LA-UR- 03 -0859

Approved for public release; distribution is unlimited.

Title: SEU MITIGATION FOR HALF-LATCHES IN XILINX VIRTEX FPGAS

Author(s):
Paul Stanley Graham, NIS-3
Michael Paul Caffrey, NIS-3
Michael Joseph Wirthlin, Brigham Young University
Darrel Eric Johnson, Brigham Young University
Nathan Hatley Rollins, Brigham Young University

Submitted to: IEEE Nuclear and Space Radiation Effects Conference, July 21-25, 2003, Monterey, CA, USA

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Form 836 (8/00)

SEU Mitigation for Half-Latches in Xilinx Virtex FPGAs*

Paul Graham and Michael Caffrey
NIS-3, Space Data Systems
Los Alamos National Laboratory
Los Alamos, NM 87545, USA

Michael Wirthlin, D. Eric Johnson, and Nathan Rollins
Department of Electrical and Computer Engineering
Brigham Young University
Provo, UT 84602, USA

Abstract

We introduce the half-latch single-event upset issue found in Xilinx Virtex FPGAs and describe methods for mitigating the effects of half-latch SEUs. One mitigation method's effectiveness is then illustrated through experimental data.

Authors

- *Corresponding and Presenting Author:* Paul Graham, NIS-3, Space Data Systems, Mailstop D440, Los Alamos National Laboratory, Los Alamos, NM 87545; phone: (505) 667-7024; fax: (505) 665-4197; e-mail: grahamp@lanl.gov
- *Contributing Authors:*
  - Michael Caffrey, NIS-3, Space Data Systems, Mailstop D440, Los Alamos National Laboratory, Los Alamos, NM 87545; phone: (505) 667-2422; fax: (505) 665-4197; e-mail: mpc@lanl.gov
  - Michael Wirthlin, Department of Electrical and Computer Engineering, 459 Clyde Building, Brigham Young University, Provo, UT 84602; phone: (801) 422-7601; e-mail: wirthlin@ee.byu.edu
  - D. Eric Johnson, Department of Electrical and Computer Engineering, 459 Clyde Building, Brigham Young University, Provo, UT 84602; phone: (801) 422-7206; e-mail: dej23@email.byu.edu
  - Nathan Rollins, Department of Electrical and Computer Engineering, 459 Clyde Building, Brigham Young University, Provo, UT 84602; phone: (801) 422-7206; e-mail: nhr2@ee.byu.edu

*This work is supported by the Deployable Adaptive Processing Systems project (DAPS) at the Los Alamos National Laboratory, United States Department of Energy

1 Introduction

The performance, in-system reprogrammability, flexibility, and reduced costs of SRAM-based field-programmable gate arrays (FPGAs) make them very interesting for high-speed, on-orbit data processing, but, because the current generation of radiation-tolerant SRAM-based FPGAs are derived directly from COTS versions of the chips, their memory structures are still susceptible to single-event upsets (SEUs). While previous papers have described the SEU characteristics and mitigation techniques for the configuration and user memory structures on the Xilinx Virtex family of FPGAs ([1] [2] [3] [4] [5] [6] [7] [8]), we will concentrate on the effects of SEUs on "half-latch" structures within the Virtex architecture, describe techniques for mitigating these effects, and provide new experimental data which illustrate the effectiveness of one of these mitigation techniques under proton radiation.

2 Half-Latches and SEUs

Xilinx Virtex FPGAs use a little-known resource, called the "half-latch", to generate many of the constant "0" and "1" logic values used internally by Virtex FPGA designs. By using half-latches to produce constant logic values, Xilinx can avoid using more expensive logic resources such as look-up tables (LUTs) to generate constant logic values. At an architectural level, half-latches drive many of the internal inputs to I/O, logic, RAM, clocking, and other resources when there are no direct sources for the inputs, i.e., when the inputs are left unconnected. Hence, half-latches are very efficient and ubiquitous sources of constant "0" and "1" values that are available throughout these devices and, consequently, they are used frequently by the Xilinx implementation tools in most designs. Our experience has been that large Virtex XCV1000 designs commonly use hundreds or thousands of half-latches.

Unfortunately, half-latch resources are susceptible to radiation-induced SEUs when exposed to proton or other forms of radiation. When upset, the output values of these circuits will remain inverted until the device is fully reprogrammed or another upset occurs. Further, unlike upsets in the FPGA device's programming data (frequently called a *configuration bitstream*), these inversions are not directly observable through device mechanisms for reading back the device's programming data. Thus, an SEU affecting a half-latch is not as easily detectable or correctable as SEUs in the configuration bitstream.

An example of the half-latch issue is shown in Figure 1. In this case, the unused clock enable input is driven by a half-latch. If the half-latch is upset, it will disable the flip-flop, modifying the intended function of the circuit. This modification cannot be observed through reading back and checking the configuration bitstream.

