

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2012,
17–21 SEPTEMBER 2012,
OXFORD, U.K.

FPGAs operating in a radiation environment: lessons learned from FPGAs in space

M.J. Wirthlin

*NSF Center for High Performance Reconfigurable Computing,
Department of Electrical and Computer Engineering, Brigham Young University,
Provo, Utah, U.S.A.*

E-mail: wirthlin@byu.edu

ABSTRACT: Field Programmable Gate Arrays (FPGAs) are increasingly being used as a key component of digital systems because of their in-field reprogrammability, low non-recurring engineering costs (NRE), and relatively short design cycle. Recently, there has been great interest in using FPGAs within spacecraft. FPGAs, like all semiconductor devices, are susceptible to the effects of radiation. There is an active research community investigating the effects of radiation on FPGAs and developing methods to mitigate against these effects. There has been significant progress over the last decade in the understanding and developing FPGA technology that is resistant to the effects of radiation. The success of FPGAs within spacecraft suggests that FPGAs may be used in particle physics experiments where radiation levels are considerably higher than the conventional terrestrial earth environment. This paper will summarize the effects of radiation on FPGAs, methods to mitigate against these effects, provide a case study of a successful FPGA system operating in space, and discuss the issues that will affect the use of FPGAs within particle physics experiments.

KEYWORDS: Radiation-hard electronics; VLSI circuits; Digital electronic circuits

Contents

1	Introduction	1
2	Radiation effects and FPGAs	2
3	Single event effect mitigation approaches	3
4	Cibola Flight Experiment (CFE)	4
5	FPGAs in particle physics experiments	6
5.1	FPGAs and ATLAS Liquid Argon Calorimeter	6
6	Conclusions	7

1 Introduction

Field Programmable Gate Arrays (FPGAs) are an attractive alternative to application specific integrated circuits (ASICs) because of their in-field reprogrammability, low non-recurring engineering costs (NRE), and relatively short design cycle. They provide high logic density, access to the latest I/O standards, and can be designed with a variety of low-cost tools.

FPGAs are increasingly used in non-traditional applications such as harsh environments and in safety critical systems. There has been great interest in using FPGAs within spacecraft to perform computationally demanding tasks such as remote sensing [1]. The use of reconfigurable FPGAs within a spacecraft allows the use of application-specific hardware in place of programmable processors. The ability to customize the datapath within an FPGA to an application-specific computation allows the FPGA to perform many operations faster and more efficiently than the use of traditional programmable processors.

In addition to improved computational efficiency, the use of SRAM-based FPGAs within a spacecraft allows the programmable hardware to perform any user-specified operation. Unlike application-specific integrated circuits (ASICs), FPGAs can be configured *after* the spacecraft has been launched. This flexibility allows the same FPGA resources to be used for multiple instruments, missions, or changing spacecraft objectives. Errors in an FPGA design can be resolved by fixing the incorrect design and reconfiguring the FPGA with an updated configuration bitstream. Further, custom circuit designs can be created to avoid FPGA resources that have failed during the course of the spacecraft mission.

While the use of FPGAs within a spacecraft several advantages over conventional computing methods, SRAM-based FPGAs are sensitive to the radiation found in most satellite orbits. FPGAs are sensitive to both heavy ion and proton induced single event upsets (SEUs) [2]. Single-event upsets in the FPGA affect the user design flip-flops, the FPGA configuration bitstream, and any hidden FPGA registers, latches, or internal state. There is an active research community investigating the effects of radiation on FPGAs and developing methods to mitigate against these effects.

There has been significant progress over the last decade in the understanding and development of FPGA technology that is resistant to and tolerant of the effects of radiation. The success of these efforts has facilitated the use of FPGAs in many existing spacecraft systems.

While there is still more to learn about the behavior of FPGAs in space, the success of FPGA-based spacecraft suggests that FPGAs can be used in other environments with high energy radiation such as particle physics experiments. The same benefits offered to spacecraft are available to particle physics experiments if the effects of radiation-induced single-events are properly addressed. Although the radiation environment and operating constraints of particle physics experiments are different from those of space, many of the same issues apply. This paper will summarize the effects of radiation on FPGAs and describe common methods for mitigating these effects. A case study of successful FPGA-based systems operating in space will be given. Finally, this paper will discuss an ongoing effort to investigate the feasibility of using FPGAs within the Liquid Argon Calorimeter (LAr) of the ATLAS experiment.

