Soft Error Study of ARM SoC at 28 Nanometers

Austin Lesea, Wojciech Koszek, Glenn Steiner, Gary Swift[1], and Dagan White
Xilinx, Inc.

***Abstract*— During the past year Xilinx, for the first time ever, set out to quantify the soft error rate of a multi-core microprocessor. This work extends on Xilinx's 10+ years of heritage in FPGA radiation testing. Built on the 28 nanometer technology node, Xilinx's Zynq™ family of devices integrate a processor subsystem with programmable logic. The processor subsystem includes two 32 bit ARM Cortex™-A9 CPU's, two NEON™ floating point units, two SIMD processing units, an L1 and L2 cache, on chip SRAM memory and various peripherals. The programmable logic is directly connected with the processing subsystem via ARM's AMBA™ 4 AXI interface. This programmable logic is based on the 7 Series FPGA fabric, consisting of 6-input LUTs and DFFs along with Block RAM, DSP slices, multi-gigabit transceivers, and other blocks. Tests were performed using a proton beam to analyze the soft error susceptibility of the new device. Proton beam testing was deemed acceptable since previous neutron beam and proton beam testing had shown virtually identical cross-sections for 7 Series programmable logic. The results are promising and yield a solid baseline for a typical embedded application targeting any of the Zynq SoC devices. As a foray into processor testing, this Zynq work has laid a solid foundation for future Xilinx SoC test campaigns.**

***Index Terms*—alpha particle upset, architectural vulnerability factor, cross-section, FIT rate, FPGA, SoC, Processor, NSEU, proton beam testing, silent data corruption, single event upset, soft error, SRAM.**

I. Introduction

Determining the SEU rate, or SEU cross-section, for a processor system is relevant for safety critical systems developers. SEU rates are a challenge in design of automotive systems, medical devices, avionics equipment, and high speed rail systems, amongst others. Each of these markets has applicable standards for safety and reliability yet methods for quantifying and analyzing the soft error rates in processors is not well defined. There are no standard methods for testing soft error rates in processor devices. Such a task is made difficult by the lack of published information on how such neutron strikes affect the processor system and its peripherals, and the methods used to gather useful data.

The programmable logic in the system on chip (SoC) device must also be measured and understood so that appropriate mitigation steps may be taken to provide the required levels of reliability and availability for the complete SoC device with all intended functionality. The programmable logic testing and mitigation is where Xilinx has significant history and expertise. This paper mainly focuses however on the testing of the processor subsystem, as this is a new area for Xilinx.

The Xilinx ZC702 demonstration kit printed circuit board [1] is the subject of these proton beam tests to measure the user application failure rates of the processor subsystem, and the programmable logic for the Xilinx Z-7020 device [2].

In previous tests of the TSMC 28 nm HPL process [3], on test chips and the 7 Series FPGA product line [4], the neutron testing conducted at LANSCE [5] and the proton tests conducted at UC Davis Crocker facility [6] yielded the same cross-sections for both the configuration memory cells (CRAM), and block memory (BRAM) within +/- 30%. As testing at LANSCE from visit to visit has yielded a +/-16% (95% confidence interval) variance over the years for the same device in the beam at each visit which is used to calibrate the neutron flux, we decided that the use of protons was adequate to determine the failure rates for the new Zynq family of SoC devices. The advantages of using protons are that the flux can be very high, or very low, depending on the beam current. This allows data to be gathered on very small cross-section targeted functions of the device. Where gathering such data could otherwise take weeks of beam time at LANSCE, data may be collected in days, or even hours at Crocker.

II. Vulnerability Factor and Data Corruption

Intel defined the Architectural Vulnerability Factor (AVF) in [7] as the ratio of observable functional failures to the soft error events (bit flips or transients). Not every soft error upset (SEU) or soft error transient (SET) results in an observable error due to masking in time or space. IBM has also published on the subject of the AVF [8]. Another term commonly used is silent data corruption (SDC), which relates to an operation returning a wrong result without any indication (for all intents and purposes the result is valid, yet is actually not correct). SDC is especially difficult to mitigate in a safety critical system, as to do so requires that every result be checked by both repeating the operation and checking that the result matches, or by including some kind of algorithmic check [9].

Both AVFs and the SDC rate reported here were measured for 64 MeV protons, a reasonable surrogate for broad spectrum atmospheric neutrons.

