficult to manage due to skew.
- Triplicating all of the logic might not be possible due to the chosen device's size.
 - BL-TMR can automatically apply partial TMR for this scenario through prioritized redundancy based on device size.
- Without full triplication of logic and signals, some unprotected cross-section will exist.



How Do You Know Your Mitigation Scheme Is Working?

- **Model, fault inject and/or radiation test the design**
- **Current “gold standard” is to do pre-launch testing of user designs through radiation experiments at a particle accelerator.**
 - Space-qualifying a design could take days worth of time and thousands of dollars at an accelerator.
 - Radiation-induced faults are statistical in nature which further complicates the time and expense of radiation-experiments
 - Hard to correlate errors to flaws in the user design.
- **Designers need faster, cheaper and more uniform methods of testing user designs.**
 - Modeling tools and fault injection tools can be useful in these regards, and
 - Radiation experiments used only to validate these results.
- **If you do not have a fault injection tool or a modeling tool, ask us to help you get one**
 - Not only incredibly helpful but cheaper in the long run

Scrubbing (1 of 2)

- **Memory scrubbing is a standard practice with memory devices to remove detectable errors from memory so that errors do not accumulate on the device**
 - Scrubbing is a necessary part of a good mitigation scheme
- **For FPGAs, this process is done by doing partial or on-line reconfiguration**
 - In the Virtex-4, it is possible to do a on-line reconfiguration of the device without taking the circuit down
 - On-line configuration provides the ability to operate through SEUs in the configuration memory
- **On-line reconfiguration of the user FFs, SRL16s and BRAM are not possible**
 - In the Virtex-4, use of the GLUTMASK allows the scrubber to “jump over” the SRL16s
 - Xilinx’s BRAM Scrubber can be used to remove errors in the BRAM

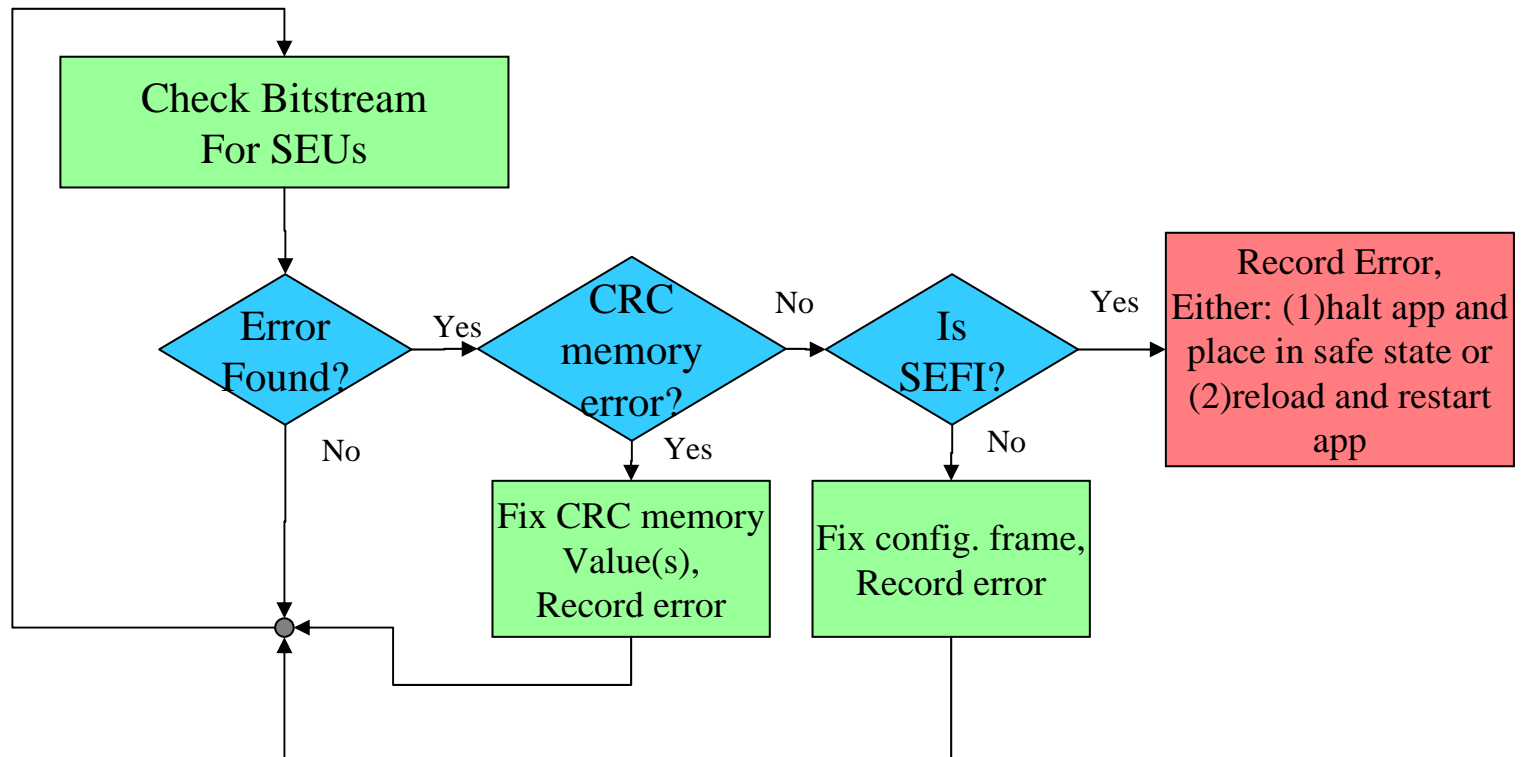
Scrubbing (2 of 2)