3 Mitigation Techniques

The solution to the hidden half-latch inversion problem is to use explicit constant sources which can be configured directly through FPGA programming data instead of using the FPGA's implicit half-latch resources. As a result, any SEUs that affect constant generation can be both observed and repaired through existing SEU mitigation techniques for configuration bitstreams ([2] [3] [5] [7]). Figure 2 illustrates one option for an observable, correctable constant valued source—an FPGA input pin driven externally by a logic "0" or "1". LUTs and flip-flops can also be effectively used as constant sources. While this solution to the half-latch problem does require additional FPGA resources and uses constant sources that are still vulnerable to SEUs, the results are FPGA designs which can be suc-

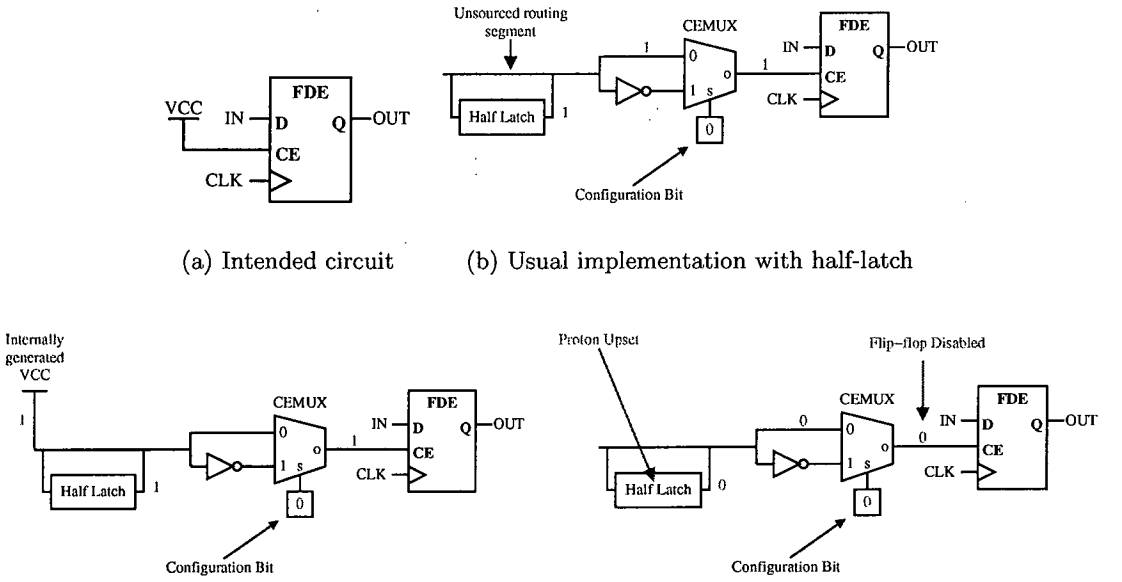

(a) Intended circuit (b) Usual implementation with half-latch

(c) Initialization during full configuration (d) Circuit with half-latch SEU

Figure 1: Half-latch Illustration Showing the Intended Circuit, Half-latch Initialization, and the Result of a Half-latch Upset

cessfully repaired while they are still executing as opposed to having to stop the system completely to perform a full FPGA reconfiguration.

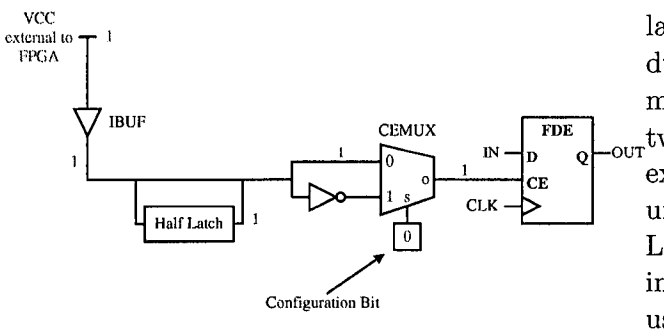

Figure 2: An Observable, Correctable Alternative to Half-latches

Even for small designs, adding explicit constant sources throughout the design can be tedious and error prone so automating the half-latch replacement process is essential. Los Alamos has developed a tool named *RadDRC* which replaces half-latches in the final placed-and-routed, low-level design, ensuring that half-latches do not reappear later in the design flow due to complex design synthesis or technology mapping processes. *RadDRC* currently provides two options for replacing half-latches: using an externally sourced constant "1" value as in Figure 2 or using a "1" value generated by internal LUTs. Since resources' inputs can be selectively inverted, these constant "1" values can also be used to generate "0" values for the inputs that require it.

4 Radiation Experiment

Los Alamos National Laboratory and Brigham Young University conducted a proton radiation experiment at Crocker Nuclear Laboratory in

November 2002 to validate our SEU simulator and the *RadDRC* tool. Regarding half-latches, the experiment's goal was to measure the average fluence to failure (FtF) under proton radiation for both a normal FPGA design and a version of the same design which had undergone *RadDRC*'s half-latch replacement scheme. "Failure" in this case is when a design does not operate properly even when its configuration data is error-free—the signature observed when a half-latch critical to design operation has been upset. This FtF quantity provides us with a measure for comparing the relative circuit reliability for the unmitigated and mitigated designs.

