

# Upset Characterization and Test Methodology of the PowerPC405 Hard-Core Processor Embedded in Xilinx Field Programmable Gate Arrays

Gregory R. Allen, *Member, IEEE*, Gary M. Swift, *Member, IEEE*, and Greg Miller

**Abstract**— Pseudo-static upset results for memory elements in the PPC405 core embedded in a 1.5V, 130 nm Virtex-II Pro FPGA are compared to the PPC405 core embedded in a 1.2V, 90nm Virtex-4 FX FPGA. The results show consistency with earlier PowerPC processor measurements and illuminate scaling trends. While details vary, the upsettable elements consistently yield very low thresholds (below LET=2 MeV/mg/cm<sup>2</sup> for heavy ions and 10 MeV for protons), but also small per bit limiting cross sections (below 10<sup>-7</sup> cm<sup>2</sup> for heavy ions and 10<sup>-14</sup> cm<sup>2</sup> for protons) and, therefore, moderate upset rates in space radiation environments.

**Index Terms**— field programmable gate array, heavy ion upsets, single event upset, microprocessor faults.

## I. INTRODUCTION

FIELD programmable gate arrays (FPGAs) have gained great interest in recent years as an enabling technology providing designers with once unseen capability. The performance, capacity, and application of FPGAs have steadily increased over time to the point that the devices are not only capable of replacing the glue logic and management duties of ASICs, but can replace the functionality of entire processor boards. The trend of advanced reconfigurable blocks has evolved from simple programmable SRAM, to integrated hardware multipliers, to embedded hard core processors. Xilinx's Virtex-II Pro and Virtex-4 FPGAs are two such families of devices that provide designers with a

Some of the work described herein was performed by the Jet Propulsion Laboratory, California Institute of Technology under contract with the National Aeronautics and Space Administration (NASA) with funding from the NASA Electronics Parts Program (NEPP) and NASA ESR&T.

G. R. Allen is with the Jet Propulsion Laboratory / California Institute of Technology, Pasadena, CA 91109 USA (telephone: 818-353-7558; fax: 818-393-4559; e-mail: [Gregory\\_r.allen@jpl.nasa.gov](mailto:Gregory_r.allen@jpl.nasa.gov))

G. M. Swift was with the Jet Propulsion Laboratory / California Institute of Technology, Pasadena, CA 91190 USA. His is now with Xilinx Inc., San Jose, CA 95124 USA. (telephone: 408-879-2751, e-mail: [gswift@xilinx.com](mailto:gswift@xilinx.com)).

G. Miller is with Xilinx Inc., San Jose, CA 95124 (telephone: 720-652-3493, e-mail: [greg.miller@xilinx.com](mailto:greg.miller@xilinx.com)).

choice of one or two embedded PowerPC405 processors. These hard-core processors provide designers with a very powerful computing capability surrounded by programmable, reconfigurable logic gates to provide a virtually limitless array of interfaces to the rest of the spacecraft.

Commercial scaling trends continue to drive down feature size and reduce core voltage. Previous work [1] has shown that such sizing trends have reduced the static, saturated cross-sections of the processor registers and cache. This paper presents results from two generations of PowerPC processors.

Much work has been done into defining testing the inherently complicated task of testing processors [2]-[4]. In the case of the Xilinx devices, processor testing is further complicated as the processors are embedded in upsettable, programmable logic. This programmable logic is both the device's greatest strength and Achilles' heel if not properly mitigated. Proper mitigation techniques such as Triple Modular Redundancy (TMR) and partial reconfiguration must be correctly implemented in order to take advantage of the processing power [5]-[6]. If properly mitigated, the surrounding FPGA fabric provides a means to extract a great deal of processor information quickly and efficiently. After proper mitigation, a system error can be probabilistically reduced to the domain of the processors, which cannot be TMR'ed.

Most microprocessor testing falls into either the domain of static or dynamic. A test falls into one of these categories based upon the level of processor activity. Static testing is done with minimal or no processor activity. At an abstract level, static testing of a processor is done by writing to available registers and cache, irradiating while inactive (unclocked), and reading back the data after irradiation. Conversely, dynamic testing is done while clocking the processor, preferably at full speed, with known inputs while recording expected outputs during irradiation, usually using processor benchmarking software to evaluate overall performance and device characteristics. Although specific upsets cannot be segregated and counted during dynamic testing, output behaviors seen during irradiation are counted and categorized (e.g. "hangs"). Neither method of static or generic dynamic testing will produce accurate space error rate predictions. Static testing will produce a worst-case, conservative bound, which is defined in this paper. Dynamic

testing provides insight into the level of conservatism, but will be completely application specific.

