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

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

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

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 defectivity, 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 projet 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

2. References

[1] A. Charif, N.E. Zergainoh, A. Coelho, M. Nicolaidis: Rout3D: A Lightweight Adaptive Routing Algorithm for Tolerating Faulty Vertical Links in 3D-NoCs, 22th IEEE European Test Symposium (ETS), pp. 1-6, Limassol, Cyprus, 22- 26 May 2017

[2] A. Charif, A. Coelho, N.E. Zergainoh, M. Nicolaidis: MINI-ESPADA: A Low-Cost Fully Adaptive Routing Mechanism for Networks-on-Chips, IEEE Latin-American Test Symposium (LATS), pp. 1-4, Bogota, Colombia, 13-15 March 2017

[3] T. Bonnoit, A. Coelho, N.E. Zergainoh, R. Velazco: SEU Impact in Processor's Control-Unit: Preliminary Results Obtained for LEON3 Soft-Core, 18th IEEE Latin American Test Symposium (LATS), pp. 1-4, Bogota, Colombia, 13-15 March 2017

[4] M. Solinas, A. Coelho, J. Fraire, N.E. Zergainoh, P. Ferreyra, R. Velazco: Preliminary Results of NETFI-2: An Automatic Method for Fault Injection on HDL-Based Designs, 18th IEEE Latin-American Test Symposium (LATS), Bogota, Colombia, 13-15 March 2017

[5] T. Bonnoit, N.E. Zergainoh, M. Nicolaidis, R. Velazco: Low Cost Rollback to Improve Fault-Tolerance in VLSI Circuits, IEEE International Symposium on Circuits and Systems (LASCAS), pp. 1-4, Bariloche, Argentina, 20-23 February 2017

[6] A. Charif, A. Coelho, N.E. Zergainoh, M. Nicolaidis: Detailed and highly parallelizable cycle-accurate network-on-chip simulation on GPGPU, ACM/IEEE Design Automation Conference (ASPDAC), pp. 672-677, Chiba/Tokyo, Japan, 16-19 January 2017

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, and enabling the next computation turn. But the aggressive technology scaling that paves the way to the ultimate CMOS nodes has dramatic impact to: process, voltage and temperature (PVT) variations; sensitivity to electromagnetic interferences (EMI), to atmospheric radiation (neutrons and protons) and to alpha particles; and circuit aging. It also imposes stringent power dissipation constraints. The resulting high defect levels, heterogeneous behavior of identical circuits, accelerated circuit degradation over time, and extreme complexity, affect adversely fabrication yield and/or prevent fabricating reliable chips in ultimate CMOS and post-CMOS technologies. These issues are becoming the main show-stoppers in the path leading to these technologies.

The Cells framework addresses the severe issues related to the design of massively parallel tera-device processors affected by high defect densities, in which we severe issues have to be addressed 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 massively affect one or more regular blocks (RAMs, FIFOs, buses) in a large fraction of nodes. Such faults may also very 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

Clearly, no existing solution can cope with such massively defective systems, which invalidate even massive redundancy schemes (e.g. duplication, TMR), as all replicated parts may be defective. Such schemes also induce high area and power penalties. Some approaches targeting the design of reliable single-chip massively parallel processors avoids massive redundancy by using self-tests (hardware implemented or software implemented to detect failures and create routing tables that are used subsequently to avoid failed processing nodes or failed routes. However, such approaches could not cope with the issues affecting ultimate CMOS and post CMOS technologies as:

- In highly defective technologies, the vast majority of nodes (processing elements and routers) may include one or another kind of faults (e.g. timing faults produced by process, voltage and temperature variations, EMI, or aging). Thus, declaring defective the nodes affected by any kind of faults will quickly waste the computational resources of the chip
- Achieving high fault coverage for timing faults is very difficult. Thus, many of these faults may escape fabrication test and also periodic self-tests and produce run-time errors
- Faults occurring during application execution can not be covered by self-tests

In this project we develop a comprehensive approach enabling using in efficient and reliable manner all parts able to perform useful computations.

The Cells framework (On-**C**hip S**el**f-hea**l**ing Tera-Device Processor**s**) [1-2] comprises several techniques spanning at all levels of the system: circuit, processor/architectural, array/routing, 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. Developments concerning some of these components are presented in the next sections.