The experiments were conducted using 63.3 MeV protons with beam fluxes of $1.0x10^7$ and $3.5x10^7 \frac{p}{cm^2 \cdot s}$. The designs executed on a modified USC/ISI SLAAC-1V FPGA board with a free running system clock of 20 MHz for all but one trial, where a 2 MHz system clock was used. For each test, one copy of the design under test (DUT) executed in an FPGA which was exposed to proton radiation. A second, "golden" copy of the design executed synchronously with the DUT in an FPGA shielded from the protons. A third FPGA performed real-time output differencing between the two designs so we could quickly determine when the outputs for the two designs no longer matched. During the experiment, our software continuously identified and repaired configuration bitstream upsets and reset both designs when the two design's outputs differed, allowing their operation to be resynchronized—this sampling and repair process occurred at a 1–2 Hz rate. When the DUT's output continued to differ from the output of the "golden" design despite having an error-free configuration bitstream and having been reset multiple times, a half-latch upset failure was identified. For the test, we only had enough beam time to thoroughly test one design: a series of large cascaded multipliers and adders mimicking the hardware used in many signal processing applications.

Initially, we expected that observing 3 or 4 consecutive DUT output errors after bitstream repairs and design resets would be enough to identify a design failure due to half-latch upsets. The assumption was that half-latch upsets should not recover over time. As it turned out, this assumption was not valid for the unmitigated design where as many as 40 consecutive output error samples could occur but the design would still recover with or without the proton beam on. In contrast, half-latch mitigated designs never experienced more than 4 consecutive samples with output errors. Because of this, we declared a repeated failure as a full half-latch failure only after a significant number of failures occurred (often 100+) from which the design would not recover.

Table 1 provides the observed fluence to failure measurements for the sample design and its mitigated version. For the unmitigated design, three distinct sets of data were taken at different points during the test and, hence, three different ranges of accumulated dosage. As the FPGA accumulated dosage, the average fluence to failure due to half-latches for the unmitigated design decreased.

As noted in the table, the half-latch mitigated design did experience one apparent half-latch failure. After the test, the half-latch mitigated design was analyzed for possible causes of this behavior. Though the analysis indicated that two operation critical half-latches had been left in the design, the half-latch failure could not be recreated using our SEU simulator, so the actual cause of the failure is still not certain.

Finally, the most significant result is the ratio of the two average fluence-to-failure results for the two designs. The mitigated design achieves over 100 times the fluence to failure over the unmitigated design, a significant improvement in reliability. Considering that the accelerator test time was limited, this result may be conservative.

Design Test	Total Failures (p/cm^2)	Total Fluence	Ave Fluence to Failure (p/cm^2)	Accum. Dosage Range (krads)
Unmitigated Design				
Set 1	5	5.80E10	1.16E10	17.6 to 25.4
Set 2	15	5.42E10	3.62E9	57.2 to 64.5
Set 3	13	1.74E10	1.34E9	82.9 to 85.1
All	33	1.30E11	3.93E9	17.6 to 85.1
Mitigated Design				
All	1	4.10E11	4.10E11	10.7 to 86.5
Ratios of Mitigated to Unmitigated Results				
All	1/33	3.16	104.41	

Table 1: Fluence to Failure for the Unmitigated and Mitigated Versions of the Sample Design

References

[1] 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.

[2] C. Carmichael, M. Caffrey, and A. Salazar, "Correcting single-event upsets through Virtex partial configuration", Tech. Rep., Xilinx Corporation, June 1, 2000, XAPP216 (v1.0).

[3] E. Fuller, M. Caffrey, A. Salazar, C. Carmichael, and J. Fabula, "Radiation testing update, SEU mitigation, and availability analysis of the Virtex FPGA for space reconfigurable computing", in *3rd Annual Conference on Military and Aerospace Programmable Logic Devices (MAPLD)*, 2000, p. P30.

[4] 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.

[5] C. Carmichael, E. Fuller, J. Fabula, and F. D. Lima, "Proton testing of SEU mitigation methods for the Virtex FPGA", in *4th Annual Conference on Military and Aerospace Programmable Logic Devices (MAPLD)*, 2001, p. P6.

[6] C. Carmichael, "Triple module redundancy design techniques for Virtex FPGAs", Tech. Rep., Xilinx Corporation, November 1, 2001, XAPP197 (v1.0).

[7] 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. June 2002, pp. 49–53, CSREA Press.

[8] P. Graham, M. Caffrey, M. Wirthlin, D. E. Johnson, and N. Rollins, "Reconfigurable computing in space: From current technology to reconfigurable systems-on-a-chip", in *Proceedings of the 2003 IEEE Aerospace Conference*, Big Sky, MT, March 2003, p. TBD.