

# Radiation-hardening of off-the-shelf FPGA systems

Keywords : Fault tolerance; Single-Event Effect; Space applications; Low-power; SoC FPGA.

Charlotte Frenkel, Jean-Didier Legat, David Bol

**Abstract – Radiation-induced errors in electronic circuits, known as Single-Event Effects (SEEs), are critical in satellites as space is a radiative environment: ensuring the reliability of electronic devices raises a tradeoff between reliability, overhead and latency. In order to leverage high-performance and low-cost Commercial Off-The-Shelf (COTS) FPGAs in space applications, this work tackles fault tolerance along three abstraction levels: circuit, organization and control. We proposed a topology based on modern Xilinx Zynq System-on-Chip FPGAs. It offers strong circuit overhead reductions compared to the conventional Triple Modular Redundancy (TMR) and was successfully validated through fault injection simulation and proton beam-testing.**

In space, circuits are particularly exposed to Single-Event Effects (SEEs), resulting from high-energy particles striking the silicon lattice, which induce soft errors through Single-Event Upsets (SEUs) in memory elements and Single-Event Transients (SETs) in combinational logic. Circuits dedicated to space applications must thus be made radiation-hardened using specific design techniques. As technology scales down to increase resource integration and reduce power, circuits become more vulnerable to upsets, up to the point that new challenges are posed by the occurrence of Multiple-Bit Upsets (MBUs) which cannot be further neglected. Compared to dedicated ASICs, the use of FPGAs in space applications reduces development costs, improves the time-to-market and allows for on-orbit reprogrammability, hence lowering the mission risk. There is an increasing interest in using Commercial Off-The-Shelf (COTS) SRAM-based FPGAs as they drastically improve performance over traditional radiation-hardened FPGAs, leading to a narrower required satellite communication bandwidth while further reducing costs. However, as the SRAM used to store the configuration bitstream is very sensitive to upsets, proper hardening techniques must be applied to protect both the user logic and the configuration memory.

As the power consumption impacts the battery size and thus the weight of the satellite, the main objective of our work is to design cost-driven fault-tolerance topologies with minimum resource and power overheads while ensuring full error handling, up to MBU hardening. In order to optimize the resulting tradeoff between reliability, overhead and latency, we proposed in [1] a new design methodology based on three abstraction levels: circuit, organization and control. An overview of the proposed hardening strategy is shown in Fig. 1, it is based on a modern Xilinx Zynq SoC FPGA which embeds an FPGA and an ARM Cortex A9 processor in the same die.

We proposed a key innovation at each abstraction level of our methodology. At the circuit level, a new ultra-low overhead Forward Temporal Redundancy (FTR) scheme was designed to detect errors in user logic at an overhead below that of Duplication With Comparison (DWC). At the organization level, this work leveraged the opportunities brought by frame- and module-based Dynamic Partial Reconfiguration (DPR) to handle configuration memory errors. At the control level, this work fully exploited the Xilinx Zynq SoC FPGA by offloading a circuit state preservation structure based on checkpointing and rollback to the embedded Cortex A9. Choosing a five-stage pipelined MIPS processor as a benchmark, our complete topology is far more efficient than a Triple Modular Redundancy (TMR) design and requires only 85% combinational and 125% sequential overheads. The detection and correction latencies are of 4.5 ms and 320  $\mu$ s, respectively.

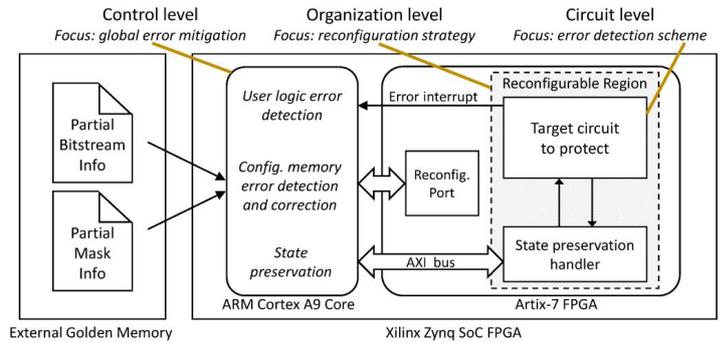


Figure 1: Global overview of the hardening strategy over three abstraction levels: circuit, organization and control.

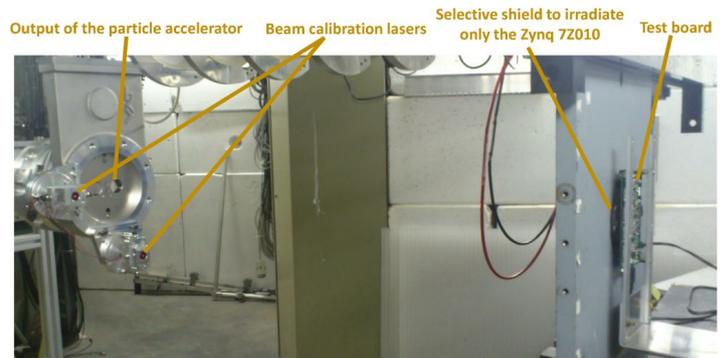


Figure 2: Proton beam-testing setup, simulating a radiative environment on a Xilinx Zynq 7Z010 SoC FPGA.

The proposed design was successfully validated in a two-fold process: fault-injection simulation was first conducted to verify the different concepts, then proton beam-testing was carried to simulate the particle strikes on the tested device (Fig. 2), which is the closest approximation to real space conditions [2]. Fault-injection predicted a 99.998% reliability, while beam testing reported only one system error over 493 logged SEEs, over which 147 were MBUs. It shows that, despite the low resource utilization of our design, full MBU hardening at a minimized latency penalty is reliably achieved.

Further research in this field targets the hardening of other critical points such as the golden bitstream memory, the clock distribution and the I/Os.

## References

- [1] C. Frenkel, J. D. Legat and D. Bol, "A Partial Reconfiguration-based scheme to mitigate Multiple-Bit Upsets for FPGAs in low-cost space applications", in Proc. of International Symposium on Reconfigurable Communication-centric Systems-on-Chip (ReCoSoC), 2015
- [2] C. Frenkel, "Soft Error Mitigation Techniques for Modern Commercial Off-The-Shelf FPGAs in Space Applications", Master Thesis, UCL, 2015.