[1] Now with Swift Engineering and Radiation Services, LLC, San Jose, CA.

III. Experiment Setup

The two ARM Cortex™-A9 processors and their subsystems executed an instruction mix of the software which is used by Xilinx to verify the design both prior to tape-out and post tape-out. The results are all known, such that if any instruction results in the wrong result, it is caught, and logged. The instruction mix is meant to imitate applications running under an operating system, such as a version of GNU/Linux.

The test program with SDC visibility has an instruction mix profile consisting of:

Instruction group	**%**
Load / Store	42
FPU	22
Branch& Conditional	10
NEON	9
ALU	8
Other	8

All exceptions and interrupts were monitored so that the parity errors generated by an upset in caches or on chip memory are caught, as well as attempts to access out of bounds memory, or attempted execution of an invalid instruction.

The board under test was connected to a control laptop computer by an extended Ethernet cable to allow monitoring, logging, and control. Watchdog timers were coded into the software in the device-under-test; this allowed the laptop to record 'hang' events where the program stopped without issuing an error or exception. The programmable fabric side of the device contained AXI interfaces to the DMA, peripherals, block memories, and other logic to aid in the exercise of the processor system. Failures in the programmable logic are not discussed here, and are detailed in [10] along with the techniques of mitigation.

Exercised elements of the processor subsystem
APU – Core 0 and Core 1
A9-MPCore
Instruction Cache
Data Cache
NEON/FPU
Snoop Control Unit (SCU)
L2 Cache
On-chip fast memory (OCM)
OCM Interconnect
MMU
GIC IOP
CAN with DMA
Ethernet with DMA
I2C
SD/SDIO
UART
GPIO
IO MUX & MIO PS
TrustZone
DDR Memory Controller
DMA Controller
Central Interconnect
Clock Generators & PLLs
Device Configuration
FPGA XADC
FPGA interconnect, LUT, DFF, and BRAM

IV. Test Results

Over 25 hours of testing in the proton beam over 3 days resulted in more than 500 experiment runs which resulted in more than 5000 documented upset signatures were logged. This is equivalent to 175,000 years of terrestrial radiation exposure at the common reference point for the JESD89A standard [11]: New York City (sea level) equating to 12.9 neutrons per square centimeter per hour. Xilinx 28nm testing shows rough proton mono-energetic beam equivalence with the broad spectrum neutron LANSCE beam and others have noted a similar equivalency [12]. See Figure 1 for Xilinx's historical proton, neutron, and Rosetta soft error rates, on vertical axis, as a function of technology node, on horizontal axis.

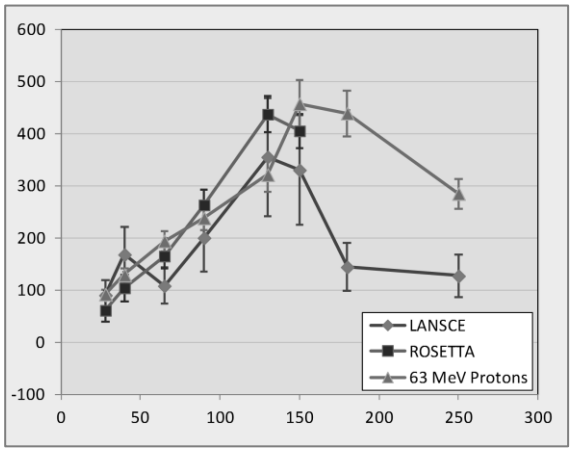

Figure 1: Proton vs. Neutron Measurements

In all testing, system 'hangs' or watchdog timers reaching their maximum wait time without a result was 1.6% of the events or 5 FIT (5 failures per billion hours). Silent data corruption was estimated at less than 15 FIT as a result of analysis of all test runs and the test duty cycle where mismatch detection is working. Together these FIT values are about 20% of the total processor sub-system FIT.

This result of total FIT is half of what the estimates were for the processor system from ARM and TSMC. The discrepancy is related to the process used (HPL vs. estimates were for HP), as well as the ultra-low alpha content of the materials that Xilinx utilizes for the fabrication, packaging, and assembly [13].

The contribution from each individual element of the processor system was noted; however, the values for most of the elements are too small (too few errors) to be accurate. For example, errors from all four of the High Performance (HP) AXI ports are 4 FIT in total. Such a result has a wide uncertainty because perhaps only ten or so such errors were present in the log.