## II. DEVICE TESTED

### A. Device Description

All Virtex-II Pro and Virtex-4 testing was performed using the Mil/Aero-grade devices, specifically the XQR2VP40-FF1152 and XQR4VFX60-FF1148, respectively. Both devices are packaged in a BGA package which uses a flip-chip geometry, in which the device circuitry is upside-down, facing the packages' ball contacts. When irradiating with protons, this does not cause great difficulty. Due to the limited range of the majority of heavy ion beams used at accelerators from Texas A&M and Lawrence Berkeley National Laboratory, we are required to remove the package lid and thin the substrate down to less than 100  $\mu\text{m}$ .

In addition to a myriad of user configurable features (described below), both devices contain embedded IBM PowerPC 405's, a 32-bit, integer-only microcontroller cousin to the PPC750 with which it shares its instruction set.

#### 1) Virtex-II Pro

The 2VP40 device from the Xilinx Virtex-II Pro family includes two 300 MHz-capable IBM PPC405 processors. The device is fabricated in a 130 nm CMOS process with commercial devices on bulk CMOS or a thick epitaxial layer and the Mil/Aero devices on thin-epitaxial CMOS. The device has an operating core voltage of 1.5V. In addition to the pair of PPC cores and almost 5,000 configurable logic blocks, this FPGA has the following features available for inclusion in a design: 3.5Mb of user RAM, 192 hardware multipliers, 804 I/O's with 12 so-called RocketIO Transceivers capable of speeds up to 3 Gb/s, and 8 digital clock manager blocks.

#### 2) Virtex-4 FX

The V4FX60 device from the Xilinx Virtex-4 family is fabricated in a 90 nm CMOS process, with an internal core voltage of 1.2V. The FX sub-family of Virtex-4 includes up to two IBM PowerPC405 processors specified to run at speeds up to 450 MHz. Commercial devices are fabricated on a thick epitaxial layer or on bulk CMOS with the Mil/Aero devices on a thin epitaxial CMOS layer. The user resources of this device are comparable to that of the 2VP40, including: approximately 4700 configuration logic blocks, 128 DSP slices, 4.2Mb of user RAM, 576 user I/O's, 12 digital clock managers.

### B. Storage Elements

An inventory of the upsetable storage elements and proposed mitigation techniques are shown in Table 1. The vast majority are outside the processor cores in a) the configuration cells, b) design-level memory, or c) design-level flip-flops. Mitigation in the form of a combination of triplication of design functions and storage plus active configuration scrubbing can potentially make an application very robust in spite of configuration SRAM upsets. Thus, an upset in one of the bits inside the processor is more likely to cause a system error or malfunction. Consider, for example, an upset in the program counter; it is clear that, although

relatively few, upsets in processor bits are much more important.

TABLE I  
INVENTORY OF UPSETABLE BITS IN THE XILINX VIRTEX II PRO  
AND VIRTEX-4 FX FPGAS

Name	Description	Mitigation
<b>User Resources Outside of the Processors</b>		
CLB	Configuration SRAM	Scrubbing
BRAM	Design-level SRAM	Triplication and Scrubbing
FF	Design-level Flip-Flops	Triplication
Control	Control Registers	SEFI Detection and Mitigation
<b>Processor Storage Elements</b>		
Cache	Data and Instruction Cache	Cache Parity Detection and Error Detection
SPR	Purpose Registers	Processor Comparison
GPR	Purpose Registers	Processor Comparison
<b>Processor to Configuration Interface</b>		
Gasket	PPC-to-CLB interface	Scrubbing

## III. Test Methodology

Testing for microprocessor upsets and calculating space error rates for arbitrary flight software is not straightforward; thus, there have been many different test methodologies utilized or proposed [2]. At present, microprocessor testing mostly falls into two categories, dubbed static and dynamic. Static testing treats the processor like a memory device (albeit accessed with some difficulty); that is, a pattern is loaded and, after irradiation, inspected for upsets. In contrast, during dynamic testing, a benchmark program with fixed inputs and known outputs is executed repeatedly and the results of each iteration are categorized with regard to correctness and other behaviors, for example, not completing. Neither set of results translates tractably and correctly into in-flight error rate predictions, but static results yield a conservative upper bound while dynamic results give an idea of the degree of that conservatism, at least, for an artificial software instance or instances.

