RIS team (Robust Integrated Systems)

Themes

Robust massively parallel single-chip architectures
Power management from the OS down to silicon
Fault tolerant and self-adaptative architectures
3D NOC Robust Architectures
Design in Reliability face to aging, process variation and soft errors
Evaluation of robustness and qualification: radiation testings, fault injection
Architectures for Nanotechnologies

Expertise

Fields of expertise
Design for Reliability, Design for Test, Self-Repair, Fault-tolerance, Design for Soft-Error Mitigation: Methodologies, Tools and Architectures

Know-how
Multilevel platforms for fault simulation and robustness automatic insertion at several abstraction levels; 3D integration solutions
Test platform for radiation faults measurement; SEE error-rate prediction of circuits and systems

Research keywords

Fault tolerance, multi-core systems robustness, 3D circuits, aging, fault-injection

Contact

Raoul VELAZCO
CNRS Research Director
(+33) 4 76 57 46 89
raoul.velazco@univ-grenoble-alpes.fr

Cells: On-Chip Self-healing Massively Parallel Tera-Device Processors in Ultimate CMOS and Beyond. From Circuit-Level to System

Members: N.E. Zergainoh, M. Nicolaidis, A. Charif, T. Bonnoit, A. Coelho

Keywords: Ultimate CMOS and post-CMOS technologies, high defect densities, reliability, yield, low-power, massively parallel single-chip tera-device computers,

Cooperation: STMicroelectronics, iRoC, Atmel

Contracts: ELESIS, RESIST

1. Goal

With every new generation of semiconductor device manufacturing process, further miniaturization is realized and more complex designs enabled. Ultimate-CMOS and post-CMOS technologies promise integrating trillions devices in a single die, leading to single-chip massively parallel architectures comprising thousands interconnected processors. While this is meant to increase functionality and performance of the target electronic product, but it makes the new chips more susceptible to defect, and result in fabrication yield, life-span, and reliability challenges. It has adverse impacts on:

- Timing skews
- Process, voltage, and temperature (PVT) variations
- Circuit aging caused by hot carrier injection (HCI), negative-bias thermal instability (NBTI)
- Sensitivity to electro-magnetic interferences (EMI) causing cross-talk and ground bounce
- Sensitivity to radiation causing single-event effects (SEUs, SETs)
- Power dissipation and thermal constraints

These issues are becoming the main showstoppers in the path leading to these technologies.

The resulting high defect levels, heterogeneous behavior of identical circuit nodes, circuit degradation over time, and integrated circuits complexity, affect adversely fabrication yield and reliability.

Cells (On-**C**hip S**el**f-hea**l**ing Tera-Device Processor**s**) project comprises several techniques spanning at all levels of the system: circuit, processor/architectural, routing and task-scheduling/allocation. Innovations are introduced at all levels of this framework, including its overall architecture, its particular components, and the way the cooperation of these components is architected to optimize the outcome. Cells addresses the several issues related to ultimate CMOS and post CMOS massively parallel Tera-Device processors, such us:

- After fabrication, all processing and routing nodes may be affected by some temporary faults such as delay faults, or clock skews.
- Fabrication faults altering persistently the circuit behavior may affect one or more regular blocks (RAMs, FIFOs, buses) in a large fraction of nodes. Such faults may also frequently occur during product life.
- Fabrication faults altering persistently the behavior of irregular blocks (thus difficult to repair) may affect a significant portion of nodes. Such faults may also frequently occur during circuit life (e.g. every few days), and thus during application execution.
- New timing faults induced by circuit aging, as well as soft errors (SEUs and transients) may frequently occur during circuit life (and thus during application execution).
- Circuit degradation is continuous and requires continuous self-regulation of circuit parameters (clock-frequency, voltage levels, body bias), to maintain operational each processor node.

Reducing Rollback Cost in VLSI Circuits to Improve Fault Tolerance:

In nanometer technologies, circuits are more and more sensitive to various kinds of perturbations. Alpha particles and atmospheric neutrons induce single-event upsets, affecting memory cells, latches, and flip-flops. They also induce single-event transients, initiated in the combinational logic and captured by the latches and flip-flops associated with the outputs of this logic. In the past, the major efforts were related on memories. However, as the whole situation is getting worse, solutions that protect the entire design are mandatory. Solutions for detecting the error in logic functions already exist, but there are only few solutions allowing the correction, leading to a lot of hardware overhead in non-processor design. In Cells project, we present a novel technique that includes several hardware architectures and an algorithm for their implementations, which reduces the cost of rollback in any kinds of circuit.

Designing reliable cores in ultimate CMOS and beyond: A double sampling solution:

The double sampling paradigm is an efficient method to protect the circuits against soft-errors. But the data that are going out of the area protected by double sampling are still vulnerable. In Cells project we proposed an architectural solution that uses three latches to remove those constraints and protect the area outside the double sampling domain without adding a buffer stage.

Robust Routing Solutions for TSV-Based Three-Dimensional Networks-on-Chip:

3D integration opens up new opportunities for future many-cores by enabling fast and highly scalable 3D Network-on-Chip (NoC) topologies. However, in an aim to reduce the cost of Through-silicon via (TSV), partially vertically connected NoCs, in which only a few vertical TSV links are available, have been gaining relevance. To reliably route packets under such conditions, we introduce a lightweight, efficient and highly resilient adaptive routing algorithm targeting partially vertically connected 3D-NoCs. It requires a very low number of virtual channels (VCs) to achieve deadlock-freedom (2 VCs in the East and North directions and 1 VC in all other directions), and guarantees packet delivery as long as one healthy TSV connecting all layers is available anywhere in the network. For Cells project, an improved version of our algorithm is also developed and shown to dramatically improve performance under low TSV availability while still using less virtual channels than state-of-the-art algorithms. A comprehensive evaluation of the cost and performance of our algorithms is performed to demonstrate their merits with respects to existing solutions.

2. References

[1] A. Charif, A. Coelho, M. Ebrahimi, N. Bagherzadeh, N.E. Zergainoh, "First-Last: A Cost-Effective Adaptive Routing Solution for TSV-Based Three-Dimensional Networks-on-Chip", IEEE Transactions on Computers, Ed. IEEE, Vol. 67, No. 10, pp. 1430-1444, October 2018

[2] T. Bonnoit, N.E. Zergainoh, M. Nicolaidis, "Reducing Rollback Cost in VLSI Circuits to Improve Fault Tolerance", Transactions on Very Large Scale Integration (VLSI) Systems, Ed. IEEE, Vol. 26, No. 8, pp. 1438-1451, août 2018

[3] A. Charif, A. Coelho, N.E. Zergainoh, M. Nicolaidis, "A Framework for Scalable TSV Assignment and Selection in Three-Dimensional Networks-on-Chip", Journal of VLSI Design, Ed. Hindawi Publishing Corporation, Vol., 2017

[4] A. Charif, A. Coelho, N.E. Zergainoh, M. Nicolaidis, "A Dynamic Sufficient Condition of Deadlock-Freedom for High-Performance Fault-Tolerant Routing in Networks-on-Chips", IEEE Transactions on Emerging Topics in Computing, Ed. IEEE, Vol. PP, No. 99, DOI: 10.1109/TETC.2017.2776909, November 2017

[5] A. Coelho, A. Charif, N.E. Zergainoh, J. Fraire, R. Velazco, "A soft-error resilient route computation unit for 3D Networks-on-Chips", Design, Automation & Test in Europe (DATE'2018), Dresden, GERMANY, March 19-23, 2018

[6] T. Bonnoit, G.F. Bouesse, N.E. Zergainoh, M. Nicolaidis, "Designing reliable processor cores in ultimate CMOS and beyond: A double sampling solution", Design, Automation & Test in Europe (DATE'2018), Dresden, GERMANY, March 19-23, 2018

Soft Errors in SRAMs at Ultra-Low Bias Voltages

Keywords: Accelerated radiation tests, high-energy neutrons, SEE, SEU, MCU, MBU

Members: J.A. Clemente, F.J. Franco, G. Hubert, R. Velazco, H. Puchner

Cooperation: University Complutense of Madrid, ONERA, Cypress Semiconductor

1. Context and goals

Ionization resulting from charged particles present in the environment where integrated circuits operate may result in a wide range of consequences gathered under the acronym SEE (Single Event Effects) and classed as soft and hard errors. Soft errors are the consequences of the modification of the contents of memory cells. The evaluation of the sensitivity of a given technology is mandatory to cope with this problem.

A common strategy to save power in modern electronic devices consists in reducing the power supply value to the lowest possible value that allow retaining information. This is known as Dynamic Voltage Scaling (DVS). Although said threshold value depends on the technology miniaturization, in practical it means that the power supply can drop from 3.3 V (nominal voltage) down to, e.g., 0.6 V. However, the main drawback of this approach is that it leads to an increase of the soft-error rate of the devices in such a way that even low-energy particles can induce single events.

In order to understand this phenomenon, TIMA/RIS collaborates with GHADIR (Group of Dynamically Reconfigurable Hardware) research group in UCM (*Universidad Complutense de Madrid)* since 2013. Other organisms / companies involved in this research line are ONERA (*Office National d'Etudes et de Recherches Aérospatiales*) at Toulouse in charge of modelling, simulations and predictions, and Cypress Semiconductor (US), who has provided proprietary information about the layout of the DUTs (Devices Under Test). CMOS COTS Cypress 130 nm, 90 nm and 65 nm SRAMs, as well as Renesas' 110 nm SRAMs have been tested under neutron radiation [1, 2]. The latter is soft-error free, according to the manufacturer.

Another challenge that arises is the correct interpretation of the results of a radiation campaign. More specifically, Single Bit Upsets (SBUs) and Multiple Cell Upsets (MCUs) must be correctly extracted from the bulk of bitflips. Even if an accurate methodology is used for this purpose, there is still a problem: the accumulation of errors, which might lead to think that 2 independent (and nearby) SBUs are part of the same MCU. To cope with this problem, a methodology has also been developed to estimate how the measured multiple-event cross sections must be corrected to remove the overestimation due to false events [3].