The individual memories: L1, L2 and OCM raw FIT rates were measured separately from the instruction mix tests. The contents of the memories were written, the beam was turned on briefly, and the resulting errors counted in each memory. The resulting measured FIT rates for the processor-coupled memories were consistent with expectations based on the foundry SEU estimate. However, the raw FIT/Mb is not observed in an actual program, as typically only 5% (or less) of the memory is in use, that is, is critical to a given application. Further, "real" operation causes caches to be routinely flushed and over-written reducing the chances that upsets in those memories result in a processor system error. Thus, for example, the instruction-mix testing observed AVFs for the L1, L2, and the OCM of 5%, 5%, and 2%, respectively.

The results of the complete processor subsystem testing can be framed against the soft error rate for the full device as conveyed in Figure 2. These values are given in relative terms where the full device upset rate equals 1.00 for the combination of processor and programmable fabric.

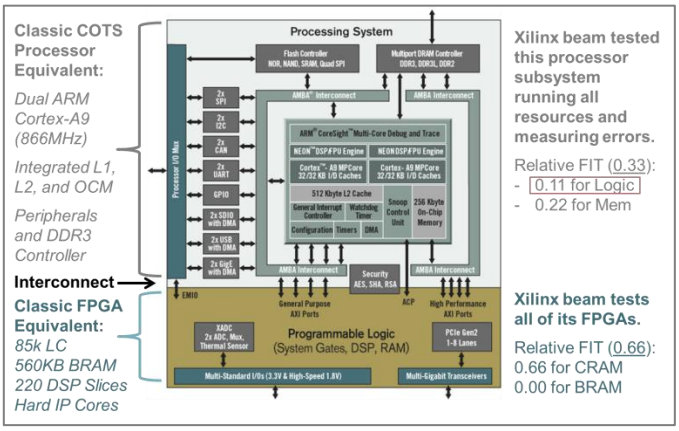

FIGURE 2: SOC AND ATTRIBUTABLE RELATIVE SEU FIT RATE

Given the nature of testing results having dependency on the embedded software, which is proprietary testing software in this case, we can only discuss absolute FIT values through NDA discussions. The results however indicate that the processor logic contributes about 11% of the FIT rate for the full device, with the processor memories contributing about 22%, and the programmable fabric contributing 67% when using Xilinx Essential Bits technology with the Single Event Upset Monitor IP [15] (SEM IP). In this discussion the system hangs and the SDC are included within the processor logic's 11% contribution. Aside from hangs and SDC, other errors do contribute to the 11%.

V. DISCUSSION

To make a safety critical application incorporating a processor system and programmable logic, one must make performance tradeoffs and architectural choices, including choosing how to mitigate failure modes.

Consider a straightforward baseline that ignores soft error detection and mitigation and three possible alternate scenarios:

Scenario	Name	FIT (arbitrary units)
1	baseline	100
2	reboot-on-error	6
3	restricted-dual	43
4	restricted-single	28

1. *baseline* with essentially no mitigation, explicitly defined as using all available resources with no re-boot allowed
2. *reboot on error detection*, again using all resources plus the processor system enables parity and other possible SEU-induced exceptions and responds with a re-boot. Alternatively, the last two scenarios apply if the occasional re-boot outage is not acceptable
3. *resource-restricted dual processor* mode with the largest soft-error targets off (L2 cache disabled and OCM not used)
4. *resource-restricted single processor* mode: similarly to scenario (3), but without 2nd CPU core.

The tradeoff for the gains in robustness of scenarios (3) and (4) relative to the baseline case is a significant performance hit for not using the caches and not using the fast on chip memory. To meet a targeted FIT rate, such decisions might be acceptable. In any safety critical system where failure is not an option (that is, the FIT rate must be less than some low absolute value, such as 10 FIT), duplication or even triplication of systems is the only solution. Essentially, the redundancy forced by the probability of a hard failure inherently eliminates soft errors as a side benefit when done correctly. In such redundant systems, common single points of failure determine the functional system failure rate.

Using the resources of the programmable logic sub-system would allow the performance penalty of scenarios (3) and (4) to be significantly reduced. In particular, ECC protected block memories (BRAM) which may be utilized as a replacement for the OCM with no performance lost because the BRAM and AXI controller can be operated at the same bandwidth by utilizing twice the data width at half the clock rate. Of course, this modification of scenarios (3) and (4) still gives up the performance advantages of the processor caches.