The static results presented here are really pseudo-static meaning clocked at full speed, but with processor activity minimized. The processor is in a very tight infinite loop more than 99.9% of the irradiation time. A tiny fraction of the time is spent storing snapshots of the registers in a strip chart in external memory; additional details of this test method, a.k.a. "do little," can be found in [4]. The reason for pseudo-static testing (as opposed to strictly static) is to provide visibility into the evolution of the registers (e.g. a processor event may occur due to an upset in the internal state machine that writes to one or several registers invalidating the upset data). Moreover, a device event such as a Single Event Functional Interrupt (see [7]-[8] for full definitions and rates) may cause a

loss of communication to the FPGA, or potentially force a system reconfiguration. Events such as this emphasize the need to periodically strip chart the registers in order to gain confidence in the register data.

The focus during static testing is usually on a particular portion of the storage elements within the processor, as opposed to the system as a whole. Although the processor is a hard core, it is surrounded by soft/configurable Intellectual Property (IP) that is necessary to ensure correct functionality of the processor. Some of these components include: a reset block, digital clock manager (DCM), processor local bus (PLB), general purpose I/O, and processor JTAG. Certain IP, such as the reset block and DCM can be removed from the device under test (DUT) and placed into a service FPGA that exercises the DUT. IP that has timing related constraints such as the PLB must remain in the DUT and be triplicated and scrubbed. However it is accomplished, a general rule is that any soft IP should either be removed or mitigated. Isolating and mitigating the system as much as possible reduces system errors and increases the efficiency and reliability of the desired data.

#### IV. Static Test Results

Heavy ion results for Virtex-II Pro are shown in Fig. 1 and 2 while proton results are shown in Fig. 3 and 4, for the general purpose registers (GPRs) and the data cache respectively. Note that the statistical error bars for the cache data are much smaller than for the registers, almost as small as the plotting symbols; this is the result of over an order of magnitude more upsets in the cache data set. Also it is interesting that, for the cache, the susceptibility of cells storing ones is indistinguishable from that of cells storing zeros, but register bits storing ones are almost an order of magnitude more susceptible than they are when storing zeros. Qualitatively, these observations are in agreement with previous work on the IBM PPC750FX [1].

Several models for converting heavy ion data to predicted proton data were considered and the best results compared to the actual proton results similar to the work in [9]. As seen in Fig. 5, the PROFIT model [10] and Edmonds model [11] both underestimate the actual proton results for the data cache. While the PROFIT model has two adjustable parameters helping the fit, the Edmonds model has none and very rarely under predicts – these data are exceptional. In fact, a similar comparison for the GPR results, shown in Fig 5, yields closer agreement and the model is somewhat conservative when compared to the actual data.

Heavy ion results for Virtex-4 FX general purpose registers are shown in Fig. 7. Note that the saturated cross sections and LET thresholds remained nearly identical to the Virtex-II Pro GPR results. It is also noteworthy that the asymmetry seen in the susceptibility of storing ones and zeros in the GPRs was flipped, i.e. zeros are more susceptible than ones.

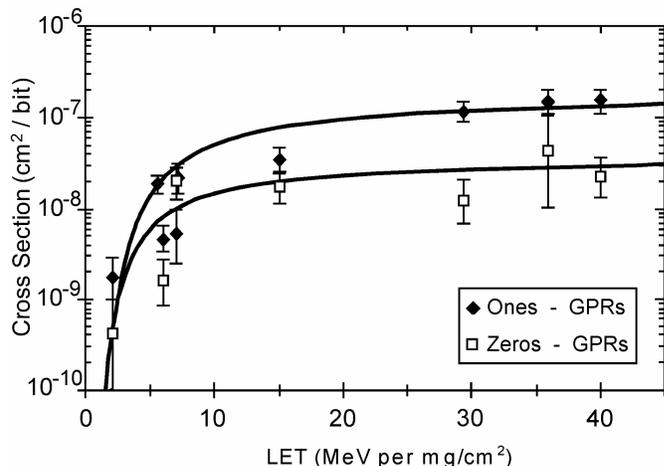


Fig. 1. Heavy ion results for the general purpose registers (GPRs). Solid lines are Edmonds' curves with the following fitting parameters: for the upper (ones) curve,  $L(1/e) = 12.8$  MeV/mg/cm<sup>2</sup> and  $\sigma_{inf} = 1.8 \cdot 10^{-7}$  cm<sup>2</sup> and, for the lower (zeros) curve,  $L(1/e) = 9$  MeV/mg/cm<sup>2</sup> and  $\sigma_{inf} = 3.6 \cdot 10^{-8}$  cm<sup>2</sup>. Counting statistic error bars are shown (approximately two sigma).