Radiation campaigns were carried out in the GENEPI-2 (*GEnérateur de NEutrons Pulsés Intense*) facility, which is a neutron accelerator located at *Laboratoire de Physique Subatomique et de Cosmologie* (LPSC, Grenoble - France). It produces 14.2-MeV neutrons that are produced under the impact of a deuteron beam onto a tritium (T) or deuterium (D) target by fusion reactions (Figure 1). The DUT faces the centre of the beam line at a distance adjusted to the desired neutron flux. Typically the DUT is placed to reach a neutron flux of ~2.10^7 $n.s^{-1}.cm^{-2}$. Under these conditions, the DUT is exposed to a dose of ~10^{10}-10^{11} neutrons within one hour.

Figure 1: The GENEPI-2 neutron facility

GENEPI-2 was proved to be an efficient irradiation facility to study the neutron impact on integrated circuits: relevant statistics on induced SEUs can be reached within five minutes. This irradiation flux is a significant improvement with respect to the radiation values previous to the facility upgrade that was carried out recently.

2. Recent outcomes

The following 2Mx8-bit SRAMs from Cypress Semiconductor and Renesas were tested:

- CY62167DV30LL: 130 nm (Cypress Semic.)
- CY62167EV30LL: 90 nm (Cypress Semic.)
- CY62167GE30-4: 65 nm (Cypress Semic.)
- RMLV1616AGSA-5S2: 110 nm (Renesas)

Results obtained for SRAMs at ultra-low bias voltage

Experimental results clearly indicate that the soft-error cross sections increase as the power supply value decreases [1]. Figure 2 presents the raw numbers of bitflips that were found, although a finer analysis distinguishing MCUs and SBUs was also carried out in that work (not included in this report for simplicity). Events with larger multiplicities were very common for the 90-nm SRAM, but much less likely in the other ones (see the peaks between 1,5 and 2,0 V for that SRAM in Figure 2). Also, the soft-error cross sections show a sensitivity decrease as the technology shrinks (this can also be observed in Figure 2, roughly).

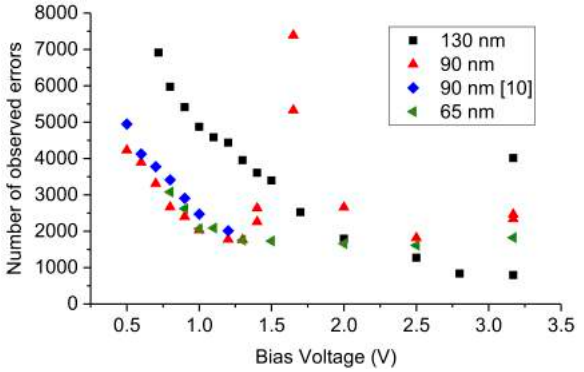

Figure 2: Errors observed at different voltages (taken from [1])

The classification of the different kinds of errors was done by using proprietary information provided by Cypress Semiconductor. Another important detail is that the 65-nm device has Error Correction Code (ECC) that had to be deactivated for a better interpretation of the experiments.

The results from the experiments were used to feed simulation tools developed by ONERA that allowed estimating the expected number of errors due to radiation in several natural environments such as sea level or high-altitude flights. Matching between predictions and experimental results was almost perfect.

Experiments also showed that micro latch-ups and SEFIs (Single Event Function Interrupts) may occur in these SRAMs at intermediate voltages (1.1-1.7 V) yielding strange of errors on the order of 25 bitflips and MBUs (see Figure 3). This phenomenon was not observed neither at nominal voltages (~3.3 V) nor at ultra-low ones (~0.5 V).

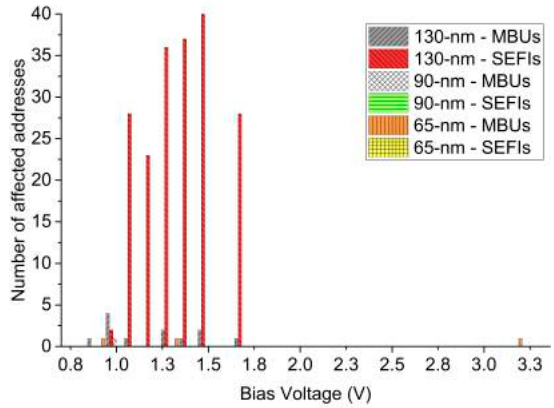

Figure 3: Errors affected by MBUs and SEFIs (taken from [1])

Concerning the RMLV1616AGSA-5S2 SRAM, an old version of this device (CMOS 150-nm) was tested in previous campaigns. Recently, a more advanced model was released by Renesas in CMOS 110-nm technology, so different samples were purchased and tested with nominal and low bias voltages (from 0.5V to 3.3V). Experimental results seem to show that, at nominal voltage, this SRAM is extremely tolerant to 14-MeV neutrons, backing up the manufacturer's claim of being soft-error free at typical environments. However, isolated SBUs, along with clusters of errors that were only visible at low bias voltages (which were named "Low Voltage Stuck Bits") were observed. There is a slight difference in the signature of these errors between the new device and the previous one. Indeed, in the 150-nm device, said errors usually occurred at ultra-low bias voltages, close to the retention limit. On the contrary, in the new model, the phenomena occur at intermediate voltages. As Renesas has not yet provided information about the internal structure of the device, these phenomena will be postulated in the near future.

SBU/MCU refinement technique

Finally, it is well known that, if a radiation-ground experiment has enough bitflips, it is possible that unrelated addresses affecting nearby cells are erroneously taken as a multiple event. We have also studied radiation experiments as a special case of the urn-and-balls problem in probability theory to estimate how the measured multiple-event cross sections must be corrected to remove the overestimation due to the false events [3]. That work proposes equations to estimate the number of false MBUs and 2-bit MCUs.

These equations were verified with Monte-Carlo simulations, where only SBUs occurred randomly in all the SRAM. The match between the predictions of the equation and said simulations was perfect (Figure 4, referring to the Manhattan Distance, a metric to group bitflips into MCUs).

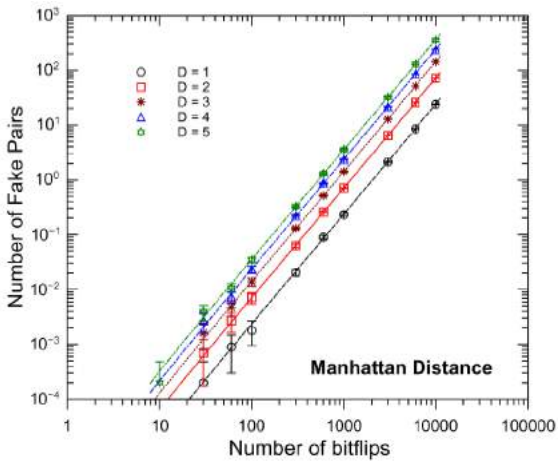

Figure 4: Monte-Carlo simulations vs. predictions of false MCUs, when using Manhattan distance metric (taken from [3])

3. References

[1] J.A. Clemente et al., "*SEU Characterization of Three Successive Generations of COTS SRAMs at Ultralow Bias Voltage to 14.2 MeV Neutrons*," IEEE Trans. on Nuclear Science, Vol. 65, Issue 8, 2018

[2] M. Rezaei, J.C. Fabero, H. Mecha, F.J. Franco, S. Rey, B. Cheymol, M. Baylac, R. Velazco, J.A. Clemente, "*Impact of Neutron Irradiation on a 110-nm COTS Soft-Error free SRAM at Low Bias Voltage*", RADECS 2018

[3] F.J. Franco, J.A. Clemente, H. Mecha, R. Velazco, "*Influence of Randomness during the Interpretation of Results from Single-Event Experiments on SRAMs*", IEEE Transactions on Device and Materials Reliability (Early Access), December 2018

Evaluation of the SEE sensitivity and methodology for error rate prediction of applications implemented in Multi-core and Many-core processors

Keywords: Reliability, Radiation ground test, Fault Injection, Many-core processor, Multi-core Processor, Single Event Effect, Single Event Upset, Robust

Members: P. Ramos, V. Vargas, R. Velazco, N.E. Zergainoh

Cooperations: KALRAY

Contracts: CAPACITES (Calcul Applications Critiques en Temps et Sûreté), projet LEOC (Logiciel et Objets Connectés)

1. Introduction

For selecting the appropriate device intended to be implemented in a system operating in harsh radiation environment, it is imperative to evaluate the SEE sensitivity of the candidates, and then to establish a trade-off between costs and reliability depending on the application and the operating environment. However, the evaluation based on dynamic radiation tests is costly in terms of time and money. For this reason, a prediction approach is required. Furthermore, the widespread use of multi/many-core processors in embedded systems requires a prediction error-rate suitable for these devices.

This work proposes an error-rate prediction approach and the evaluation of the sensitivity of applications implemented in multi and many-core processors exposed to harsh radiation environments. To validate the generality of this approach, three different Commercial-Off-The-Shelf devices were targeted aiming at representing the most relevant technological and architectural aspects of multi/many-core processors: the Freescale P2041 quad-core processor, the Adapteva Epiphany E16G301 microprocessor and the Kalray MPPA-256 many-core processor. The Single Event Effect (SEE) sensitivity evaluation and the error-rate prediction was accomplished by combining radiation experiments with 14 Mev neutrons in particle accelerators to emulate a harsh radiation environment, and fault injection in cache memories, shared memories or processor registers, to emulate the consequences of Single Event Upset (SEU) in the program execution.