2 Radiation effects and FPGAs

Electronic circuits operating outside the earth's atmosphere are exposed to a radiation environment much different than the radiation found on Earth. High-levels of radiation may damage or upset the operation of a conventional semiconductor device. With increased interest in exploiting programmable logic in space related applications, several researchers have investigated the suitability of commercially available FPGAs in radiation environments [3].

Radiation has both long-term and single particle effects on electronic components. Long-term effects include accumulated total ionizing dose (TID). Single particle effects include destructive events such as single-event latchup (SEL), single-event gate rupture, and single-event burnout. Non-destructive single particle events include single event upsets (SEU) that change the state of a digital memory element and single event transients (SET). Although single-event effects do not cause any permanent damage to an FPGA, they cause a variety of problems to the run-time operation of the circuit implemented in the FPGA.

To address the need for radiation tolerant FPGAs, Xilinx has introduced a number of radiation tolerant FPGAs including the space-grade Virtex-4QV line of high-reliability FPGAs [4]. This radiation tolerant FPGA is manufactured on a thin-epitaxial 90 nm CMOS process and based on the commercially available Virtex4 FPGA architecture. The Virtex-4QV FPGA has been tested extensively for radiation tolerance and has been shown to tolerate a total dose in range of 300 krad(Si) which is more than acceptable for many space applications. In addition, this device is immune to latchup up to an LET of $100 \text{ MeV}=\text{cm}^2/\text{mg}$.

While the Virtex-4QV FPGA is immune to latch-up and has an acceptable total-dose tolerance, it is sensitive to single-event upsets (SEUs). Like all FPGAs, the Virtex-4QV contains a dense array of memory cells including user flip-flops, internal user memory, and configuration memory.

User flip-flops. An important architectural component of all FPGAs are user programmable flip-flops. User designs exploit these flip-flops to implement common sequential logic circuits such as state machines, counters, and registers. User flip-flops in most digital technologies are susceptible to radiation-induced single-event upsets.

User memory. Modern FPGAs provide blocks of internal memory larger than the typical look-up table. This block memory is used for traditional random access memory functions such as data storage, buffering, FIFO, etc. The Virtex4 family includes a set of internal dual-ported BlockRAM memories that provide 18 Kb of randomly accessible memory. Dense static memory such as the BlockRAM is especially susceptible to radiation-induced SEUs. Well-known error-correction coding techniques are often used within a user design to detect and correct such upsets.

Configuration memory. A large amount of memory cells are required to define the operation of the user-designed FPGA circuit. These memory cells define the operation of the configurable logic blocks, routing resources, input/output blocks, and other programmable FPGA resources. The use of static memory cells for configuration storage allows the device to be reprogrammed as often as necessary by reloading a new configuration memory. Like other static memory cells, configuration memory is susceptible to single-event upsets. Upsets within the configuration memory are especially troublesome as they may *change* the operation of the circuit. Several techniques have been proposed for detecting and mitigating such upsets.

Because of the large amount of configuration memory within the device, the configuration memory is especially susceptible to single-event upsets (SEUs). While upsets in the user flip-flops or memory may alter the state and output of the circuit, upsets within the configuration memory actually *change* the user design. Upsets within the configuration memory may alter the function of the configurable logic blocks, upset the routing network, or modify the operation of the input/output blocks. The XQR4VSX55 member of the Virtex4 family contains 22,702,848 configuration bits. The expected upset rate of a single configuration memory cell in a GEO orbit for this FPGA family in a quite solar maximum is $1.88\text{E-}7$ upsets/day [5] resulting in an overall upset rate for the device of 4.28 upsets/day.

3 Single event effect mitigation approaches

Many techniques have been proposed and tested for mitigating SEUs in FPGAs. Most techniques use hardware redundancy to reduce the probability of failure [6]. By replicating the desired circuitry and comparing the results, faults in the configuration can be detected and reported. The most common form of redundancy is to apply triple-modular redundancy (TMR). As shown in figure 1, TMR involves the triplication of all circuit resources and the addition of majority voters at the appropriate circuit outputs. In principle, TMR allows a circuit to operate with any single fault — the presence of a single fault will not cause any operational problem as there are two working copies of the circuit that operate correctly and selected by the majority voter. In practice, however, not all SEUs within the configuration memory can be protected by TMR [7]. A variety of FPGA-specific TMR implementation methods have been introduced to address these issues. Configuration memory fault-injection can also be used to verify the effectiveness of these techniques.