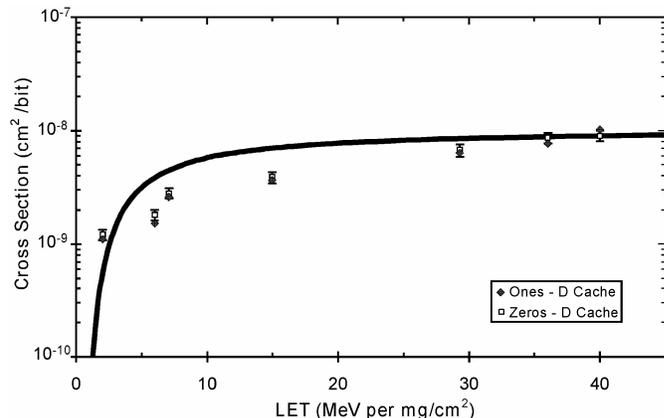


Fig. 2. Heavy ion upset results for the data cache; solid line is an Edmonds fit of the data with  $L(1/e) = 6$  MeV/mg/cm<sup>2</sup> and  $\sigma_{inf} = 1.06 \cdot 10^{-8}$  cm<sup>2</sup>. The bit susceptibility appears identical regardless of whether ones or zeros are stored. Note that the main contribution to the error bars is uncertainty in the beam fluence measurement.

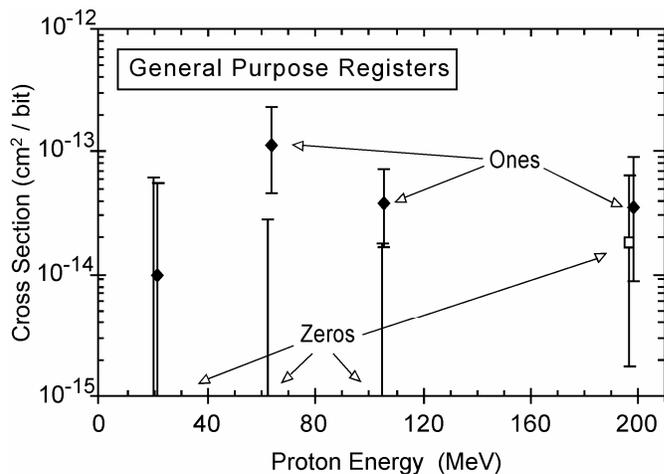


Fig. 3. Proton results for the general purpose registers (GPRs). Points are offset slightly in energy for readability. Counting statistic error bars are shown (approximately two sigma); note that for observations without any GPR upsets (i.e., the three lower energies when zeros were the contents), only the tops of the error bars can be shown.

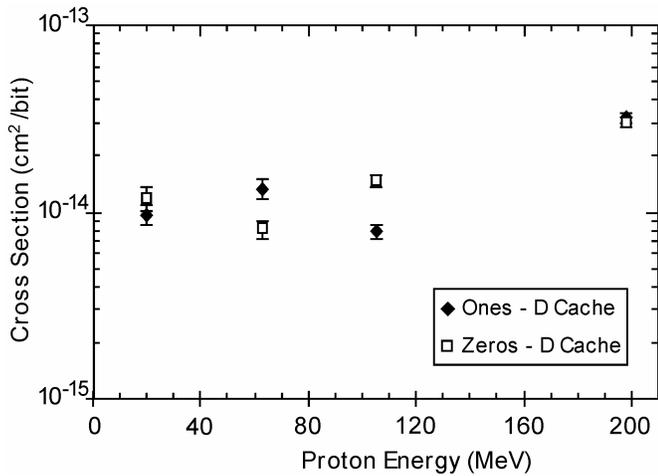


Fig. 4. Proton results for the data cache. Counting statistic error bars are shown (approximately two sigma).