2. Error-rate prediction approach

The proposed approach is based on the principles of the Code Emulating Upset (CEU) approach. Up to now, the CEU approach has been successfully applied and validated for mono-core processors. However, the complexity of the processors has significantly increased due to the manufacturing technology, device architecture, number of cores, interconnections, functionalities, etc. Therefore, it is reasonable to validate a new approach for complex devices such as multi/many-core processors. Due to the large number of functionalities and pins that complex processors implement, it is not further possible to use the ASTERICS platform for injecting fault in this kind of devices. It is thus convenient to extend the CEU approach to multi/many-core processors benefiting of the multiplicity of cores by using one of them as fault injector while the others execute the chosen application. In order to isolate the fault injector, the device has to be configured in Asymmetric multi-processing mode. For performing the fault-injection, inter-core interrupts are used. Considering the architecture of the multi/many-core processors, this work proposes the addition of derating factors to the contribution of shared and cache memories for improving the accuracy of the prediction. These factors depend on the memory used by the application and the exposure time to radiation of shared and cache memories. By adding these derating factors, the equation that defines the approach is the following:

$$\tau_{SEU} = \tau_{inj} * \sigma_{STATIC} * \mathrm{Mf} * \mathrm{Etf}$$

Where τ_{inj} is obtained from fault injection campaigns. It is defined as average number of injected faults needed to cause an error in the result of the application. σ_{STATIC} is obtained from radiation tests and provides the average number of particles needed to cause a bit-flip in the device memory cells. Mf is the memory utilization factor. It is the amount of memory used by the application with respect to the total memory of the device. Lastly, Etf is the exposure time factor, which is applied when the multi/many-core processor performs as a co-processor of a development board and in order to log the results, it needs synchronization between the co-processor and the Host.

3. Evaluation of the target devices

Freescale P2041

This device is a quad-core processor manufactured in 45nm SOI technology which implements ECC and parity in their cache memories. Obtained results from the evaluation of the P2041 multi-core demonstrate that fault injection allows identifying vulnerabilities in the application, and improving the programming strategy for reducing the impact of faults in the results. From the static test, it was confirmed that SOI process technology is more robust than traditional bulk CMOS. On the other hand, dynamic tests have demonstrated that in spite of the parity and ECC protection mechanisms, there were errors in the result of the application caused by

MBUs in the address tags and data array. Finally, results show an underestimation of the predicted error-rate, since not all sensitive zones were targeted during the static test and fault injection campaigns. Furthermore, the implementation of ECC and parity in the device's cache memories may affect the error-rate prediction.

Adapteva Epiphany E16G301

This microprocessor is manufactured in 65nm CMOS process which integrates 16 processor cores and do not implement any protection mechanism. From its evaluation, it can be seen that the proposed approach was effective for predicting the application error-rate. The fact that this device does not implement protection mechanisms has allowed a good estimation of the error-rate, confirming that protection mechanisms affect the testing and error-rate prediction. During the dynamic radiation test, input matrices were also checked to identify silent faults. It was done in order to obtain the experimental error-rate of the application which has a good correlation with the error rate obtained from fault injection.

Kalray MPPA-256

The MPPA-256 many-core processor is manufactured in 28nm TSMC CMOS technology which integrates 16 compute clusters each one with 17 processor cores, and implements ECC in its static memories and parity in its cache memories.

Its evaluation shows that both, ECC and interleaving implemented in the SMEMs of the clusters are very effective to mitigate SEU type errors, since all the detected SEUs in the SMEMs were corrected during the static test. In addition, dynamic tests have demonstrated that by enabling the cache memories it is possible to gain in application performance without a reliability penalty, since cache memories implement an effective parity protection. Regarding the radiation experimental results, the prediction of the error-rate was based only on registers' contribution since they do not implement any protection mechanism. Despite the complexity of this many-core processor, the prediction of the error-rate has a small underestimation that confirms the applicability of the approach to these devices. The possible reasons for this underestimation are: only accessible registers were targeted, communication infrastructure was not targeted, protection mechanisms may affect the error-rate prediction. Table 1 shows a summary of the neutron radiation experiments targeting the three multi/many-core processors.

DEVICE	σ_{STATIC} ***[cm²/device]***	***Measured error rate [cm²/device]***	***Predicted error rate [cm²/device]***
Freescale P2041	$8.51\text{x}10^{-9}$	$4.25\text{x}10^{-9}$	$3.27\text{x}10^{-9}$
Adapteva E16G301	$9.27x10^{-8}$	$4.63\text{x}10^{-8}$	$4.02\text{x}10^{-8}$
Kalray MPPA-256	$12.71x10^{-9}$	$5.78x10^{-9}$	$4.73\text{x}10^{-9}$

Table 1: Radiation experiments on multi/many-core processors

4. Future work

The current work has presented a first insight in the vast study of the sensitivity to radiation of multi-core and many-core processors. For continuing with this work, the following topics can be explored: validation of the proposed approach using different programming models, validation of a real space application, validation by exposing to heavy-ions, evaluation of the communication infrastructure, and application of redundancy techniques to improve the reliability of the device.

5. References

[1] P. Ramos, V. Vargas, M. Baylac, F. Villa, S. Rey, J.A. Clemente, N.E. Zergainoh, J.F. Méhaut, R. Velazco, "Evaluating the SEE Sensitivity of a 45 nm SOI Multi-Core Processor Due to 14 MeV Neutrons", IEEE Trans. Nucl. Sci., vol. 63, no 4, pp. 2193-2200, August 2016

[2] V. Vargas, P. Ramos, V. Ray, C. Jalier, R. Stevens, B. Dupont de Dinechin, M. Baylac, F. Villa, S. Rey, N.E. Zergainoh, J.F. Méhaut, R. Velazco, "Radiation Experiments on a 28nm Single-Chip Many-core Processor and SEU error-rate prediction", IEEE Trans. Nucl. Sci., vol. 64, no 1, pp. 483-490, 2017

[3] P. Ramos, V. Vargas, M. Baylac, F. Villa, S. Rey, J.A. Clemente, N.E. Zergainoh, J.F. Méhaut, R. Velazco, "Sensitivity to Neutron Radiation of a 45 nm SOI Multi-Core Processor", Proc. Radiation and its Effects on Components and Systems (RADECS), pp. 135-138, September 2015.

[4] P. Ramos, V. Vargas, M. Baylac, F. Villa, S. Rey, N.E. Zergainoh, R. Velazco (2017), "Error-rate prediction for applications implemented in Multi-core and Many-core processors", Accepted for publication in Proc. Radiation and its Effects on Components and Systems (RADECS 2017), Geneva - Switzerland, 2-6 October 2017

[5] V. Vargas, P. Ramos, R. Velazco, J.F. Méhaut, N.E. Zergainoh, "Evaluating SEU fault-injection on parallel applications implemented on multicore processors", In Proc. 6th Latin American Symposium on Circuits & Systems (LASCAS), Bariloche - Argentina, February 20-23, 2015

[6] V. Vargas, P. Ramos, W. Mansour, R. Velazco, N.E. Zergainoh, J.F. Méhaut, "Preliminary results of SEU fault-injection on multicore processors in AMP mode", In Proc. 20th International On-Line Testing, pp. 194-197. September 2014

[7] W. Mansour, P. Ramos, R. Ayoubi, R. Velazco (2014), "SEU fault-injection at system level: method, tools and preliminary results". 15th Latin American Test Workshop - LATW, pp. 1-5

[8] F. Villa, M. Baylac, S. Rey, O. Rossetto, W. Mansour, P. Ramos, R. Velazco, G. Hubert, "Accelerator-Based Neutron Irradiation of Integrated Circuits at GENEPI2 (France)," in Proc. Radiation Effects Data Workshop, July, pp. 1-5. July 2014

[9] A. Kchaou, Y. El Hadj, R. Tourki, F. Bouesse, P. Ramos, R. Velazco, "A deep analysis of SEU consequences in the internal memory of LEON3 processor", in Proc. 17th Latin American Test Symposium (LATS), pp. 178, April 2016

[10] V. Vargas, P. Ramos, V. Ray, C. Jalier, B. Dupont de Dinechin, M. Baylac, F. Villa, S. Rey, B. Minguez, R. Velazco (2016), "First Results of Radiation Experiments on a 28nm Single-Chip Many-core Processor", IEEE 53rd Nuclear and Space Radiation Effects Conference (NSREC), July 11-15, Portland, Oregon

[11] R. Velazco, S. Rezgui, R. Ecoffet, "Predicting Error Rate for Microprocessor-Based Digital Architectures through C.E.U. (Code Emulating Upsets) Injection," IEEE Trans. Nucl. Sci., vol. 47, pp. 2405-2411, December 2000

NMR-MPar: A Fault-Tolerance Approach for Multi-Core and Many-Core Processors

Keywords: fault tolerance; many-core; multi-core; partitioning; redundancy; reliability; fault injection

Members: V. Vargas, P. Ramos, J.F. Méhaut, R. Velazco

Cooperations: KALRAY

Contracts: CAPACITES (Calcul Applications Critiques en Temps et Sûreté), projet LEOC (Logiciel et Objets Connectés)

1. Introduction

Multi-core and many-core processors are a promising solution to achieve high performance by maintaining a lower power consumption. However, the degree of miniaturization makes them more sensitive to soft-errors. To improve the system reliability, this work proposes a fault-tolerance approach based on redundancy and partitioning principles called N-Modular Redundancy and M-Partitions (NMR-MPar). By combining both principles, this approach allows multi-/many-core processors to perform critical functions in mixed-criticality systems. Benefiting from the capabilities of these devices, NMR-MPar creates different partitions that perform independent functions. For critical functions, it is proposed that N partitions with the same configuration participate of an N-modular redundancy system. In order to validate the approach, a case study is implemented on the KALRAY Multi-Purpose Processing Array (MPPA)-256 many-core processor running two parallel benchmark applications. The traveling salesman problem and matrix multiplication applications were selected to test different device's resources. The effectiveness of NMR-MPar is assessed by software-implemented fault-injection. For evaluation purposes, it is considered that the system is intended to be used in avionics. Results show the improvement of the application reliability by two orders of magnitude when implementing NMR-MPar on the system. Finally, this work opens the possibility to use massive parallelism for dependable applications in embedded systems.

2. N-Modular Redundancy and M-Partitions approach

The partitioning principle is used to create different partitions that co-exist in the same device. This is typically done by the hypervisor, which provides virtual CPUs to each partition. In contrast, the NMR-MPar approach proposes a physical resource distribution to each partition, to minimize the propagation of faults producing dysfunctions in other resources or cores. Each partition can be setup as a mono- or multi-core running on different multiprocessing modes: Asymmetric Multi-Processing (AMP) or Symmetric Multi-Processing (SMP) mode with different programming models: bare-metal, OpenMP, Portable Operating System Interface uniX (POSIX), etc. Consequently, there is a considerable versatility for the system configuration that is enhanced by the number of cores comprising the device. It is important to note that each partition can be configured in bare-metal or with an independent OS.