The hard failure rate for Xilinx 28nm HPL process devices at TSMC is published in the quarterly reliability report [14], and is 11 FIT (February 14, 2014). Arguably, once the soft failure rate is as low as the hard failure rate, it makes little sense to spend additional effort and resources to improve the soft failure rate, as hardware failure rate becomes the

dominant factor.

By utilizing the programmable logic of the Zynq device, one is also able to add hardware watchdog timers, as well as other means of mitigation such as the SEM IP which not only reports upsets, but also corrects them in the programmable logic. By mitigation of upsets in the programmable logic, and adding capabilities to the processor system, the overall processor FIT rate can be taken to the same order as the device hardware failure rate. At that point, there is no more reliable single device architecture; any further significant decrease in functional failure rate requires device redundancy and the elimination or, at least, minimization of single points of failure in the total system.

VI. CONCLUSIONS

Measurements of upset susceptibility of the Zynq processor sub-system to a 64 MeV proton beam testing were presented; the tested device was fabricated in the TSMC 28 nanometer HPL process. Previous testing on the logic sub-system suggests that LANSCE neutron beam tests and long-term atmospheric testing results would yield the same results albeit with larger statistical error bars. Those results indicate that the most impactful problems seen, hangs at 5 FIT and SDCs below 15 FIT, are very near the device hard failure rate of 11 FIT.

Thus, the system soft error functional failure rate can be brought down to the same level as the failure rate for the hardware itself. As such, a single device may be used in many critical system applications where the system functional failure rates are acceptable. In applications where the requirements are more stringent, the use of more than one device may be required.

REFERENCES

[1] Xilinx Zynq-7000 All Programmable SoC ZC702 Evaluation Kit, see http://www.xilinx.com/products/boards-and-kits/EK-Z7-ZC702-G.htm

[2] Z7020 Device Datasheet, see http://www.xilinx.com/support/documentation/data_sheets/ds190-Zynq-7000-Overview.pdf

[3] "TSMC 28 Nanometer Process Technology." Online at: http://www.tsmc.com/download/brochures/2010_28_Nanometer_Process_Technology.pdf

[4] *"Xilinx series 7 FPGA device family"* http://www.xilinx.com/innovation/7-series-fpgas.htm

[5] *LANSCE*, see http://lansce.lanl.gov/

[6] *Crocker Nuclear Laboratory*, see http://crocker.ucdavis.edu/cyclotron/usage-radiation-effects/

[7] Architectural Vulnerability Factor (or, does a soft error matter?), Shubu Mukherjee, Intel, IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems, 2008.

[8] Xiaodong Liy, Sarita V. Advey, Pradip Bosez, Jude A. Rivers "Online Estimation of Architectural Vulnerability Factor for Soft Errors," 2008. Online at rsim.cs.illinois.edu/Pubs/08ISCA.pdf

[9] S. A. Seshia, W. Li, and S. Mitra, "Verification-guided soft error resilience", in Proc. Design, automation and test in Europe (DATE '07). EDA Consortium, San Jose, CA, USA, pp. 1442 --1447, 2007. Online at: www.eecs.berkeley.edu/~sseshia/pubdir/vgser-date07.pdf

[10] Xilinx Soft Error Mitigation http://www.xilinx.com/products/quality/single-event-upsets.htm

[11] JESD89A: Measurement And Reporting Of Alpha Particle And Terrestrial Cosmic Ray-Induced Soft Errors in. Semiconductor Devices. Online at www.jedec.org/sites/default/files/docs/jesd89a.pdf

[12] Seifert et al., "Soft Error Susceptibilities of 22nm Tri-Gate Devices," *IEEE Trans.Nucl.Sci*., v59, n6, pp. 2666 - 2673, 2012

[13] Gerry Malone, Ultra-Low Alpha Packaging Materials, Xilinx, IEEE Santa Clara Valley CPMT Society Chapter Workshop, 2011 http://ewh.ieee.org/soc/cpmt/presentations/cpmt1110w-7.pdf

[14] Xilinx Quarterly Reliability Report, February 2014. Online at http://www.xilinx.com/support/documentation/user_guides/ug116.pdf

[15] Soft Error Mitigation IP Core. Online at http://www.xilinx.com/products/intellectual-property/SEM.htm