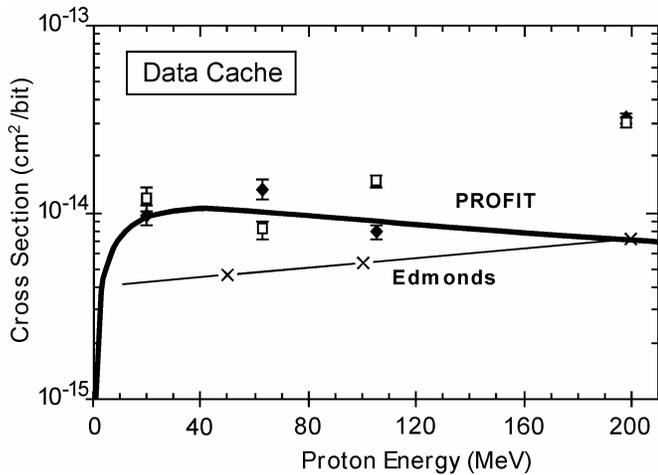


Fig. 5. Model predictions of proton results derived from the heavy ion data for the cache bit susceptibility compared to the actual proton data. The PROFIT model is from Ref. [7]; the Edmonds model is from Ref. [8]

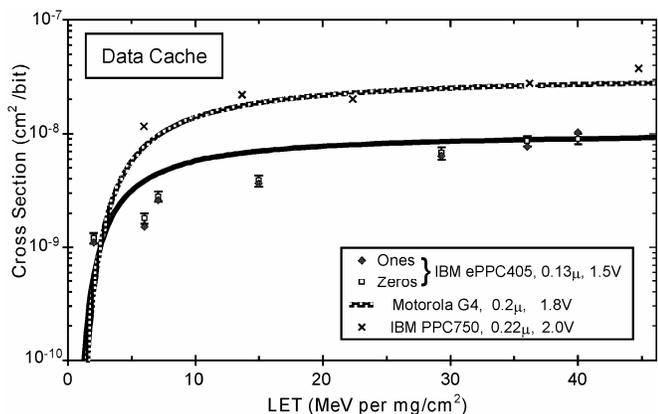


Fig. 6. Comparison of the upset susceptibility of the data cache of older PowerPC technology nodes with the present results (from Fig. 3); older nodes from Ref. [9].

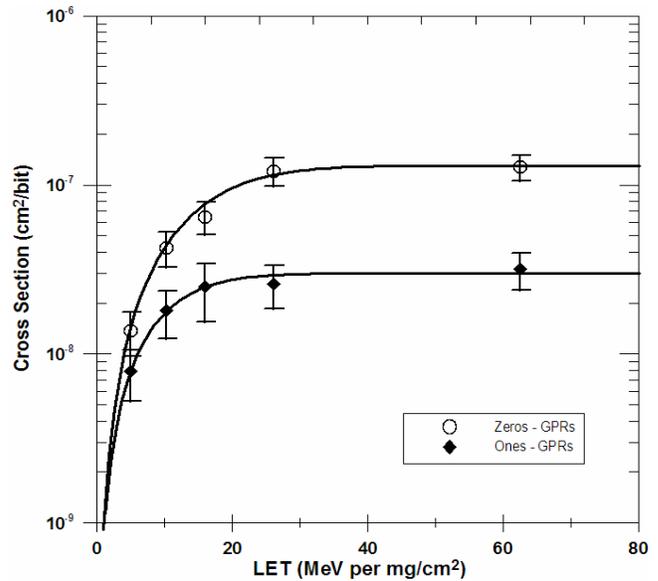


Fig. 7. Heavy ion results for the PPC's general purpose registers (GPRs) in the Virtex-4 FX. Solid lines are Weibull curves with the following fitting parameters: for the upper (zeros) curve,  $L_{th} = 1.75$  MeV per  $mg/cm^2$ ,  $\sigma_{sat} = 1.3 \cdot 10^{-7}$   $cm^2/bit$ , and  $W = 17$ ; for the lower (ones) curve,  $L_{th} = 1.5$  MeV per  $mg/cm^2$ ,  $\sigma_{sat} = 3.0 \cdot 10^{-8}$   $cm^2$ , and  $W = 11$ . Counting statistic error bars are shown (approximately two sigma).