In bare-metal, no OS is used, then the programmer uses the functions provided by the manufacturer to access hardware resources. There is no abstraction layer from the hardware architecture.

Furthermore, it is proposed for critical functions that N partitions with the same configuration participate in an N-Modular redundancy system. Thus, several partitions execute the same application, and the results are used by a voter to build a fault-tolerant system. The voter system can include one or more cores of the device, or an external one if desired. Depending on the partition of the device, several N-modular redundancy systems may run on it concurrently. Figure 1 illustrates an example of this approach implemented on a multi-core having 16 processor cores. The example establishes seven partitions (M = 7). Partitions P0-P2 are quad-core, while P3-P6 are mono-core. The quad-core partitions are part of a TMR system. Each one of them runs in SMP mode executing the parallel Application 1. On the other side, mono-core partitions (P3-P5) are part of another TMR running Application 2. The last mono-core partition P6 is the voter of both TMRs.

The latter is configured in bare-metal to reduce the use of resources, minimizing the impact of faults. By using this approach, temporal and spatial isolation of the applications is guaranteed by this type of partitioning. The security of each partition against unauthorized access and data modification can be obtained by the configuration of the system.

3. Case-Study: NMR-MPar Implemented on the MPPA-256 Processor

This case-study implements the NMR-MPar approach to improve the reliability on parallel applications running on the MPPA-256 many-core processor. The latter is a clusterized device that implements two Input Output cluster (IOs) for external communication and minimal processing and sixteen Compute Clusters (CCs) exclusively for processing. The IO cluster is comprised by eight cores called Resource Managers (RMs), while each CC consists of one RM in charge of managing the resources and sixteen Processing Engines (PEs) cores for computing. Both types of clusters include a private Static Memory (SMEM). The intra-cluster communication is achieved by buses, while the inter-cluster communication is performed by using two

Network-on-Chips (NoCs). The great configuration flexibility of the MPPA allows running independent applications per cluster or programming multi-cluster applications in a classic master/slave scheme, where the IO cluster performs as the master. For inter-cluster communication, the manufacturer provides a library with a set of functions for data exchanges through the MPPA NoC. Since, the MPPA is a coprocessor, a CPU host (HOST) is needed to manage it. The communication between the HOST and the MPPA is achieved using specific drivers provided by the manufacturer. Two types of distributed applications were evaluated: CPU-bound and memory bound, the Traveling Salesman Problem (TSP) and the Matrix Multiplication (MM), respectively. The importance of selecting two different natures of application lies in testing different complementary on-chip resources, which allows verifying the effectiveness of the approach under different scenarios.

4. Results

The validation of the proposed approach was done by SWIFI, which is a useful technique to emulate SEU effects. The adopted approach was proven in our previous works concerning fault-injection for multi-core processors. It is important to note that the objective of this approach is to reproduce the SEUs effects, which are bit-flips in memory cells caused by natural radiation. This case-study implements one core of the device as the fault-injector, the RM3 core being of IO Cluster 0. Fault-injection campaigns for both applications are devoted to inject faults in the General Purpose Registers (GPRs) and 15 System Function Registers (SFRs) of the PEs belonging to the compute clusters. The fault-injection campaigns consider the emulation of one SEU per execution.

For the scenario including a supervisor, only the runs resulting in exceptions with the same fault-injection parameters were re-executed, this to evaluate the behavior of the system in this scenario under the same conditions. It was thus necessary to inject 204 additional faults in the TSP and 4393 in the MM. Results show that the number of exceptions was drastically reduced from two hundred four to one and from four thousands three hundred ninety three to one hundred forty one, respectively. The details regarding the exceptions for both scenarios are summarized in following table.

Scenario	Core Stuck	Segmentation Fault	Device Exit	Total Exceptions
TSP-no-tty	39	165	0	204
TSP-tty	0	0	1	1
MM-no-tty	402	3948	43	4393
MM-tty	0	95	46	141

Table 1: Exceptions produced by fault-injection on the 4MR-8Par applications

5. Future work

The NMR-MPar can be implemented in other multi-/many-core processors, the implementation being specific for each device. Evaluation results of the implemented case study are encouraging since they open the possibility to use massive parallelism for dependable applications in embedded systems.

6. References

[1] TOP 500 Supercomputer list November 2017. Available online: https://www.top500.org/lists/2017/11/ (accessed on 16 March 2018)

[2] Across. Advanced Cockpit for Reduction Of StreSs andWorkload. Available online: http://www.acrossfp7. eu (accessed on 16 March 2016)

[3] MultiPARTES. Multi-Cores Partitioning for Trusted Embedded Systems. Available online: http://www. multipartes.eu (accessed on 14 April 2017)

[4] S. Guertin, "Initial SEE Test of Maestro". Presented at the 5th Workshop on Fault-Tolerant Spaceborne Computing Employing New Technologies, Albuquerque, NM, USA 29 May 1-June 2012. Available online: https://trs.jpl.nasa.gov/bitstream/handle/2014/42682/12-2145_A1b.pdf?sequence=1 (accessed on 16 March 2018)

[5] S.S. Stolt, E. Normand, "A Multicore Server SEE Cross Section Model", IEEE Trans. Nucl. Sci. 2012, 59, 2803-2810

[6] D.A.G. Oliveira, P. Rech, H.M. Quinn, T.D. Fairbanks, L. Monroe, S.E. Michalak, C. Anderson-Cook, P.O.A Navaux, L. Carro, "Modern GPUs Radiation Sensitivity Evaluation and Mitigation Through Duplication With Comparison". IEEE Trans. Nucl. Sci. 2014, 61, 3115-3122

[7] P. Ramos, V. Vargas, M. Baylac, F. Villa, S. Rey, J.A. Clemente, N.E. Zergainoh, J.F. Méhaut, R. Velazco, "Evaluating the SEE sensitivity of a 45 nm SOI Multi-core Processor due to 14 MeV Neutrons", IEEE Trans. Nucl. Sci. 2016, 63, 2193-2200

[8] V. Vargas, P. Ramos, V. Ray, C. Jalier, R. Stevens, B. Dupont de Dinechin, M. Baylac, F. Villa, S. Rey, N.E. Zergainoh et al., "Radiation Experiments on a 28 nm Single-Chip Many-core Processor and SEU error-rate prediction", IEEE Trans. Nucl. Sci. 2016, 99, 1-8

[9] R. Hyman, K. Bhattacharya, N. Ranganathan, "Redundancy Mining for Soft Error Detection in Multicore Processors", IEEE Trans. Comp. 2011, 60, 1114-1125

[10] T.N. Vijaykumar, I. Pomeranz, K. Cheng, "Transient Fault Recovery using Simultaneous Multithreading", In Proceedings of the 29th Annual International Symposium on Computer Architecture, Istanbul, Turkey, November 18-22, 2002; pp. 87-98

[11] C. Basile, Z. Kalbarczyk, R.K. Iyer, "Active Replication of Multithreaded Applications", IEEE Trans. Parallel Distrib. Syst. 2006, 17, 448-465

[12] H. Mushtaq, Z. Al-Ars, K. Bertels, "Efficient Software-Based Fault Tolerance Approach on Multicore Platforms", In Proceedings of the Design, Automation & Test in Europe Conference, Grenoble, France, March 18-22, 2013, pp. 921-926

[13] S. Reinhardt, S. Mukherjee, "Transient fault detection via simultaneous multithreading", In Proceedings of the 27th International Symposium on Computer Architecture (IEEE Cat. No.RS00201), Vancouver, BC, Canada, 10-14 June 2000; pp. 25-36

[14] S. Mukherjee, M. Kontz, S. Reinhardt, "Detailed design and evaluation of redundant multi-threading alternatives", In Proceedings of the 29th Annual International Symposium on Computer Architecture, Anchorage, AK, USA, 25-29 May 2002; pp. 99-110

[15] A. Holler, T. Rauter, J. Iber, G. Macher, C. Kreiner, "Software-Based Fault Recovery via Adaptive Diversity for Reliable COTS Multi-Core Processors", In Proceedings of the 6th International Workshop on Adaptive Self-tuning Computing Systems, Prague, Czech Republic, January 18, 2016; pp. 1-6

Assessing the Static and Dynamic Sensitivity of a Commercial Off-the-Shelf Multicore Processor for Noncritical Avionic Applications

Keywords: Reliability, Radiation ground test, Multi-core Processor, Single Event Effect, Single Event Upset

Members: P. Ramos, V. Vargas, N.E. Zergainoh, R. Velazco

Cooperations:

Contracts:

1. Introduction

Multicore processors are a suitable solution for achieving high performance and reliability without increasing significantly the power consumption. Its processing capacity and redundancy capabilities make them appropriate devices for implementing fault-tolerant mechanisms [1, 2]. Hence, avionic industries are interested in incorporating these devices in their systems [3]. However, the high degree of miniaturization (nanometer-scale) of multicores increases their vulnerability to the effects of natural radiation. This radiation may result in transient and permanent failures called single event effects (SEEs). Among them, the single event upset (SEU) is the most representative, since it may produce the modification of the content of a memory cell [4]. For this reason, manufacturers are enhancing fabrication processes and architectural designs. Silicon-on-Insulator (SOI) is a clear example of technology improvements, implemented to face traditional bulk CMOS drawbacks [5, 6]. Radiation hardening by design (RHBD) techniques are also used to mitigate SEU consequences [7]. The implementation of error correcting codes (ECCs) and parity to protect the internal memory of the processors is useful but not enough in presence of multiple bit upsets (MBUs). Another well-known RHBD technique is the triple modular redundancy (TMR) which significantly improves the reliability of the system. Nevertheless, having more robust or dedicated components implies a considerable increase in costs. Consequently, an important challenge for aircraft industries is the integration of commercial-off-the-shelf (COTS) multicore processors due to budget and availability issues [8]. The current work assesses the effects of neutron radiation on a multicore processor which does not implement protection mechanisms in its internal memories. This is achieved by means of two accelerated radiation experiments. The first one aims at evaluating the device (hardware) sensitivity, while the second one evaluates the application (software) sensitivity. Part of the results of current research has been presented in [9].