Another important technique relies device reconfiguration to continually “scrub” the configuration bitstream [8]. By repeatedly configuring the device, SEUs occurring within the configuration bitstream are repaired with the correct value. Configuration scrubbing prevents the build-up of configuration errors within the configuration memory and significantly reduces the possibility that multiple upsets will occur within the configuration memory and “break” TMR.




Figure 1. (a) A typical circuit with Flip-Flops and logic and (b) the same system operating with triple modular redundancy.

Although expensive, the most successful SEU mitigation approach for FPGAs operating within space is the combination of TMR and frequent configuration scrubbing [9]. Extensive radiation testing, reliability models, and on-orbit experience suggest that FPGAs can operate successfully in space with proper mitigation.

4 Cibola Flight Experiment (CFE)

The Cibola Flight Experiment Satellite (CFESat), funded by the Department of Energy and developed by Los Alamos National Laboratory, tests the suitability of FPGAs as an on-board reconfigurable processors in a spacecraft [10]. The CFE instrument uses the reconfigurable logic to perform high-throughput RF sensor processing. The real-time processing demands of this system are immense and cannot be performed using multi-processing with traditional radiation hardened processor architectures. This platform also incorporates a number of techniques for detecting SEUs and mitigating the effects of SEUs within the FPGAs. The architecture of the processing payload of CFE is shown in figure 2. The CFE payload includes an R6000 microprocessor, spacecraft communications interface, a digitally controlled radio tuner, a two channel, 12-bit, 100 MHz analog to digital converter, three reconfigurable computing processors using Xilinx Virtex FPGAs, and non-volatile memory to store program and FPGA configuration data.

The processing payload was built around three reconfigurable computer (RCC) modules used to perform processing duties for a variety of experiments. Each RCC module uses three Xilinx Virtex XQVR1000 FPGAs as the data processors. The nine FPGAs provide over 9 million system gates and over 1 Megabyte of block RAM memory. While the use of Virtex FPGAs in this system may seem old when compared to FPGAs available today, the Virtex 1000 FPGA family was the most complex and dense FPGA available when the CFE system was first conceived. Since all satellite systems go through an extensive design, qualification, and testing procedure, the components used on orbiting satellites typically lag far behind the components available commercially. Furthermore, the Xilinx Virtex FPGA was the first SRAM-based FPGA to go through extensive reliability qualification for radiation environments.




Figure 2. CFE Payload Block Diagram.

The CFE team investigated and developed techniques at the system level and application level for providing reliable operation of the SRAM-based FPGAs. First, the RCC boards used the QPro Virtex radiation hardened FPGAs [11]. These FPGAs use an epitaxial process that provides immunity against single event latch-up (SEL, an acute destructive failure) to an LET of 125 MeV/mg/cm^2 . The .25 micron process provides for approximately 100 K Rad(Si) total ionizing dose (a chronic and destructive radiation effect). These FPGAs are not, however, immune to SEUs within the user flip-flops, memories, and configuration data.

To detect and repair SEUs within the configuration memory, the CFE system employs a form of configuration scrubbing [8]. Configuration scrubbing is accomplished at the system level with the use of the radiation tolerant fused-based Actel FPGA. This device detects configuration SEUs by continuously reading the bitstream on each FPGA device through configuration readback. A cyclic-redundancy check (CRC) is calculated “on-the-fly” for each frame of a configuration bitstream. This calculated CRC is compared against the codebook CRCs that are precomputed on the ground. When an upset is detected by a CRC mismatch, a microprocessor interrupt is generated causing a reconfiguration of the upset frame from the onboard flash memory. The use of configuration scrubbing prevents the accumulation of configuration upsets in order to significantly reduce the probability of having concurrent multiple configuration upsets.

In addition to configuration scrubbing, a variety of application-specific mitigation techniques have been developed for CFE. Specific techniques that have been applied include half-latch removal [12] and triple modular redundancy (TMR) [6, 13]. A tool for automatically applying TMR to a user design was also created [14]. This tool triplicates circuit resources and inserts majority voters to isolate any single upset caused by a configuration SEU. This tool provides the ability to “partially” mitigate user circuits when TMR proves too costly in hardware resources.