References

Publications on the Overall Cells Framework

[1] M. Nicolaidis: Designing Single-Chip Massively Parallel Tera-Device Processors: Towards the Terminator Chip, Keynote–Plenary Session, IEEE VLSI Test Symposium, May 1-5, 2011, Dana Point, CA

[2] M. Nicolaidis: "Biologically Inspired Robust Tera-Device Processors", IEEE Design & Test of Computers, Volume 29, No 5, September/October 2012

[3] M. Nicolaidis, L. Anghel, N.E. Zergainoh, D. Avresky: "Designing Single-Chip Massively Parallel Processors Affected by Extreme Failure Rates", Proc. Design Automation and Test in Europe Conference (DATE), March 12 – 16, 2012, Dresden, Germany

Recent Publications on Cells' Components

[4] A. Charif, N.E. Zergainoh, M. Nicolaidis: "Addressing Transient Routing Errors in Fault- Tolerant Networks-on-Chips", Proc. 21th IEEE European Test Symposium, (ETS), May 24-37, 2016, Amsterdam, The Netherlands

[5] P. Papavramidou, M. Nicolaidis: **"**Test Algorithms for ECC-based Memory Repair in Ultimate CMOS and Post-CMOS", accepted for publication IEEE Transactions on Computers, to appear August 2016

[6] M. Nicolaidis, P. Papavramidou: "Memory Repair for High Defect Densities", Proc. 33rd IEEE VLSI Test Symposium (VTS), April 27-29, 2015, Napa Valley, California

[7] P. Papavramidou, M. Nicolaidis: "Low Power Memory Repair for High Defect Densities, P.Papavramidou, M.Nicolaidis", Proc. 21st IEEE International On-Line Testing Symposium (IOLTS), July 6-8, 2015, Halkidiki, Greece

[8] A. Charif, N.E. Zergainoh, M. Nicolaidis: "MUGEN: A High-Performance Fault-Tolerant Routing Algorithm for Unreliable Networks-on- Chip", Proc. 21st IEEE International On-Line Testing Symposium (IOLTS), July 6-8, 2015, Halkidiki, Greece

[9] M. Nicolaidis, P. Papavramidou: "Memory Repair for High Defect Densities", Proc. 33rd IEEE VLSI Test Symposium (VTS), April 27-29, 2015, Napa Valley, California

[10] P. Papavramidou, M. Nicolaidis: "Low Power Memory Repair for High Defect Densities, P.Papavramidou, M.Nicolaidis", Proc. 21st IEEE International On-Line Testing Symposium (IOLTS), July 6-8, 2015, Halkidiki, Greece

[11] M. Nicolaidis: "Double-Sampling Design Paradigm - a Compendium of Architectures", IEEE Transactions on Device and Materials Reliability, Vol. 15, No. 1, March 2015

[12] D. Avresky, F. Chaix, M. Nicolaidis: "Congestion-Aware Adaptive Routing in 2D-Mesh Multicores", Proc. 13th IEEE International Symposium on Network Computing and Applications (NCA), August 21-23, 2014, Cambridge, MA USA

[13] F. Chaix, N.E. Zergainoh, M. Nicolaidis: "A generic and high-level model of large unreliable NoCs for fault tolerance and performance analysis", Proc. 19th IEEE European Test Symposium, (ETS), May 26-30, 2014, Paterborn, Germany

[14] M. Nicolaidis : "Double-sampling Achitectures", Proc. IEEE International Reliability Physics Symposium (IRPS), June 1-5, 2014, Waikoloa, Hawaii, US

Early Publications on Cells' Components

[15] M. Nicolaidis, P. Papavramidou: "Transparent BIST for ECC-based memory repair" Proc. 19th IEEE International On-Line Testing Symposium (IOLTS), July 6-8, 2013, Chania, Greece pp. 216 - 223

[16] P. Papavramidou, M. Nicolaidis: "An iterative diagnosis approach for ECC-based memory repair", Proc. 31st IEEE VLSI Test Symposium (VTS), pp. 1-6, April 29 – May 2, 2013, Berkeley, CA, US.