2. Methodology

Accelerated radiation ground testing allows performing the analysis of the sensitivity to radiation of electronic devices through artificial radiation environments. It is the fastest way to obtain statistically meaningful data in a short period of time, since the more particles hit the component, the more SEEs are observed [16]. The reproducibility of the experiment is also another major advantage of this strategy. Consequently, this work considers two models of tests for evaluating the sensitivity of a multicore processor: a static test in order to obtain the intrinsic sensitivity of the device's memory cells and a dynamic test for evaluating the dynamic response of the implemented application [17].

In this work, experimental tests have been conducted with 14MeV neutron radiation to emulate the effects of high-energy neutrons present at avionic altitudes, since neutrons are the most representative particles in the Earth's atmosphere. Reference [18] discusses the relevance of using the 14MeV neutron test to characterize the SEU sensitivity of digital devices. Sections 3 and 6 of the JESD89A document of the *JEDEC standard* were used as a base protocol for the experimental tests [19].

Static Test. This test aims at estimating the intrinsic sensitivity to SEE of the memory cells of a processor. The device under test (DUT) is placed facing the center of the target perpendicularly to the beam axis at a distance depending on the required radiation flux. Typically, the method consists in writing a predefined pattern in the memory and accessible registers of the processor via the instruction set (load and store). Once finished the initialization, the DUT is irradiated and the program checks periodically the memory locations along the radiation test to detect upset events. If an upset is detected, the program writes the correct pattern in the associated memory location and logs the results to an external host via Ethernet ports.

During the static test, all the sensitive zones are exposed to radiation at the same time, which do not represent the real behavior of the circuit since not all the memory resources are used simultaneously when an application is executed. For this reason, the static test provides the worst-case estimation of the device sensitivity. As a result of this test, the static cross section (σ_{STATIC}) of the device is obtained. It is defined as the number of detected upset events divided by the fluence, which is the neutron flux integrated in time.

The elementary data pattern for memory circuits is

a logical checkerboard. All zeros and all ones is also a common pattern used during the radiation test. However, some memories such as DRAMs usually have a favorite error failure, either 0->1 or 1->0. For this reason, for testing when there is no a priori information about the component, the test pattern have to balance the number of 0's and 1's. Thus, the selected pattern for the static test was 0x55AA55AA.

Regarding the exposure time to radiation of the device, it is important to consider that the probability of having an upset event during a given period of time is a stochastic process that follows a Poisson distribution. Thus, the waiting time between the read operations in the static test can be validated by analyzing the distribution of the number of events per unit of time. If the obtained distribution does not follow the Poisson law, the waiting time should be adjusted.

Dynamic Test. The goal of this test is to estimate the SEE dynamic response of an application running on a processor. As a result of the experiment, the dynamic cross section (σ_{DYN}) is obtained. Unlike the static test, it only evaluates the memory cells used by the application. The method consists in the periodical execution of an application while the processor is being irradiated to induce SEE. Once finished each execution of the program, results are compared with a set of correct values previously obtained, in order to detect errors. The experiment is launched and monitored using a *host* computer located outside the armored chamber. The communication between the external host and the Parallella host is achieved by means of the Ethernet port using the Linux *ssh* communication command. All detected errors are logged and transmitted to the external host which stores the results.

3. Results

Radiation experiments performed on the Epiphany E16G301 are very interesting compared to similar works targeting other multicore processors, since errors produced by SEE are clearly identified as they are not masked by protection mechanisms such as ECC or parity. This fact allows a better analysis of the behavior of the device in presence of SEEs.

Concerning the limitations of the experiments, there are two points to consider:
(i) The E16G301 processor does not have direct access to *printf* function for logging results. For this reason it has to write the information about observed events in the external DDR memory of the board. *is information is logged by the *host* processor (ARM).
(ii) The physical distance between the E16G301 multicore and the *host* processor in the Parallella board is less than one centimeter. It was thus necessary to limit the neutron flux for avoiding particles affecting the *host* processor and other circuitries.

Experimental Setup. The DUTwas placed at a distance of 38.5 ± 0.5cm to the target. The neutron beam energy was 14MeV with an estimated flux of 7.2×10^4 $n \cdot cm^{-2} \cdot s^{-1}$ with an error of $\pm 0.1 \times 10^4$ $n \cdot cm^{-2} \cdot s^{-1}$. Special attention was required to protect the rest of the platform components from radiation. For that, the E16G301 multicore was irradiated through a small window on a 5 cm thickness polypropylene block intended to protect the readout platform.

SEE Type	Test 1	Test 2	Test 3	Consequences
SBU	29	17	23	Bit-flip
MCU (2)	1	5	1	Bit-flip
SEFI	3	2	0	Hangs
Total	33	24	24	—

Table 1: Results of the static radiation test campaigns

4. Future work

The Adapteva Epiphany multicore processor will be proposed to be used for image processing in a military aircraft. In parallel, another module containing the Epiphany processor will execute a memory-bound application in order to detect SEUs produced in real operating environment.

5. References

[1] R. Hyman, K. Bhattacharya, N. Ranganathan, "Redundancy mining for soft error detection in multicore processors", IEEE Transactions on Computers, vol. 60, no. 8, pp. 1114-1125, 2011

[2] E. Kim, N. Shanbhag, "Soft N-modular redundancy" IEEE Transactions on Computers, vol. 61, no. 3, pp. 323-336, 2012

[3] Certification Authorities Software Team (CAST), "Position paper cast-32 multi-core processors," November 2016, https:// www.faa.gov/aircraft/air_cert/design_approvals /air_software/cast/cast_papers/media/cast-32A.pdf

[4] R. Baumann, "Soft errors in advanced computer systems," IEEE Design and Test of Computers, vol. 22, no. 3, pp. 258-266, 2005

[5] B. Vandana,M. Siva Mukar, "Study of VLSI bulk CMOS and SOI technologies", International Journal of Scientific and Research Publications, vol. 3, no. 2, 2013

[6] G. Gasiot, V. Ferlet-Cavrois, J. Baggio et al., "SEU sensitivity of bulk and SOI technologies to 14 MeV neutrons," IEEE Transactions on Nuclear Science, vol. 49, no. 6, pp. 3032-3037, 2002

[7] W. Hughes, "Single Event Effects Mitigation Techniques Report", DOT/FAA/TC-15/62, Federal Aviation Administration, Washington, DC, USA, 2016

[8] P. Bieber, F. Boniol, M. Boyer, E. Noulard, C. Pagetti, "New Challenges for Future Avionic Architectures", Onera, Palaiseau,France, AerospaceLab, 2012

[9] P. Ramos, V. Vargas, R. Velazco, N.E. Zergainoh, "Assessing the Static and Dynamic Sensitivity of a Commercial-Off-De-Shelf Multi-Core Processor for Non-Critical Avionic Applications", Nano Science and Nanotechnology Week,Quito, Ecuador, 2017

[10] S.S. Stolt, E. Normand, "A multicore server SEE cross section model," IEEE Transactions on Nuclear Science, vol. 59, n° 6, pp. 2803-2810, 2012

[11] S. Guertin, "Initial SEE Test of Maestro", National Aeronautics and Space Administration, Jet Propulsion

Laboratory, Pasadena, CA, USA, 2012, http://trs-new.jpl.nasa.gov/dspace/bitstream/ 2014/42682/1/12-2145_A1b.pdf

[12] T. Santini, P. Rech, G. Nazar, L. Carro, F.R. Wagner, "Reducing embedded software radiation-induced failures through cache memories" in Proceedings of the 19th IEEE European Test Symposium (ETS), pp. 1-6, Paderborn, Germany, May 2014

[13] D.A.G. Oliveira, P. Rech, H.M. Quinn et al., "Modern GPUs radiation sensitivity evaluation and mitigation through duplication with comparison," IEEE Transactions on Nuclear Science, vol. 61, no. 6, pp. 3115-3122, 2014

[14] P. Ramos, V. Vargas, M. Baylac et al., "Evaluating the SEE sensitivity of a 45 nm SOI multicore processor due to 14 MeV neutrons," IEEE Transactions on Nuclear Science, vol. 63, no. 4, pp. 2193-2200, 2016

[15] V. Vargas, P. Ramos, V. Ray et al., "Radiation experiments on a 28 nm single-chip many-core processor and SEU error-rate prediction," IEEE Transactions on Nuclear Science, vol. 64, no. 1, pp. 483-490, 2017

[16] H. Quinn, "*e use of accelerated radiation testing for avionics," AIP Conference Proceedings, vol. 1525, 2013

[17] S. Rezgui, R. Velazco, R. Ecoffet, S. Rodriguez, J. Mingo, "Estimating error rates in processor-based architectures", IEEE Transactions on Nuclear Science, vol. 48, no. 5, pp. 1680-1687, 2001

[18] F. Miller, C. Weulersse, T. Carriere et al., "Investigation of 14 MeV neutron capabilities for SEU hardness evaluation", IEEE Transactions on Nuclear Science, vol. 60, no. 4, pp. 2789-2796, 2013

[19] Jesd89A, Measurement and Reporting of Alpha Particle and terrestrial Cosmic Ray-Induced Soft Errors in Semiconductor Devices, JEDEC Solid State Technology Association, Arlington,VA, USA, December 2016

Evaluation by Neutron Radiation of the NMR-MPar Fault-Tolerance Approach Applied to Applications Running on a 28-nm Many-Core Processor

Keywords: radiation ground testing; many-core processor; single event effect; single event upset; partitioning; redundancy

Members: V. Vargas, P. Ramos, R. Velazco

Cooperations: KALRAY

Contracts: CAPACITES (Calcul Applications Critiques en Temps et Sûreté), projet LEOC (Logiciel et Objets Connectés)

1. Introduction

Currently, there is a special interest in validating the use of Commercial-Off-The-Shelf (COTS) multi/many-core processors for critical applications thanks to their high performance, low power consumption and affordability. However, the continuous shrinking of transistor geometry and the increasing complexity of these devices dramatically affect their sensitivity to natural radiation, and thus diminish their reliability. One of the most common effects produced by natural radiation is the Single Event Upset which is the bit-flip of a memory content producing unexpected results at application-level. For this reason, manufacturers and users implement hardware and software error-mitigation techniques on multi/many-core processors. In this context, the present work aims at evaluating a new fault-tolerance approach based on N-Modular redundancy (NMR) and partitioning called NMR-MPar by means of 14 MeV neutron radiation ground testing in order to emulate the effects of high-energy neutrons present at avionics altitudes. For evaluation purposes, a case-study is implemented on the 28 nm CMOS KALRAY MPPA-256 many-core processor running two complementary benchmarks applications: a distributed Matrix Multiplication and the Travel Salesman Problem. Radiation experiments were conducted in GENEPI2 particle-accelerator.

