

Design of Finite State Machines for SRAM-based FPGAs operated in radiation field

Matteo Lupi^{*a,b}, P. Giubilato^c, M. Bonora^{a,d}, K. M. Sielewicz^{e,f}

^aCERN, Geneva (CH)

^bGoethe University, Frankfurt am Main (DE)

^cUniversità e INFN, Padova (IT)

^dSalzburg University, Salzburg (AT)

^epreviously at CERN, Geneva (CH)

^fpreviously at Warsaw University of Technology, Warsaw (PL)

E-mail: matteo.lupi@cern.ch

For the CERN LHC Run 3, the ALICE experiment completely redesigned the Inner Tracking System (ITS), which now consists of seven cylindrical layers instrumented with 24 120 Monolithic Active Pixel Sensors, covering an area of 10 m². The ITS is controlled and read out by 192 custom Readout Units, which employ commercial SRAM-based FPGAs and will operate in a radiation field, requiring specific FPGA design to ensure system reliability. This contribution focuses on the techniques developed for designing radiation tolerant finite state machines, discussing the theoretical background, the actual implementation, and their validation with fault injections and proton irradiation tests.

Topical Workshop on Electronics for Particle Physics TWEPP2019
2-6 September 2019
Santiago de Compostela - Spain

*Speaker.

1. Introduction

During the current CERN Large Hadron Collider shutdown period, the ALICE experiment is upgrading its Inner Tracking System (ITS) [1]. The new ITS uses 24120 ALPIDE chips: custom large area CMOS Monolithic Active Pixel Sensors (MAPS) [2]. The readout system, composed of 192 identical Readout Units (RUs), has complete control over all sensor operations, including power management. Its reliability is, therefore, critical for the correct operation of the entire ITS. The RUs will be placed at about five meters from the interaction point, along the beam axis, and at a radial distance of about one meter. The expected Total Ionising Dose (TID) for the entire detector life cycle is about 100 Gy (Si), including a safety factor of ten, which does not raise concerns since all the system components have been verified against this dose. Conversely, the expected flux of high energy hadrons with sufficient energy (> 20 MeV) to induce Single-Event Effects (SEEs) in microelectronic devices is of the order of $1 \cdot 10^3 \text{ cm}^{-2} \text{ s}^{-1}$, posing a challenge to the utilisation of commercial SRAM-based FPGAs. Irradiation tests showed that the whole system of 192 RUs, each employing a Xilinx Kintex UltraScale XCKU060 FPGA, will experience on average an SEE affecting the FPGA configuration RAM (CRAM) every eight seconds. An external scrubbing sub-system, driven by a flash-based FPGA, ensures the long-term stability of the FPGA design. However, the SRAM-based FPGA design needs to deal with errors until the scrubbing corrects them. Previous publications [3] showed how to efficiently protect combinatorial and sequential networks without feedback. Finite State Machines (FSMs) represent a fundamental building block of an FPGA design: since they need to retain a state, it is mandatory that they can continue operating between scrub cycles.

2. Theoretical approach




Figure 1: Base FSM. The input Σ and the current state S generate the next state in the δ process. The output process ω generates the output Γ from the current state (Moore FSM).

The reliability of a state machine, as the one in Fig. 1, can be modelled using time-continuous Markov chains [4]. The model proposed, first studies a single bit FSM, protected using different mitigation techniques: no mitigation, Triple Modular Redundancy (TMR) with single and triplicated voter, and Hamming with Single Error Correction (SEC). It takes into account different failure rates for logic λ_f and voters λ_v , and it can also account for the scrubbing correction rate if the state of the FSM is not lost due to the upset.

In the base example provided in Fig. 2, a triplicated FSM protected with a single voter with non-negligible cross-section is shown. The associated Markov chain can be drawn assuming three different states: one where the FSM is working, one where one of the three TMR branches is broken, and one where the FSM is not working any longer. The latter happens either when the voter is broken or when two branches have an upset. From the Markov chain a state transition matrix, \underline{T} can be derived, in Eq. (2.1), indicating the system evolution as a function of time. Solving




Figure 2: Triuplicated FSM. Left: Block diagram of the connections. Right: Markov chain used to calculate the reliability model.

the differential-equation system, the probability of the system being in a determinate state can be found. The reliability of the system, $R(t)$ is the probability of the system being in the state 0 or 1. As it can be observed in Eq (2.2) the reliability is the product of the reliabilities of the two sub-components: this is the expected behaviour of a series of sub-systems. The first term is the reliability of the voter, the second term the reliability of a triuplicated system.

$$\underline{T}(\Delta t) = \begin{bmatrix} 1 - 3\lambda_f\Delta t - \lambda_v\Delta t & 3\lambda_f\Delta t & \lambda_v\Delta t \\ 0 & 1 - 2\lambda_f\Delta t - \lambda_v\Delta t & 2\lambda_f\Delta t + \lambda_v\Delta t \\ 0 & 0 & 1 \end{bmatrix} \quad (2.1)$$

$$R(t) = e^{-\lambda_v t} (3e^{-2\lambda_f t} - 2e^{-3\lambda_f t}) \quad (2.2)$$

The reliability of multi-bit FSMs can be calculated as the reliability of the series (if one fails, the FSM fails) of single-bit FSMs. The assumption made is that the FPGA LUT inputs are not saturated, outside this approximation, the single bit upset rate will change, in addition to the effect of the series of reliability. The evaluation of the reliability of an FSMs protected with Hamming requires a similar approach to the one described, but it demands studying the FSM knowing the number of state bits. The plot in Fig. 3 shows the reliability function of an eight-state FSM protected with different radiation-effect mitigation techniques and a logic upset rate of $\lambda_f = 1 \cdot 10^{-12} \text{ s}^{-1}$. In particular, the effect of the voter cross-section and the scrubbing are studied. In the case of the unprotected FSM, scrubbing has no effect. The ideal TMR has a negligible voter cross-section, whereas, for the real TMR the voter has the same cross-section as the logic, which represents a realistic scenario in SRAM-based FPGAs. The ideal TMR shows no difference between single and triuplicated voter implementation, whereas, the difference is significant when scrubbing is active.