[17] P. Papavramidou, M. Nicolaidis: "Reducing power dissipation in memory repair for high defect densities", 18th IEEE European Test Symposium (ETS), pp. 1-7, May 2013, Avignon, France

[18] F. Chaix, G. Bizot, M. Nicolaidis, N.E. Zergainoh: "Variability-Aware and Fault-tolerant Self-Adaptive applications for Many-Core chips", Proc. IEEE International On-Line Testing Symposium, Chania-Crete, Greece, July 8-10 2013

[19] G. Bizot, F. Chaix, M. Nicolaidis, N.E. Zergainoh: "Variability-Aware and Fault-tolerant Self-Adaptive Applications for Many-Core Chips", Proc. IEEE European Test Symposium, Avignon, France, Mai 27-31, 2013

[20] P. Papavramidou, M. Nicolaidis: "Test Algorithms for ECC-based Memory Repair", Proc. IEEE VLSI Test Symposium (VTS), April 23-26, 2012, Maui, Hawaii, US

[21] F. Chaix, D. Avresky, N.E. Zergainoh, M. Nicolaidis: "A Fault-Tolerant Deadlock-Free Adaptive Routing for On Chip Interconnects", Design Automation and Test in Europe Conference (DATE), March 14 – 18, 2011, Grenoble, France, (BEST IP AWARD)

[22] M. Nicolaidis, T. Bonnoit, N.E. Zergainoh: "Eliminating Speed Penalty in ECC Protected Memories", Proc. Design Automation and Test in Europe Conference (DATE), March 14 – 18, 2011, Grenoble, France

[23] G. Bizot, D. Avresky, F. Chaix, N.E. Zergainoh, M. Nicolaidis: "Self-Recovering Parallel Applications in Multi-Core Systems", 10th IEEE International Symposium on Network Computing and Applications (IEEE NCA11), 25 - 27 August 2011 Cambridge, MA USA

[24] F. Chaix, G. Bizot, M. Nicolaidis, N.E. Zergainoh: "Variability-aware Task Mapping Strategies for Many-core Processor Chips", IEEE International On-line Testing Symposium, July 2011, Athens, Greece

[25] F. Chaix, D. Avresky, N.E. Zergainoh, M. Nicolaidis: "Fault-Tolerant Deadlock-Free Adaptive Routing for Any Set of Link and Node Failures in Multi-Cores Systems", 9th IEEE International Symposium on Network Computing and Applications (NCA10), July 15-17 2010, Cambridge, MA

Resilient Integrated Systems

Keywords: Double sampling, GRAAL, single event transients, single bit upsets, soft errors

Members: T. Bonnoit, F. Bouesse, M. Nicolaidis, N.E. Zergainoh

Contracts: RESIST (RESilient Integrated SysTems) is a CATRENE project funded by the ministry of the sustainable development. Duration September 2014 – December 2017.

1. Context and goals

In massively parallel processor chips, a possible (and rather straightforward) approach may consist in exploiting the existence of large numbers of processing cores to implement fault tolerance, by using two processor cores to duplicated the execution of each task and comparing their outcome to detect errors (double modular redundancy-DMR), or by using three processor cores to triplicate the execution of each task and voting their outcome to correct errors (triple modular redundancy-TMR). These approaches have two major flaws: drastic decrease of processing power, and drastic increase of energy dissipation per task. Similarly, software implemented fault-tolerance drastically impacts performance and power, as it replicates the execution of each task. Hence, a breakthrough in fault-tolerant design is required for reducing drastically performance and energy dissipation penalties, and improving reliability. Such a breakthrough was achieved by a scheme (referred hereafter as double-sampling [1][2]). This scheme reduces drastically hardware and power costs, by avoiding both hardware replication and operations replication. Instead, to check the correctness of an output signal of a logic block; this scheme observes this signal at two different instants, by adding a redundant latch and driving it by means of a delayed clock signal (Figure 1).

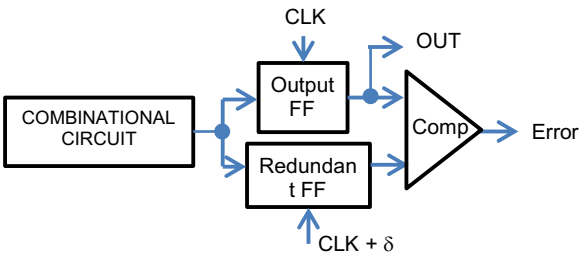

Figure 1: Double Sampling

While the double-sampling scheme reduces drastically the area and power cost with respect to traditional fault-tolerant schemes, some area and power penalties are still introduced due to the redundant latches. Furthermore, it suffers from two drawbacks:

- To avoid false alarms and miss-corrections, the path delays of the combinational circuits signals checked by the schemes in figure 1, must exceed the delay δ of the delayed clock. To ensure this constraint, buffers should be needed in the paths with delays shorter than δ, inducing no negligible area and power cost.
- Also, due to this issue, to avoid significant area and power penalties, we have to use moderate values for δ, limiting the duration of detectable faults.