The correctness of the results of the application when an error s detected confirms the approach's effectiveness and boosts their usage on avionics applications.

2. Materials and Methods

The selected device is the KALRAY MPPA-256 many-core processor due to its advanced CMOS 28 nm manufacturing technology and its architecture, which is similar to the Shen Wei SW26010 (260 cores) many-core processor, which is the base processor of the second TOP500 super computer list (June 2018). In addition, the MPPA many-core is focus of interest of embedded community to study the possibility of its use in critical real-time embedded systems. For instance, CAPACITES is a project that gathers French academics and industrial partners to analyse the possibility of using MPPA many-core processors for critical embedded systems.

The MPPA-256 considered in this work is the second version called Bostan. This processor is manufactured in TSMC CMOS 28HP technology. The processor operates between 100 MHz and 600 MHz, for a typical power ranging between 15 W and 25 W. Its peak floating-point performances at 600 MHz are 634 GFLOPS and 316 GFLOPS for single and double-precision respectively. It integrates 256 Processing Engines (PEs) cores and 32 Resource Managers (RMs) cores distributed in a clustered architecture. All cores are based on the same VLIW 32-bit/64-bit architecture. The MPPA-256 comprises two Input Output cluster (IO) for external communication and minimal processing, and sixteen Compute Clusters (CC) exclusively for processing. The communication intra-cluster is achieved by buses while the communication inter-cluster is performed by a wormhole switching network-on-chip (NoC) with 32 nodes and a 2D torus topology.

This work proposes the use of two types of distributed applications: a CPU-bound and a memory-bound. The selected CPU-bound application is the Traveling Salesman Problem (TSP), a Non-deterministic Polynomial (NP) hard problem very used for evaluating computing system optimization. This application aims at finding the shortest possible route to visit n cities, visiting each city exactly once and returning to the departure city. To solve the problem there are several proposals. This work uses a brute force exact algorithm based on a simple heuristic. The implemented version of TSP on the MPPA by authors of was used as a basis. The light dark format in the illustration of the solution means that this possible route was not completely explored. It was discarded when the algorithm arrive to city 4, because the path distance until this point (65) is larger than other complete solution (62).

On the other hand, the Matrix Multiplication (MM) was chosen as memory-bound application. The MM is widely used for solving scientific problems related to linear algebra, such as systems of equations, calculus of structures and determinants

among others. Concerning avionic applications, MM is used for image processing, filtering, adaptive control, and navigation and tracking. In addition, the parallelism of MM is one of the most fundamental problems in distributed and High-Performance Computing. There are many approaches proposed to optimize performance. The present work implements the approach divide and conquer. The selected application is a collaborative 256 x256 matrix multiplication.

This work evaluates the fault-tolerance approach called NMR-MPar under neutron radiation. This approach uses redundancy and partitioning as basic concepts to improve the reliability of applications running on multi-core and many-core processors. NMR-MPar takes advantage of the multiplicity of cores to implement redundancy techniques that allow masking faults. Complementary, partitioning protects against not authorized access and data modification by temporal and spatial isolation of each partition. The proposed approach was presented in a previous work [17].

3. Results

Soft error consequences of the six scenarios are illustrated in Figure 1. From the results, it is possible to observe a different behavior of the NMR-MPar depending on the application. For the MM, all the errors were masked, while for the TSP, there remain to be some timeouts and exceptions that cannot be masked. This can be explained by the fact that parity errors in cache memories were not masked by the approach since the memory was not invalidated. As it was aforementioned, the default code for handling-trap-code was not overwritten by the user-programmed wrap-code, so when a parity error is detected, the operating system produces a timeout or an exception. Since the execution time of TSP is larger than the MM, and the use of cache memories is different because of the nature of each application, TSP is more sensible to parity errors in cache memories. Results show more exceptions and timeouts for the TSP scenarios solving the 17 cities problem caused by the longer exposure time of the application which is around seven times the other applications (TSP-16 cities and MM).

Furthermore, an increase in the number of exceptions when using NMR-MPar in the TSP-17 cities is observed. This can be explained by both a longer exposure time and an increase in the sensitive zone, so the possibility of having parity errors in cache memories has increased. Therefore, in order to decrease the number of timeouts and exceptions of the application, it is necessary that the handling-trap-code considers the invalidation of cache memories when a parity error is detected. This politic of invalidation should be considered by manufacturers in proprietary systems or by programmers in open systems. In all the cases, by using the approach, the application reliability was increased. The most dangerous consequences “application erroneous result” were masked, while timeouts and exceptions were detected by the system. Timeouts were produced when an IO Trap was executed. The possibility of having IO Traps depends on the use of the Resource Manager cores of the IO Cluster. By minimizing its use, the system reliability could be improved.

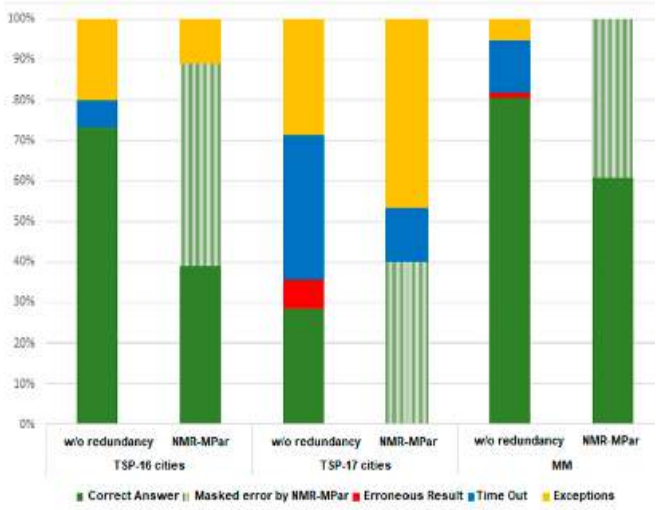

Figure 6: Comparison of the consequences of the evaluated applications

4. References

[1] R. Gaillard, "Single Event Effects: Mechanisms and Classification", In Soft Errors in Modern Electronic Systems; Springer: New York, NY, USA, 2011

[2] S. Mittal, J.S. Vetter, "A survey of techniques for modeling and improving reliability of computing systems", IEEE Trans. Parallel Distrib. Syst. 2016, 27, 1226-1238

[3] T. Vijaykumar, I. Pomeranz, K. Cheng, "Transient Fault Recovery using Simultaneous Multithreading", In Proceedings of the 29th Annual International Symposium on Computer Architecture, Anchorage, AK, USA, 25-29 May 2002

[4] C. Basile, Z. Kalbarczyk, R.K. Iyer, "Active Replication of Multithreaded Applications", IEEE Trans. Parallel Distrib. Syst. 2006, 17, 448-465

[5] H. Mushtaq, Z. Al-Ars, K. Bertels, "Efficient Software-Based Fault Tolerance Approach on Multicore Platforms", In Proceedings of the Design, Automation & Test in Europe Conference, Grenoble, France, March 18-22, 2013

[6] S. Reinhardt, S. Mukherjee, "Transient fault detection via simultaneous multithreading", In Proceedings of the 27th International Symposium on Computer Architecture (IEEE Cat. No.RS00201), Vancouver, BC, Canada, 14 June 2000

[7] S. Mukherjee, M. Kontz, S. Reinhardt, "Detailed design and evaluation of redundant multi-threading alternatives", In Proceedings of the 29th Annual International Symposium on Computer Architecture, Anchorage, AK, USA, May 25-29, 2002

[8] A. Shye, J. Blomstedt, T. Moseley, V.J. Reddi, D. Connors, "PLR: A Software Approach to Transient Fault Tolerance for Multicore Architectures", IEEE Trans. Dependable Secure Comput. 2009, 6, 135-148

[9] C. Bolchini, A. Miele, D. Sciuto, "An adaptive approach for online fault management in many-core architectures", In Proceedings of the 2012 Design, Automation Testing

Europe Conference Exhibition (DATE), Dresden, Germany, March 12-16, 2012

[10] M. Alhakeem, P. Munk, R. Lisicki, H. Parzyjegla, H. Parzyjegla, G. Muehl, "Framework for Adaptive Software-Based Reliability in COTS Many-Core Processors", In Proceedings of the ARCS2015 - 28th International Conference on Architecture of Computing Systems, Porto, Portugal, March 24-27, 2015

[11] E. Kim, N. Shanbhag, "Soft N-Modular Redundancy", IEEE Trans. Comput. 2012, 61, 323-336

[12] J. Walters, R. Kost, K. Singh, J. Suh, S. Crago, "Software-based fault tolerance for the Maestro many-core processor", In Proceedings of the 2011 Aerospace Conference, Big Sky, MT, USA, March 5-12, 2011

[13] A. Löfwenmark, S. Nadjm-Tehrani, "Challenges in Future Avionic Systems on Multi-Core Platforms", In Proceedings of the IEEE International Symposium on Software Reliability Engineering Workshops, Naples, Italy, November 3-6, 2014

[14] M. Mollison, J. Erickson, J. Anderson, S. Baruah, J. Scoredos, "Mixed-Criticality Real-Time Scheduling for Multicore Systems", In Proceedings of the 10th IEEE International Conference on Computer and Information Technology, Bradford, UK, 29 June-1 July 2010