3. Experimental approach

The different radiation-effect mitigation techniques were validated on a Xilinx Kintex-7 XC7 K325T FPGA by fault injection and proton irradiation. The fault injection [3], was executed in the laboratory by randomly toggling a CRAM bit during the FPGA operation and subsequently restoring its value. The number of errors corresponds to an equivalent fluence of $\Phi_{equivalent} = 8.2 \cdot 10^8 \text{ cm}^{-2}$, i.e. 227 h of operation of a single RU, or 1.18 h of the whole system of 192 RUs.




Figure 3: Reliability of different TMR implementations. The ideal TMR assumes a voter with a neglectable cross-section. The real TMR assumes a voter with a cross-section equal to the cross-section of the logic is protecting.




Figure 4: Block diagram of the irradiation setup.

The 20 MeV proton irradiation was run with a flux of $1 \cdot 10^6 \text{ cm}^{-2} \text{ s}^{-1}$, in order for the mean time between CRAM upset to match the scrubbing cycle. The test flux was 1000 time higher than the target application flux, but it was used to collect statistics in a reasonable time and to study how the design operates when CRAM upsets accumulate.

The results in Fig. 5 show a significant difference between the different radiation-effect mitigation techniques. The differences between the fault injection and the proton irradiation are mainly in the different amount of statistics gathered, i.e. in the confidence interval of the survivability plot. The non-mitigated design has, as expected, the highest cross-section. The Hamming with SEC and the TMR with a single voter have similar performance: the one with the lowest resources occupancy, i.e. TMR is preferable. When the number of accumulated errors increases, however, they start failing since the feedback mechanism is not protected. The TMR with a triple voter, however, shows a significantly better performance than the other solutions. This is due to the higher level of protection guaranteed by the implementation, which can keep working despite one out of three voters and FSMs being affected by an error. The triple voter also guarantees that the correct value is fed back to the δ combinatorial logic. The Mean Time To Failures (MTTF) for a single FSM potentially used in the ITS readout can be calculated using the target flux at the ITS readout electronics position, i.e. $1 \cdot 10^3 \text{ cm}^{-2} \text{ s}^{-1}$. The values in Table 1 are relative to the cross-sections estimated during the proton irradiation. The TMR with triple voter guarantees a much higher MTTF if compared to the other techniques. This method has a cross-section of at least one order of magnitude lower while having similar resources utilisation to the Hamming implementation. Assuming

Table 1: Mean fluence (Φ) and time to failure as estimated from the proton irradiation results for 8-state FSMs protected with different radiation-effect mitigation techniques.

Mitigation technique	M Φ TF $1 \cdot 10^{12}$ [cm $^{-2}$]	MTTF (1 RU) $1 \cdot 10^4$ [day]	MTTF (192 RUs) $1 \cdot 10^2$ [day]
No mit.	1.58 ± 0.64	1.83 ± 0.75	0.95 ± 0.39
Ham. SEC	1.67 ± 0.92	1.94 ± 1.07	1.01 ± 0.56
TMR $_{sv}$	3.60 ± 2.24	4.17 ± 2.60	2.17 ± 1.35
TMR $_{tv}^a$	26.70 ± 2.67	30.90 ± 3.09	16.10 ± 1.61

^a Lower limit with CL = 0.99

approximately 100 FSMs per RU, we can estimate that TMR with triple voter provides more than a full day of MTTF, considering the whole system. For comparison, the average data acquisition run in ALICE is approximately ten hours.




Figure 5: Survivability vs fluence plots in different conditions. Left: fault injection results. Right: proton irradiation results.

4. Conclusion

An existing methodology for studying the reliability of FSM has been extended to the use-case of SRAM-based FPGAs in radiation field. The performance of this theoretical approach are shown, highlighting that the effect of the voting logic must be taken into account. The approach has been experimentally validated with a beam test showing an improvement in MTTF of the design protected with TMR with triple voter of 15 times *w.r.t.* the non-protected one.

References

- [1] B. Abelev et. al [ALICE Collaboration]. Technical Design Report for the Upgrade of the ALICE Inner Tracking System. *Journal of Physics G: Nuclear and Particle Physics*, 41(8):087002, jul 2014.
- [2] Gianluca Aglieri Rinella. The ALPIDE pixel sensor chip for the upgrade of the ALICE Inner Tracking System. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 845:583 – 587, 2017. Proceedings of the Vienna Conference on Instrumentation 2016.
- [3] K. M. Sielewicz, G. A. Rinella, M. Bonora, P. Giubilato, M. Lupi, M. J. Rossewij, J. Schambach, and T. Vanat. Experimental Methods and Results for the Evaluation of Triple Modular Redundancy SEU Mitigation Techniques with the Xilinx Kintex-7 FPGA. In *2017 IEEE Radiation Effects Data Workshop (REDW)*, pages 1–7, July 2017.
- [4] Daniel McMutrey and Keith S. Morgan. Estimating TMR Reliability on FPGA Using Markov Models. Technical Report 149, Brigham Young University, Dec 2008.