## V. CONCLUSION

The presented pseudo-static results on PPC405 hard-core processor(s) embedded in the Virtex-II Pro and Virtex-4 FX families of Xilinx FPGAs show consistency with models and earlier data on other PPC technology nodes. An upper bound on the error rates in geosynchronous orbit can be calculated assuming all the register and cache bits are utilized with a duty cycle of 100% and yields approximately two register upset per PPC405-year and two cache upset per month. Because reloading and rebooting should take at most a few seconds, this rate is low enough that, at least for non-critical flight applications, the PPC405 core's viability is clear - upsets are unlikely to be a significant operational intrusion. In combination with higher level mitigation schemes like running two instances in lockstep [14]-[16], even most critical applications may be robust enough as well.

## ACKNOWLEDGMENT

The authors gratefully acknowledge the contributions of beam time and/or support provided by Rocky Koga of The Aerospace Corp. (at LBNL), and Carl Carmichael of Xilinx (at TRIUMF and Texas A&M).

## REFERENCES

- [1] F. Irom, F. Farmanesh, et al., "Single-event upset in evolving commercial silicon-on-insulator microprocessor technologies," IEEE Trans. Nucl. Sci., vol. NS-50, no. 6, pp. 2107-2112, Dec. 2003.
- [2] R. Koga, W.A. Kolasinski, M.T. Marra, and W.A. Hanna, "Techniques of microprocessor testing and SEU-rate prediction," IEEE Trans. Nucl. Sci., vol. NS-32, no. 6, pp. 4219-4224, Dec. 1985.
- [3] F. Bezerra et al., "Commercial processor single event tests," RADECS Conf. Data Workshop Record, 1997, pp. 41-46.
- [4] G. Swift, F. Farmanesh, et al., "Single Event Upset in the PowerPC750 Processor," IEEE Trans. Nucl. Sci., vol. NS-47, no. 6, pp. 1822-1827, Dec. 2001.

- [5] C. Carmichael, "Correcting Single-Event Upsets Through Virtex Partial Configuration". Xilinx Application Note XAPP216 (v1.0). June 1, 2000.
- [6] C. Carmichael, "Triple Module Redundancy Design Techniques for Virtex FPGAs", Xilinx Application Note XAPP197, Nov. 2001.
- [7] J. George et al., "Initial Single-Event Effects Testing and Mitigation in the Xilinx Virtex II-Pro FPGA," Paper 211, *MAPLD 2005*.
- [8] J. George et. al., "Single Event Upsets in Xilinx Virtex-4 FPGA Devices," NSREC Conf. Data Workshop Record, 2006, pp. 109-114.
- [9] R. Koga, J. George, et al., "Comparison of Xilinx Virtex-II SEE sensitivities to protons and heavy ions," IEEE Trans. Nucl. Sci., vol. NS-51, no. 5, pp. 2825-2833, Oct. 2005.
- [10] P. Calvel, C. Barillot, P. Lamothe, R. Ecoffet, S. Duzellier, and D. Falguere, "An empirical model for predicting proton induced upset," IEEE Trans. Nucl. Sci., vol. NS-43, no. 6, pp. 2827-2832, 1996.
- [11] Larry D. Edmonds, "Proton SEU cross sections derived from heavy-ion test data," IEEE Trans. Nucl. Sci., vol 47, no. 5, 2000.
- [12] F. Irom, F. Farmanesh, A. Johnston, G. Swift, and D. Millward, "Single-Event Upset in Commercial Silicon-on-Insulator PowerPC Microprocessors," IEEE Trans. Nucl. Sci., vol. NS-49, no. 6, pp. 3148-3155, Dec. 2002.
- [13] Using Proton Irradiation," IEEE REDW Record, pp 51-56, 2005.
- [14] M. Wang and G. Bolotin, "SEU Mitigation Techniques for Xilinx Virtex-II Pro FPGA," Paper D110, MAPLD 2004. Available: [http://klabs.org/mapld04/presentations/session\\_d/1\\_d110\\_wang\\_s.ppt](http://klabs.org/mapld04/presentations/session_d/1_d110_wang_s.ppt)
- [15] H.H. Ng, "PPC405 Lockstep System on ML310," Xilinx Application Note 564, Octpber 2004. Available: <http://www.xilinx.com/bvdocs/appnotes/xapp564.pdf>
- [16] Yutao He, "*smCore* – Reusable SEU Mitigation Technology for FPGA-based Mitigation: Design and Implementation Documentation", JPL Document, Rev 1.1, September 2005, unpublished.