[15] M. Panic, E. Quiñones, P. Zavkov, C. Hernandez, J. Abella, F. Cazorla, "Parallel many-core avionics systems", In Proceedings of the International Conference on Embedded Software (EMSOFT), Jaypee Greens, India, October 12-17, 2014

[16] S. Trujillo, A. Crespo, A. Alonso, J. Pérez, "MultiPARTES: Multi-core partitioning and virtualization for easing the certification of mixed-criticality systems", Microprocess. Microsyst. 2014, 38, 921-932

[17] V. Vargas, P. Ramos, J.F. Méhaut, R. Velazco, "NMR-MPar: A fault-tolerance approach for multi-core and many-core processors", Appl. Sci. 2018, 8, 465

Onboard computer of a nanosatellite implemented in a many-core processor

Keywords: On-Board Computer, Many-core processors, Single Event Upset, Fault-Tolerance

Members: R. Velazco, D.C. Ardiles Saravia, I. Bhatia

Cooperation: Concordia University (Canada), Spectrum Aerospace (Germany), Kalray (Grenoble)

Contracts: OVNIPROM-SCUSI (Auvergne-Rhône-Alpes)

1. Context and goals

This project aims at building an On-Board Computer (OBC) for nanosatellites applications benefitting of the performances and fault-tolerance possibilities offered by the many-core processors. Indeed, in the average nanosatellites have an expected lifetime of one year and a half, but sometimes some nanosatellites live only a few months. This situation can be explained by the fact that OBC of nanosatellites are built from circuits Commercial Off-the-Shelf (COTS) and thus radiation effects may result in errors which are not detected/corrected and can in some cases make miss the satellite.

The continuous advances in integrated circuit manufacturing technologies allow nowadays building processors with a huge number of computing nodes, which may be used to form complete systems in a single chip. These circuits, called many-core processors, enable the implementation of parallel or distributed applications with high computing power and performance [1, 2]. This potential opens interesting doors in a wide range of domain areas, with special emphasis on aerospace projects, avionics, scientific instrumentation, standalone transport and the Internet of Things (IoT). Miniaturization, however, has a significant impact on reliability and this must be considered for any application where failures can have critical consequences. In fact, since the early 1990´s it has been extensively studied and demonstrated that integrated circuits are sensitive to the effects of energetic particles (heavy ions, protons, neutrons, etc.) present in the environment in which they will operate [3]. This which in the past strictly concerned space applications, due to the high energies of the cosmic rays which can perturb the functioning of the circuits, must nowadays be considered for any application implemented in circuits issued from advanced nanometric manufacturing technologies, whose failure can have critical consequences, even if it is intended to operate in the Earth's atmosphere at sea level. As an example it can be mentioned that several experiments aiming to highlight this issue were made in aircrafts, satellites and mountains, by the scientific communities related to this topic and as part of the researches of TIMA/RIS. These experiments have shown that the reliability of these components must be considered in any application implemented in very advanced devices such as multi and many-core processors [4, 5, 6].

Currently, tests performed in particle accelerators are the standard way to estimate the sensitivity to radiation of integrated circuits. This type of tests were the objective of the last two doctoral thesis developed and supported by TIMA/RIS and have been focused on estimating the intrinsic sensitivity of a many-core processors, the MPPA-256 (Massively Parallel Processor Array), produced by Kalray (Grenoble, France). A deep study of fault tolerance techniques that can be implemented, benefiting from the multiplicity of cores was achieved. It is important to mention that these researches lead to the first references concerning the radiation effects on a many-core processor that were disseminated in the scientific community related to the effects of radiation in circuits and integrated systems [7, 8].

Due to the significant growth of academic activities on nano-satellites, it was decided at TIMA/RIS to objectively explore the possibility of using many-core processors to implement the OBC of a nano-satellite considering the effects of radiation as a "threat" to reliability. This context has motivated the objective of the OVNIPROM (Ordinateur de Vol d'un Nanosatellite Implémenté dans un Processeur Many-core) project, to develop and validate the on-board computer of a nano-satellite with the MPPA-256 processor, including the fault-tolerance techniques necessary to guarantee the useful lifetime and efficiency of the mission. The advantages of MPPA-256 mentioned above are due to the fact that its a computer performance to 0.7 Tera FLOP (FLOP: floating-point operations per second).

2. Recent outcomes

The first year activities OVNIPROM were focused to the development of a manual, which describes in detail the programming environment of the MPPA-256 through the Eclipse software [Java Integrated Development Environment (IDE) and C/C++] containing information necessary for parallel and threaded programming of the MPPA-256.

A bibliographic and technical review of different On-Board Computer (OBC) commercial systems was also carried out.

Subsystems were studied to complement and to integrate an OBC, such as energy control (batteries and solar panels) and different topologies of magnetorquer. A study was also made on sensors (solar, temperature and altitude), elaborating a library of available electronic devices and circuits tolerant to radiation.

3. Future work

Starting from a Kalray development kit, adapt its configuration for the elaboration of an OBC, adding interfaces and subsystems that allow its application in the construction of a nano satellite.

Include fault-tolerant software and a real-time communication system, e.g. DTN (Delay Tolerant Network).

Implementation of an algorithm in MPPA-256 for an Energy Efficiency system.

Projects will be carried out to make possible the development of the nanosatellite and the possibility of its launch in 2020 in collaboration with the National Commission for Space Activities of Argentina (CONAE).

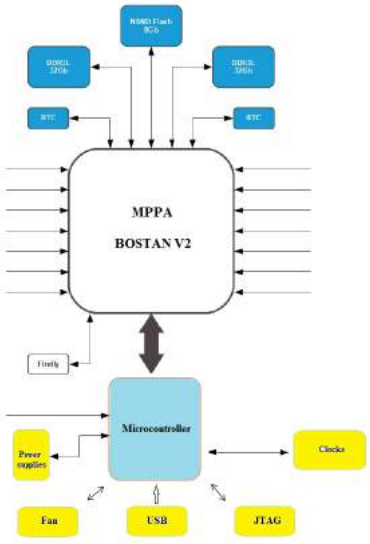

Figure 1: Kalray kit for applications

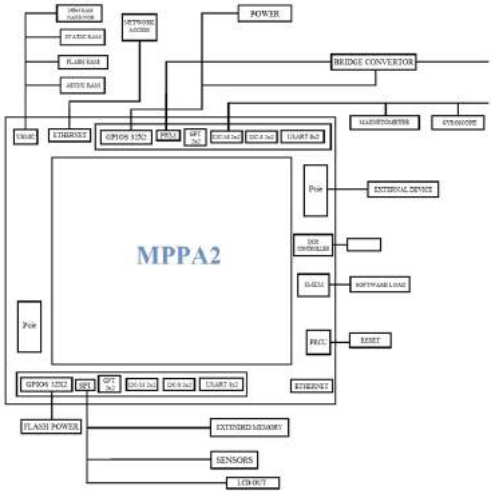

Figure 2: Configuration of an OBC containing the MPPA

4. References

[1] M.L. Shooman, "Reliability of Computer Systems and Networks: Fault Tolerance, Analysis and Design", John Wiley & Sons, Inc., New York, NY, USA, 2002

[2] B. Dupont de Dinechin, P. Guironnet de Massas, G. Lager, C. Leger, B. Orgogozo, J. Reybert, T. Strudel, "A Distributed Run-Time Environment for the Kalray MPPA-256 Integrated Manycore Processor", International Conference on Computational Science, ICCS 2013. Procedia Computer Science *Volume* 18, 2013, *Pages* 1654-1663

[3] V. Vargas Vallejo, PhD, RIS-TIMA, University Grenoble Alps - France. http://repositorio.educacionsuperior.gob.ec/bitstream/28000/4613/1/T-SENESCYT-01622. Accessed on: 05 Jan. 2018

[4] J.A. Clemente, F.J. Franco, F. Villa, M. Baylac, P. Ramos, V. Vargas, H. Mecha, J.A. Agapito, R. Velazco, "Neutron-Induced Single Events in a COTS Soft-Error Free SRAM at Low Bias Voltage", In Proceedings of IEEE European Conferences on Radiation Effects on Components and Systems (RADECS2015), pages 162–165, Sep. 2015

[5] J.A. Clemente, F.J. Franco, F. Villa, M. Baylac, P. Ramos, V. Vargas, H. Mecha, J.A. Agapito, R. Velazco, "Single Events in a COTS Soft- Error Free SRAM at Low Bias Voltage Induced by 15-MeV Neutrons", IEEE Trans. Nucl. Sci., 63(4):2072 – 2079, Aug. 2016

[6] J. Autran, P. Munteanu, P. Roche, G Gasiot, "Real-time Soft- Error Rate Measurements: A Review", Microelectronics Reliability, 54(8):1455– 1476, Feb. 2014

[7] V. Vargas, P. Ramos, R. Velazco, "Evaluation by Neutron Radiation of the NMR-MPar Fault-Tolerance Approach Applied to Applications Running on a 28-nm Many-Core Processor", ELECTRONICS 2018, 7(11), 312

[8] P. Ramos, V. Vargas, N.E. Zergainoh, R. Velazco, "Assessing the Static and Dynamic Sensytivity of a Commercial Off-the-Shelf Multicore Processor for Noncritical Avionic Applications", HINDAWI, Journal of Technology, Volume 2018, article ID 2926392 https:/doi.org/101155/2018/2926392

RuNS: A Runtime Fault-Tolerant Scheme for 3D Networks-on-Chip

Keywords: Fault-tolerant routing algorithm, 3D NetworksOn-Chip, TSV failure, partially connected 3D-NoC.

Members: A. Coelho, A. Charif, N.E. Zergainoh, R. Velazco

1. Context and goals

Global bus-based interconnect has become a bottleneck for future high-performance System-on-Chip. To alleviate the limitation of bus-based solution, Network-on-Chip has emerged as a promising alternative thanks to its scalability, high bandwidth, better throughput and lower power consumption. However, the conventional two-dimensional integrated circuit (2D-IC) has limited floor-planning choices with increasing number of processing elements attached, limiting the potential performance of two-dimensional Networks-on-Chip architecture. In order to extend 2D-IC capabilities, multiple layers of active devices are integrated using vertical high-speed interconnection in a three-dimensional integrated circuit (3D-IC) architecture. This approach permits the use of three-dimensional Networks-on-Chip (3D-NoC), which reduces the interconnection lengths and improve the overall Networks-on-Chip performance.

