Evaluating Xilinx SEU Controller Macro for Fault Injection

José Luís Nunes, João Carlos Cunha
Polytechnic Institute of Coimbra/CISUC
Coimbra, Portugal
{jnunes, jcunha}@isec.pt

Raul Barbosa, Mário Zenha-Rela
University of Coimbra/CISUC
Coimbra, Portugal
{rbarbosa, mzrela}@dei.uc.pt

***Abstract*—This paper presents a preliminary evaluation of the SEU Controller Macro, a VHDL component developed by Xilinx for the detection and recovery of single event upsets, as a building block of an FPGA fault-injector. We found that this SEU Controller Macro is extremely effective for injecting faults into the FPGA configuration memory, as single and double bit-flips, with precise location, virtually no intrusiveness, and coarse timing accuracy. We present some clues on how to extend its functionalities to build a fully-fledge FPGA fault injector.**

Keywords— fault injection; FPGA; SEU; embedded systems

I. INTRODUCTION

Field Programmable Gate Arrays (FPGA), first introduced as prototyping platforms or as low-cost replacements for the small demand of Integrated Circuits (IC) in high-density applications, are increasingly being used for production in industrial, automotive and space areas, due to their flexibility, reconfigurability and power efficiency. Recent advances in SRAM-based FPGAs, taking advantage of fast memory cells for storing configuration data, allow shorter access times and unlimited reconfigurations. However, these same SRAM cells are very sensitive to the effects of radiations, manifested as single event upsets (SEU), thus hindering its applicability to the just mentioned areas. In order to improve the resilience of these devices it is necessary to add some mechanisms for fault tolerance. Fault injection is a well-known technology to assess the effectiveness of such mechanisms. A few FPGA software-based fault-injectors have already been developed, resulting from academia research projects, e.g. FLIPPER [1] and FT-UNSHADES [2]. Due to the strong dependency on the underlying technology, they target very specific devices.

One of the two major FPGA manufacturers, Xilinx (www.xilinx.com), has been introducing several fault tolerance technologies, such as radiation-hardening, or readback CRC checks. More recently, this manufacturer designed the *SEU Controller Macro* (SEU-CM) [3], a VHDL component that uses error correction codes for double bit error detection and single bit error correction in the configuration memory frames of Virtex-5 FPGAs. In order to test this mechanism, Xilinx extended this macro to be able to introduce errors in the same configuration memory frames, in a non-intrusive way, through Partial Dynamic Reconfiguration.

This fast abstract presents a preliminary study of the capabilities of this SEU-CM to perform fault injection in FPGA-based embedded systems.

II. THE EFFECTS OF FAULTS IN FPGAS

FPGA upsets can be classified in three categories [4]: configuration, user logic, and architectural. These upsets are produced in the same way, but have different effects on the user design. Configuration upsets affect the configuration memory, resulting in permanent errors in the logic and routing of the implemented design. Errors such as a change in the logic function of the designed system (e.g. a logic OR becomes an XOR) or a short (e.g. two separate wires become connected) may result from faults affecting lookup tables (LUT), multiplexer (MUX) select lines or programmable interconnection points (PIP). These faults represent more than 80% [5] of the total faults affecting FPGAs. It is worth mentioning that a large percentage of cells in the configuration memory are not used in the implementation of the designed system, thus many faults do not produce changes in this design.

User logic upsets affect logic elements that store user data or system state, such as user memory or registers. As the contents of these cells change throughout normal system execution, a fault in a memory cell can be easily mistaken by a regular write operation. Finally, architectural upsets affect FPGA control elements (e.g. ICAP, JTAG, reset, etc.). The effects of these faults are likely to be catastrophic.

III. EVALUATING THE SEU CONTROLLER MACRO

The SEU-CM provides two interfaces to simulate the effects of SEUs in Virtex-5 devices. The main interface uses two signal lines to define the error type and the injection instant. Each time one of these lines is asserted, a Single-Bit Error (SBE) or a Multi-Bit Error (MBE) (depending on the asserted line) is generated in a randomly selected configuration memory location, by flipping a single bit or two adjacent bits in the same frame [6]. This fault model was based on the information gathered from an on-going project to measure the effects of radiation in SRAM-based FPGA – Rosetta Experiment [7].

The second interface, an UART serial port, was added for debugging purposes, in order to inform the user when an error was detected and, if possible and when enabled, automatically corrected. Besides this functionality, this interface also allows to trigger the fault injection, by specifying the type, as SBE or MBE, and location, as the address of the target frame and bit offset. The injection instant occurs immediately after a toggle bit command is sent through the UART and decoded by the

SEU-CM. This means that the time accuracy of fault injection is strongly dependent on the UART settings, in particular the baud rate and the SEU-CM finite state machine. This time delay, from the moment the fault injection order is triggered until the bit-flip(s) occurs, may be significantly different depending on the location of the triggering device: external to the FPGA, connected through a physical serial port, or internal to the FPGA, using dedicated logic.