- **There are many different types of scrubbers:**
 - Internal scrubbers using the ICAP interface
 - External scrubbers using the JTAG or SMAP interface
- **Most external scrubbers can also provide SEFI detection and recovery**
 - Most SEFIs require an off-line reconfiguration of the device, which affects the operation of the circuit
- **The current recommendation from Xilinx is that so called “blind” scrubbing should not be done**
 - While the Scrub SEFI does not have a very large cross-section, it still exists and could destroy your device in the event that it occurs
 - Blind scrubbing also does not provide a manner for determining whether the device is scrubbing properly
- **What follows is a description of the LANL V4 External Scrubber that is based on the idea of on-demand reconfiguration**

LANL Virtex-4 Scrubbing Methodology

- **Current approach is based on the proven approach used by LANL's Cibola Flight Experiment for Virtex-I (much of the same HDL code and approach)**
 - Uses pre-calculated CRCs for each configuration data frame to identify any changes in the configuration data
 - Only the frames with errors are refreshed or “scrubbed” (the rest of the bitstream is left alone—i.e., partial reconfiguration is used for only the frames in error)
 - Repeated errors or large numbers of frame errors are used to identify SEFI conditions
 - Accounts for errors in the “rad-hard” SRAM holding the CRCs

Virtex-4 Scrubbing Algorithm Overview



Virtex-4 Configuration Data SEU Detection Algorithm (Hardware)

I. While scrubbing enabled, do:

- A. Readback the configuration bitstream data for the Logic (CLK, CLB, DSP, etc.), BlockRAM interconnect and “BRAM remapping register” areas of the device
 - 1. For each 1312-bit frame
 - a. Calculate the 32-bit CRC for the frame
 - b. Compare the actual CRC to the pre-calculated expected CRC stored in “rad-hard” SRAM
 - c. If the expected and actual CRCs are not equal
 - i. If it is the first error detected, record the expected and actual frame CRCs and the frame number
 - ii. Increment the frame error count
- B. If a frame error was found, halt; otherwise, continue