Among the vertical interconnection technologies, Through Silicon-Via (TSV) has been accepted as one of the most viable technology since it enables faster and more power efficient inter-layer communication across multiple stacked layers. However, a large area overhead imposed by the TSV interconnect pitch and the extra manufacturing cost limit the number of TSV into 3D-IC. Additionally, since TSVs present lower reliability, the risk of failure in a 3D-IC will increase as the number of TSVs increases. These limitations force the designer to adopt an architecture of 3D-NoC partially connected where only a subset of all vertical links will be connected either due to the high-cost of fabrication or the high rate of TSV failures.

In this context, many works have addressed the vertical link (elevator) failure problem adopting fault-tolerant techniques in the routing algorithm for 3D-NoC. However, as it is explained in the section II, most of the already existing routing scheme present one or more of the following problems: A large number of TSVs and Virtual Channels to recover from faults, specific routing rules that pose restrictions on the location and the selection of TSVs, and/or an offline mechanism to reconfigure the entire 3D-NoC after faults. All these techniques either significantly increase the hardware resources or need to drop packets in the reconfiguration phase until the network goes back to its stable condition.

These issues motivated us to develop an efficient and highly resilient routing scheme called RuNS (Runtime fault-tolerant 3D Networks-on-Chip Scheme) that can tolerate manufacture and runtime faults occurring in the vertical link of a partially connected 3D-NoC. RuNS requires one additional virtual channel along the North and East direction to be both adaptive and deadlock-free. Also, RuNS presents a set of eight bits per router which is combined with our TSVs status propagation and our fault-tolerant routing algorithm to search for a healthy elevator (vertical link). Additionally, RuNS uses a rerouting mechanism and an escape-buffer inside of each router to avoid dropping packets [1].

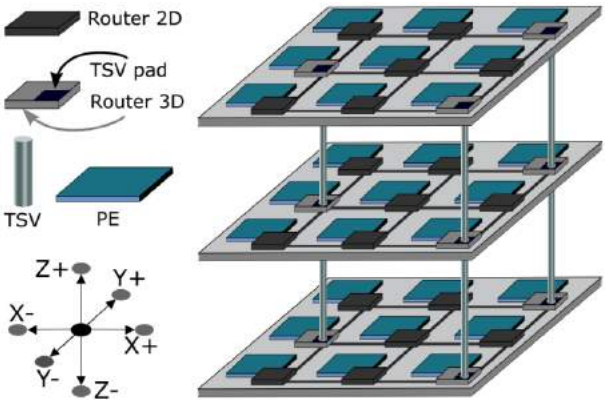

Figure 1: 3D-NoC Architecture

Target Architecture

We consider a partially connected 3D-NoC mesh architecture, wherein the routers include, in addition to the usual five ports (East/X+, West/X-, North/Y+, South/Y-, Local), either an Up/Z+ port, a Down/Z- port, or both (i.e. 5, 6 and 7 port configuration). The Up and Down ports of the router are connected vertically through TSVs, as shown in Figure 1. Each TSV pillar is connected in all layers and we call them “elevator”. That is, each elevator connects all layers and have the Up and Down port.

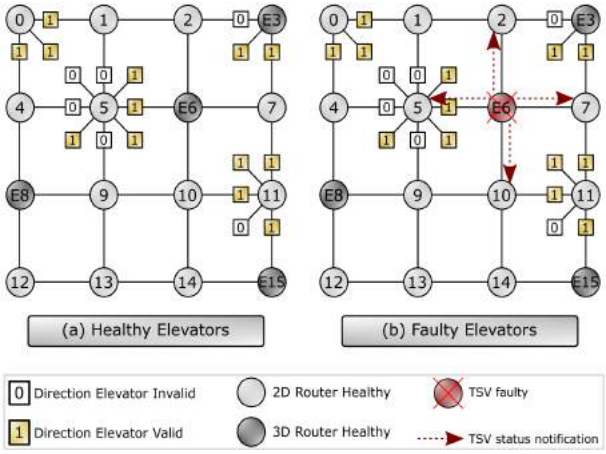

Figure 2: Elevator configurable bits for:
(a) Initially operationwith all healthy elevators and
(b) Runtime faulty TSV status propagation in one hop

Locating Healthy Elevator and Propagating Status

The challenge for the routing algorithm in a partially connected 3D-NoC is to find an elevator to deliver the packet through its source layer to its destination

layer. To deal with this challenge, we have extended the techniques proposed by [2]. Those techniques determine the possible elevator's position using some configuration bits inside of the routers. In our method, each router stores an 8-bit vector [North, East, South, West, NE, NW, SE, and SW] indicating the presence of at least an elevator in one direction. For example, the North or South bits are set if there is at least one elevator in the same column to the North or the South, respectively. The NE and NW, on the other hand, are used to indicate the existence of an elevator in the northeast or northwest directions, respectively. To illustrate how these 8-bits are set, let us consider the example shown in Figure 2(a) where only the bits pointing to elevators are set to one. In this example, the router 5 knows that there are elevators in the East(E6), NE(E3), SE(E15) and SW(E8) direction because the corresponding elevator position bits from its 8-bits vector are set to one: [01001011].

In order to propagate the elevator failures, as shown in Figure 2(b), TSV status notification signals are connected to the most adjacent routers of the same elevator's row and column allowing to share the status of the elevator. After detecting a TSV fault, each elevator transmits a signal to its closest neighbor router indicating that it cannot work as an elevator anymore. For example, let us assume that elevator E6 is faulty, as shown in Figure 2(b). In this case, the elevator E6 will send its signal statuses to routers 2, 5, 7, and 10 indicating that it is an elevator and has a fault.

Proposed Routing Algorithm

RuNS needs a total of eight virtual channels, which are distributed along X(X0+; X0−; X1+), Y(Y 0+; Y 0−; Y 1+), and Z(Z0−; Z0+). These channels are partitioned into three virtual networks, and each virtual network has an acyclic configuration to avoid deadlock. The virtual networks VN0 and VN2 use both the same positive direction distributed as follows: VN0 uses the X0+ and Y0+; VN2 uses X1+ and Y1+. The second virtual network (VN1) includes the remaining directions, i.e. the negative direction (X0−; Y0−) as well as the Up/Down direction (Z0+; Z0−). In addition to virtual network definitions, the packets must traverse virtual networks only in increasing order (VN0 → VN1 → VN2). Taking the Figure 3 as example, the proposed routing algorithm based on the virtual network definition is described as follows: When the source and destination are in the same layer, or the packet has reached its destination layer, the routing algorithm uses the negative direction first (VN1), if the destination is located at South, West, or Southwest of the source. Then, the packets are routed using the positive direction (VN2), if the destination is located at North, East, or Northeast of the source. When the source and destination are not in the same layer, the routing algorithm routes the packet using the 8-bits of Elevator location and the TSV status notification to search for a healthy elevator in the positive direction (VN0). If no elevator is found using the positive direction (VN0), the routing algorithm uses the negative direction (VN1) to search for a healthy elevator. It is worth mentioning that the packets should take VN1 to elevate to its destination layer.

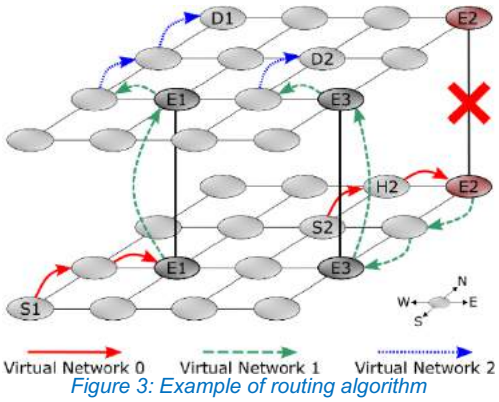

Figure 3: Example of routing algorithm with and without TSVs faulty

Simulation Results

Figure 4 shows the reliability comparison for the Elevator-First, CoBRA, and RuNS under the effect of double and triple faults. We can observe that for double faults the reliability of CoBRA is almost the same of RuNS. However, packets were lost in the CoBRA when failures occurred in the two EAST most elevators (E3 and E15 of the Figure 2(a)). It has occurred because some packets are dropped in the reconfiguration phase until the network backs to its stable condition. The analysis of triple faults occurring in the E3, E8, and E15, shows that RuNS performs better than CoBRA since RuNS can send all its packets to the healthy elevator E6, while CoBRA needs at least one elevator at East or West to avoid drop packets. The Elevator-First presents the worst scenario for double and triple faults because it cannot adapt itself to faults at runtime.

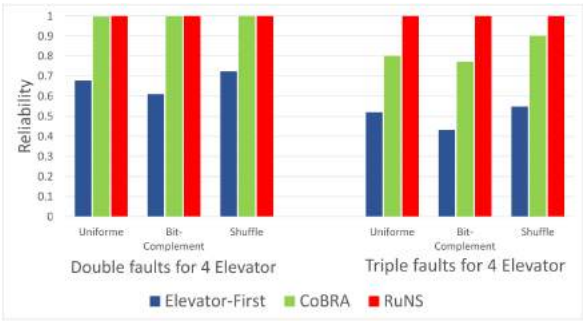

Figure 4: Reliability under double and triple faults for 4 TSVs

2. References

[1] A. Coelho, A. Charif, N.E. Zergainoh, R. Velazco, "A Runtime Fault-Tolerant Routing Scheme for Partially Connected 3D Networks-on-Chip," 2018 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), Chicago, IL, 2018, pp. 1-6

[2] A. Charif, A. Coelho, M. Ebrahimi, N. Bagherzadeh, N.E. Zergainoh, "First-Last: A Cost-Effective Adaptive Routing Solution for TSV-Based Three-Dimensional Networks-on-Chip," IEEE Transactions on Computers, pp. 1–1, 2018