The observation of the effects of the injected faults at FPGA level is achievable through the debugging facilities of the SEU-CM using the UART interface. By reading the frame contents, it is possible to obtain the internal state of the system stored in Flip-Flops, Block RAM, or Lookup Tables. However, the frame-addressing mode in the FPGA allows the existence of addresses not directly mapped to configuration memory bits, meaning that some of the randomly selected locations may not lead to real bit-flips. Furthermore, as already mentioned, some of the errors successfully introduced are located at configuration memory cells that are not being used by the implemented system design, thus are totally ineffective. The observation of the effects of the injected faults at system design level is limited by the existing information provided by Xilinx tools about the mapping between the bitstream and the higher-level view of the implemented design.

The area occupied by the SEU-CM in the FPGA configuration memory is about 206 Kbytes, which corresponds to a space overhead of 0,8% in a XUPV5 development board from Digilent, which includes a Xilinx Virtex-5 XC5VLX110T FPGA. On the other side, there is absolutely no intrusiveness of this macro with respect to the execution time of the system under test, because this family of devices is able to rewrite the frame contents without the need to previously erasing it, thus allowing the system to run without interruption.

IV. Improving FI capabilities

Albeit having the possibility of inserting bit-flips in the configuration memory of Virtex-5 FPGAs, the SEU-CM was built for early detection and fast correction of SEUs, thus lacks some capabilities to be used as a fully-fledged fault-injector. In this section we address some requirements for a fault injection tool, describe some ideas on how to use and extend the SEU-CM for FPGA fault-injection, and present some challenges.

Injecting faults in unused locations means that they will not be effective, thus the experiments are worthless. It is thus necessary to clearly identify the area of the FPGA where the system under test is located. Taking advantage of allocation resource files generated by Xilinx tools (namely the XDL output files), it is possible to identify the frames of interest for injecting faults. Anyway, even after excluding the unused regions of the FPGA, some experiments will still be ineffective since many bits inside the chosen frames are not being used.

Once the boundaries of the region of the FPGA that will be subjected to fault injection are defined, it is possible to place the hardware control logic that supports the fault injection experiments running in the same FPGA, without any interference due to the inherent parallelism of FPGA devices. It is even possible to create a reconfigurable partition that will be used exclusively by the device under test, isolating the target from the fault injection control logic and even the SEU-CM, that are located in the static region of the FPGA. Space and time intrusiveness are therefore eliminated, and even the trigger delay is significantly reduced when activated from inside the FPGA.

Increasing the observability of the system under test at low level (e.g. macro-cell) is not easily attained due to the little amount of information available from Xilinx about the bitstream internal format. At the system design level, probes can be attached to the interfaces of the reconfigurable partition to monitor the I/O and data can be logged to external storage for offline analysis.

The SEU-CM fault model only considers single or double adjacent bit-flips. Injecting multiple bit-flips in the same frame can easily be achieved by changing its VHDL code. However, simultaneously injecting multiple bits in adjacent frames, simulating an area of radiation-induced multiple bit upsets, cannot be done, as each frame is read-modified-written sequentially. Fortunately, intermittent single-bit faults are easily achieved by invoking the SEU-CM at different instants.

V. Conclusions

This fast abstract presented a brief evaluation of Xilinx SEU Controller Macro for injecting faults in the FPGA configuration memory. Some challenges and hints for improving and extending this VHDL macro for building a fully functional FPGA fault-injector have been introduced and discussed. This macro revealed to be very effective for injecting single bit-flip faults, in a non-intrusive way.

As FPGAs are gaining importance in critical systems domains, fault-injection tools targeting these devices are becoming crucial for system's verification and validation. We are currently developing a fault injector tool for the Xilinx Virtex-5 FPGA, making use of the SEU-CM capabilities.

References

[1] M. Alderighi, F. Casini, M. Citterio, S. D'Angelo, M. Mancini, S. Pastore, G. R. Sechi and G. Sorrenti, Using FLIPPER to predict irradiation results for Virtex 2 devices, In Proc. of RADECS, 2008.

[2] H. Guzman-Miranda, J. N. Tombs, M. A. Aguirre, FT-UNSHADES-uP: A platform for the analysis and optimal hardening of embedded systems in radiation environments, IEEE International Symposium on Industrial Electronics, 2008.

[3] K. Chapman, "SEU strategies for Virtex-5 devices," Xilinx XAPP864 v2.0, 2010.

[4] E. Fuller, M. Caffrey, A. Salazar, C. Carmichael, J. Fabula, Radiation characterization, and SEU mitigation, of the Virtex FPGA for space-based reconfigurable computing, In Proc. of NSREC, 2000.

[5] H. Asadi, M. B. Tahoori, B. Mullins, D. Kaeli, and K. Granlund, Soft Error Susceptibility Analysis of SRAM-Based FPGAs in High-Performance Information Systems. In IEEE Transactions on Nuclear Science, 2007, vol. 54, no. 6, pp. 2714–2726.

[6] K. Chapman, "New Generation Virtex-5 SEU Controller", Xilinx AN vA.2, 2010.

[7] A. Lesea, S. Drimer, J. Fabula, C. Carmichael, and P. Alfke, The Rosetta Experiment: atmospheric Soft Error Rate testing in different technology FPGAs, IEEE Trans. Device Mater. Rel., 5(0): 317–328, 2005.