Notes on SEU Detection Approach

- **Only reports the complete information for the first frame in error per readback (the most frequent case)**
 - If more than one frame is in error, the SEU detector is restarted to get the information for the next frame in error.
- **Reports total number of frames with errors per readback**
- **Performs a complete readback of logic/BRAM inconnect region and then a complete readback of “BRAM Remapping Register” region**
 - Each region is read completely without interruption to simplify control
- **Abort sequences are used before most if not all SelectMAP transactions to ensure the SelectMAP interface is in a known state before performing a function (readback, configuration, partial configuration, etc.)**

CRC Memory Errors

- **Though not frequent, it is expected that we will encounter occasional CRC memory errors.**
 - This has been observed on the Cibola Flight Experiment currently on orbit.
- **Simple error check approach**
 - XOR the expected and actual CRCs
 - If they differ by only a few bits (1 or 2, e.g.), then the most likely cause is an error in the CRC memory itself rather than an error in the FPGA's configuration data memory
 - This is true due to the nature of CRCs: single bit differences in the configuration data will lead to large changes in the CRC not incremental changes
 - This can be done for every SEU detected due to the low cost in software and due to the relatively high cost of other approaches (waiting to see the error X times, e.g.).

SEFI Detection

■ Simplified approach

- If an error persists for N SEU detection cycles (and it isn't a CRC memory error), then a SEFI has occurred.
- If a M errors have appeared during a single SEU detection cycle, then a SEFI has occurred.

■ Current thresholds are N=5 and M=1000. These thresholds are parameters and can be changed

- N was chosen somewhat arbitrarily: large enough not to be triggered frequently, but small enough to ensure the SEFI is identified within a reasonable time frame.
- M=1000 was chosen based on the number of frame errors we saw for SEFIs during proton radiation testing (small enough to account for all observed SEFIs).

How Do You Know If the Scrubber Is Working?

Unless you bought a scrubber off of someone, you need to take it to a radiation facility and test it.

Mission Risk:

Results of Mitigating the Virtex-4

- We have found that there is around 100 bits in the Virtex-4 that are causing single points of failure in BL-TMR-protected designs
 - Some of these are in the IOB, when two domains share the same pad
 - Some of these are in the CLBs – working with Xilinx to resolve these issues
- Even with the remaining 100 bits the Virtex-4 will be able to make the 5 “9s” availability rate

Orbit	Per Device-Day	<u>Years</u> Between Noticeable Upsets	Lost Seconds Per Year for Mitigated Circuits
LEO (400KM)	.07	888.57	.0056
LEO (800 KM)	7.56	82.27	0.0608
Polar	6.02	103.32	0.0484
CONST	23.3	26.7	0.1873
GPS	4.02	154.7	0.0323
GEO	4.28	145.3	0.0344

Let's Turn the Question Around

**How many sensitive bits can remain
and still make 5 “9s”?**

How Many Unprotected Bits Can You Have and Still Make 5 “9s”?

- Assuming it takes five seconds to mitigate an SEU, then no more than 60 noticeable errors can occur a year

Orbit	Bits
LEO (400KM)	~5M
LEO (800KM)	~500K
Polar	~600K
CONST	~150K
GPS	~900K
GEO	~850K

Let's Ask That Questions Again

Are our assumptions sound?

Caveats

- **There could be multiple-bit upsets that cause TMR to fail (“Domain Crossing Events”)**
 - Still getting our hands around it
 - Very likely single-bit upsets will continue to dominate for the next few generations
- **There could be multiple-independent upsets that cause TMR to fail**
 - How likely is that?



Harsh Radiation Environments


- **LANL and BYU studied whether MIUs could be a possible failure mode**
 - Studied the worst possible error rates – solar energetic particle (SEP) events
- **SEP events (colloquially known as solar flares)**
 - Harshesht radiation environment in space Earth orbits (Tylka et al.)
 - SEU rate can increase orders of magnitude
 - Modeled in CREME96 by week-long event in October 1989
 - Primarily affect orbits further away from the Earth's magnetic field or at the poles (e.g. GPS, GEO, Molniya, Polar)
- **An increase in SEU rate increases the likelihood of MIUs**
- ***The purpose of this work was to measure the probability of failure from MIUs during SEP events***