CFE was launched into a circular low-earth orbit (560 km) on March 8, 2007 on a Lockheed Atlas-5 Medium rocket (STP1). Since its launch, CFE has received configuration data from the ground many times, both refining and increasing the portfolio of experiments within the reconfigurable payload. Several CFE signal processing payload experiments have been uploaded to the

reconfigurable platform and executed on the FPGAs. These signal processing experiments interface directly to the on-board ADC to process sampled data from the satellite antennae. Examples of this class of circuits that have been run on the satellite include several software defined radios (SDR), demodulators, decoders, and high-throughput FFT engines that exceeded a sustained computation rate of 10 Gops per second. CFESat has effectively demonstrated the benefits of deploying high-performance reconfigurable computing in space craft and the feasibility of operating the FPGAs in the presence of single-event upsets. Single event upsets have been detected on the CFESat and compared with pre-launch SEU estimations [10].

5 FPGAs in particle physics experiments

The success of FPGAs within spacecraft suggest that they may be used in other applications that operate in a radiation environment. In particular, there is growing interest in using FPGAs within regions of particle physics experiments where there is a high level of radiation. FPGAs are ideally suited for data collection, communication, and low-level data processing that occurs near the detectors of these experiments. If the effects of radiation on FPGAs can be properly addressed within these experiments, the benefits of reconfigurability, upgradability, and programmable processing can be successfully exploited by these systems.

It is anticipated that many of the same techniques used to mitigate the single-event effects of FPGAs within spacecraft may also be used to successfully deploy FPGAs within these particle experiments (i.e., configuration scrubbing, TMR, etc.). However, there are a number of differences between an FPGA operating within a spacecraft and an FPGA within a particle physics experiment. These differences will influence how FPGAs are integrated into the system, the mitigation techniques used, and the design methods employed to generate a reliable FPGA circuit. These differences will be described in more detail below.

The first difference is that most particle experiments anticipate using *many* FPGAs. FPGAs are anticipated to be placed close to the many sensors in the system. Unlike a typical spacecraft, thousands of FPGAs may be needed within an experiment. With so many FPGAs in the system, FPGA component cost becomes an important issue. These systems do not have the luxury of using the expensive, military grade FPGAs developed for space systems. The use of commercial parts will require more testing, more aggressive SEU mitigation techniques, and a more complex FPGA integration environment. In addition, the use of a large number of FPGAs will increase the overall failure rate and will require careful reliability modeling to insure proper system availability.

The second difference is the radiation environment. The flux of ionizing radiation in space varies significantly based on the orbit, orbit location, solar cycle, etc. Although the flux is fairly constant when averaged over time, the immediate flux may vary by several orders of magnitude. System designers must decide what radiation levels they are willing to tolerate. The flux of high energy particles in particle detectors, however, is relatively constant.

5.1 FPGAs and ATLAS Liquid Argon Calorimeter

An effort has been initiated with Brookhaven National Laboratory (BNL) to investigate the feasibility of using FPGAs within the Liquid Argon Calorimeter of the ATLAS experiment. Specifically, this work is actively evaluating the suitability of using the Xilinx Kintex7 FPGA within the

ATLAS LAr environment. Initial neutron testing on this device suggests that there are no problems with single-even latchup and that the single-event upset cross section is acceptable.¹ The goals of this study are as follows:

1. Estimate the Kintex7 FPGA upset rate within the LAr environment,
2. Develop and test mitigation techniques appropriate for this environment,
3. Propose a system architecture, including SEU management approached, for the LAr based on these test results, and
4. Model the reliability and availability of this proposed system architecture.

Several radiation tests are planned to collect the data necessary for estimating the FPGA upset rate within ATLAS. A neutron test was completed at the Los Alamos National Laboratory LANSCE facility in October of 2012. A test is currently planned at the CERN H4IRAD facility in November of 2012. A test with high-energy protons will also be performed. Results from these tests will be published in future conferences and workshops. Preliminary results from these tests suggest that while no catastrophic events were observed, active SEU mitigation methods such as TMR and configuration scrubbing will be needed. SEU mitigation methods and system configuration scrubbing techniques will be investigated during early 2013. A complete system architecture based on these results and models will be proposed by the end of 2013. Previous results and experience suggest that with proper use of SEU mitigation methods and active SEU management techniques, FPGAs can be successfully deployed within the ATLAS LAr and other particle physics experiments.