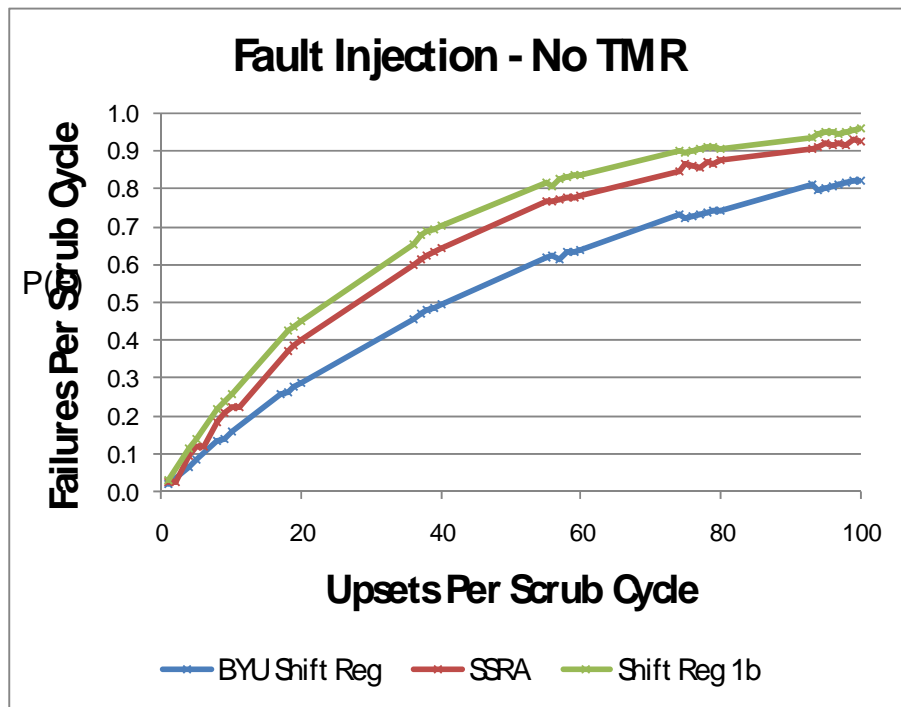
Mission Risk:

Harsh Orbits Example Circuits

- **BYU Shift Reg**
 - 32-bit wide, 250 stage shift register with arbitrary combinational logic between each stage
- **Shift Reg 1b**
 - 1-bit wide, 16,200 stage shift register
- **SSRA**
 - Digital signal processing kernel





Mission Risk: Fault Injection Results




Note: A minimum of 1000 events were observed for each data point.

Mission Risk: Accelerator Results



Note: A minimum of 30 events were observed for each data point.

Mission Risk: Fault Injection Comparison





Mission Risk:

SEU Rate Models



Orbit	Apogee (km)	Perigee (km)	Inclination (deg)	Peak 5-Minute SEU Rate (SEUs/Device/s)
GEO	35786	35786	0	3.29E-1
GPS	20200	20200	55	2.85E-1
Molniya	39305	1507	63.2	3.08E-1
Polar	833	833	98.7	7.84E-2

■ SEU Rate Models

- CREME96: Worst Week, Worst Day, Peak 5-minutes (see data above)
- What actually happens during that peak 5-minutes?
 - Real Answer: Nobody Knows!!
 - But we can bound the problem...
 - Best Case: Flux is perfectly averaged over 5 minutes
 - Worst Case: Flux all comes during one scrub cycle
 - We can also guess by extrapolating from worst week, worst day and peak 5 minutes



Mission Risk: Estimating Probability of Failure (CREME96 Best Case)



$$P(F_{1-cycle}) = \sum_{i=0}^{\infty} P(F_{1-cycle} | A_i) P(A_i)$$

$$= 9.7 \times 10^{-8}$$

$$P(F_{5-min}) = 1 - (1 - P(F_{1-cycle}))^m$$

$$= 1 - (1 - 9.7 \times 10^{-8})^{20000}$$

$$= 1.9 \times 10^{-3}$$

UNCLASSIFIED



Mission Risk:

Estimating Probability of Failure

Design	Orbit	P(F _{5-min})			P(F _{worst-day})	P(F _{worst-week})
		Best Case	Log Fit Model	Worst Case		
SSRA	GEO	*	1.97x10 ⁻³	.987	3.3x10 ⁻²	1.8x10 ⁻²

* 1.94x10⁻³ if using average λ from logarithmic fit model

* 1.99x10⁻³ if using actual CREME96 average λ

- The log fit model provides little to no extra fidelity beyond the best case (average SEU rate value directly from CREME96)
- There is a factor of 500 difference between the best case and worst case
- These are worst case numbers
 - The artificially worst, worst case rate states that the probability of failure is high
 - The best worst case rate is actually very reasonable

Mission Risk:

The Comparison So Far....

- **The Actel Flash-based RT ProASIC is probably not useful for multi-year missions due to low TID and will have many problems with SEUs and SETs**
 - Let's take that device off the table for now
- **The Xilinx Virtex-4 and the Actel Anti-Fuse-based RTAX have very comparable TID and SEL-immunity**
 - The anti-fuse device will have fewer problems with SEUs than the Virtex-4
 - The Xilinx Virtex-4 can be mitigated to meet the 5 “9s” requirement
 - MBUs not likely to be a concern
 - MIUs not likely to be a concern
- **Still need to look at device size and reconfigurability**

Mission Risk: Comparing Devices by Size

Device	XQR4VLX200	XQR4VSX55	XQR4VFX60	XQR4VFX140
Logic Cells	200K	55K	57K	142K
FFs	178K	49K	51K	126K
BRAM (kbits)	6048	5760	4176	9936
DSP	96	512	128	192
PPC	0	0	2	2
I/O	960	640	576	896

Device	RTAX250	RTAX1000	RTAX2000	RTAX4000
Equivalent System Gates	250K	1M	2M	4M
R-Cells	1,408	6,048	10,752	20,160
C-Cells	2,816	12,096	21,504	42,840
FFs	2,816	12,096	21,504	42,840
RAM (kbits)	54	162	288	540
I/O	248	516	684	840

- There are four device sizes both the RTAX-S/SL [1] and the Virtex-4 [2]
- Comparing FFs is a bad comparison, as often the logic cells will be completely utilized and the FFs will not be – logic is the important comparison
- The Virtex-4 “logic cells” are roughly equivalent to the Actel “C-Cells”
 - The largest RTAX device has 2.5 times fewer logic cells than the smallest Virtex-4 device
 - The Virtex-4 device has DSP units and PowerPCs, which can offload some of the logic into embedded cores
- Even with mitigation the designer can put more logic in the Virtex-4 device

Mission Risk: Reconfigurability

- **There are many missions for which reconfigurability is very attractive**
 - Supporting multiple missions with one satellite/payload
 - Supporting changing or emerging missions or science needs
 - Supporting graceful degradation of satellite sensors
- **Reconfigurability can also save you from**
 - Design errors
 - Launch problems
- **The RTAX-S/SL cannot provide reconfigurability – what you launch is what you get**

Conclusions:

The Comparison....

- **The Actel Flash-based RT ProASIC is probably not useful for multi-year missions due to low TID and will have many problems with SEUs and SETs**
 - Let's take that device off the table for now
- **The Xilinx Virtex-4 and the Actel Anti-Fuse-based RTAX have very comparable TID and SEL-immunity**
 - The anti-fuse device will have fewer problems with SEUs than the Virtex-4
 - The Xilinx Virtex-4 can be mitigated to meet the 5 “9s” requirement
 - MBUs not likely to be a concern
 - MIUs not likely to be a concern
 - The Xilinx Virtex-4 can be mitigated to meet the 5 “9s” requirement
 - MBUs not likely to be a concern
 - MIUs not likely to be a concern
- **The Virtex-4 provides more available user logic**
- **The Virtex-4 can meet changing mission, sensor and satellite needs**

Questions?