6 Conclusions

The successful deployment of commercial FPGAs within spacecraft suggests that FPGAs may be used in environments with high levels of radiation. Through extensive device testing, the development of novel mitigation techniques, and active SEU management methods, researchers have learned how to exploit the benefits of FPGAs while tolerating relatively frequent upsets within the configuration memory. The success of FPGAs within space suggests that they may also be used in high radiation areas of particle physics experiments. Additional efforts are being carried out to study the effects of these unique radiation environments on FPGAs and identify any unique failure mechanisms that may occur within these environments. It is anticipated that this work will demonstrate that FPGAs may be used within particle physics experiments when proper SEU mitigation and management methods are employed.

Acknowledgments

This work was supported by the I/UCRC Program of the National Science Foundation under Grant No. 0801876, the BYU Department of Electrical and Computer Engineering, and the Deployable Adaptive Processing Systems project (DAPS) at the Los Alamos National Laboratory, United States Department of Energy.

¹The neutron cross section at the Los Alamos National Laboratory for this device was measured as 6.99E-15 for configuration memory and 6.32E-15 for BRAM memory.

References

- [1] M. Caffrey, *A space-based reconfigurable radio*, in Proceedings of the *International Conference on Engineering of Reconfigurable Systems and Algorithms (ERSA)*, T.P. Plaks and P.M. Athanas eds., CSREA Press, (2002), pg. 49–53.
- [2] E. Fuller, M. Caffrey, P. Blain, C. Carmichael, N. Khalsa and A. Salazar, *Radiation test results of the Virtex FPGA and ZBT SRAM for space based reconfigurable computing*, in *MAPLD Proceedings*, September 1999.
- [3] R. Katz, *Radiation effects on current field programmable technologies*, *IEEE Trans. Nucl. Sci.* **44** (1997) 1945.
- [4] Xilinx Corporation, *Space-Grade Virtex-4QV Family Overview*, 2010, DS653 (v2.0).
- [5] G. Swift G. Allen and C. Carmichael, *Virtex-4VQ static SEU characterization summary*, Technical report, Jet Propulsion Laboratory and Xilinx, Inc., (2008), JPL Publication 08–16 4/08.
- [6] C. Carmichael, *Triple module redundancy design techniques for Virtex FPGAs*, Technical report, Xilinx Corporation, November 1 2001, XAPP197 (v1.0).
- [7] L. Sterpone and M. Violante, *Analysis of the robustness of the TMR architecture in SRAM-based FPGAs*, *IEEE Trans. Nucl. Sci.* **52** (2005) 1545.
- [8] C. Carmichael, M. Caffrey, A. Salazar, *Correcting single-event upsets through Virtex partial configuration*, Technical report, Xilinx Corporation, June 1 2000, XAPP216 (v1.0).
- [9] N. Rollins, M. Wirthlin, M. Caffrey and P. Graham, *Evaluating TMR techniques in the presence of single event upsets*, in Proceedings of the *6th Annual International Conference on Military and Aerospace Programmable Logic Devices (MAPLD)*, Washington D.C., U.S.A., September 2003, NASA Office of Logic Design, AIAA. pg. P63.
- [10] M. Caffrey et al., *On-orbit flight results from the reconfigurable Cibola Flight Experiment satellite (CFESat)*, in *IEEE Workshop on Field Programmable Gate Arrays for Custom Computing Machines (FCCM)*, 2009, pg. 3–10.
- [11] QPro Virtex 2.5V radiation hardened FPGAs, Technical report, Xilinx Corporation, November 5 2001, DS028, v1.2.
- [12] P. Graham, M. Caffrey, D.E. Johnson, N. Rollins and M. Wirthlin. *SEU mitigation for half-latches in Xilinx virtex FPGAs*, *IEEE Trans. Nucl. Sci.* **50** (2003) 2139.
- [13] F. Lima, C. Carmichael, J. Fabula, R. Padovani and R. Reis, *A fault injection analysis of Virtex FPGA TMR design methodology*, in Proceedings of the *6th European Conference on Radiation and its Effects on Components and Systems (RADECS 2001)*, 2001.
- [14] B. Pratt, M. Caffrey, P. Graham, E. Johnson, K. Morgan and M. Wirthlin, *Improving FPGA design robustness with partial TMR*, in Proceedings of the *MAPLD Conference*, September 2005.