We had developed highly efficient double-sampling architectures able to detect all errors induced by failure mechanisms affecting advanced technologies, as well as soft-errors induced by radiation in space applications but also in ground-level applications. Then, in the RESIST project, we implemented two robust processors by using two of these advanced double-sampling architectures.

In the first robust processor implementation we used the advanced double-sampling architecture based on GRAAL [3]. The scheme consist in transforming a flip-flop-based design, into its equivalent latch-based design, by moving the slave latches of the FFs to the middle of the combinational circuits, we obtain a design that can work at the same clock frequency as the original one. In addition, the master and slave latches operate at different phases. Then, the outputs of any combinational block (inputs of latches), are stable during the period in which its adjacent blocks are in computation. Thus, we can compare the outputs of the latches against their inputs to detect timing and transient faults of large duration, according to the GRAAL architecture that we have introduced few years ago. This scheme also detects SEUs as well as all kinds of latch faults (transition faults, retention faults). The outcome of this implementation on the icyflex processor from CSEM [4][5], in ST Microelectronics 40nm technology was a very high robustness that achieves detection of temporary faults of very large duration (up to 100 % of the maximum circuit delays); at very low area and power costs (less than 4 %). Thus, they can mitigate all yield, reliability, and life span issues affecting the advanced nanometric technologies, and can also be exploited for aggressively reducing the power dissipation, by aggressively reducing the supply voltage and using the fault mitigation capabilities of these architectures to cope with the delay faults induced by this voltage reduction.

In the second robust processor implementation, we succeeded to also cover all Single-Event Upsets (SEUs) and all Single-Event Transients (SETs) by employing the advanced double-

sampling architecture based on the one display in Fig.1. The main difference is that in this case $\delta < 0$, that is the redundant FF samples the signal before the main one [3]. Therefore, the parameter δ does not constrain anymore the minimum delay of the data paths of the circuit, which mean that it can be regulate to detect SETs as long as required, providing this way a robust architecture that can be used in space applications. This solution was implemented and evaluated on the Gaisler LEON3 processor [6] with the 28 FDSOI design kit of ST Microelectronics. The error detection scheme was implemented along with instruction replay to obtain a hardened integer unit version for the LEON3. The implementation of the double sampling + instruction replay increases the area by +140 %, which is 20 % lower than TMR. The respective contributions of the overhead are: 9 % for instruction replay; 53 % for the double sampling itself; 38 % for the timing constraints that ensure 100 % of SBU detection. However, as the design tools are not adapted to constraint the internal data-paths with a minimum delay, the results obtained can still be improved.

2. References

[1] M. Nicolaidis: “Time Redundancy Based Soft-Error Tolerant Circuits to Rescue Very Deep Submicron”, in Proc. 17th IEEE VLSI Test Symposium, April 1999

[2] L. Anghel, M. Nicolaidis: “Cost Reduction and Evaluation of a Temporary Faults Detecting Technique”, in Proc. IEEE DATE, March 2000
(BEST PAPER AWARD). Also, chapter in book “The Most Influential Papers of 10 Years DATE”, Springer, 2008

[3] M. Nicolaidis: “Double-Sampling Design Paradigm - a Compendium of Architectures” IEEE Transactions on Device and Materials Reliability, Vol. 15, No. 1, March 2

[4] Icyflex2 user manual, from CSEM

[5] J.L. Nagel, C. Arm, A. Corbaz, M. Morgan, V. Moser, P. Volet: “The icyflex2 Processor Architecture”, report from CSEM, p.31

Techniques and Methods to Tolerate Soft-Error in 3D Networks-on-Chip

Keywords: network-on-chip, router micro-architecture, single event transients, fault injection, and soft-errors

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

1. Context and goals

Network-on-Chip (NoC) architectures have emerged as a scalable solution to the bus-based interconnects challenges for future multi-core system-on-chip designs. Nevertheless, the conventional planar integration of integrated circuits has limited floor-planning choices with increasing number of processing elements attached, limiting the potential performance of 2D-NoC architectures. In order to extend 2D-NoC capabilities, multiple layers of active devices can be integrated in a three-dimensional NoC (3D-NoC) allowing to reduce interconnect lengths and thus, improving the overall NoC performance.

As the feature sizes of integrated circuits decrease aggressively, combinational logic becomes more susceptible to transient faults. Specifically, soft-errors provoked by the impact of energetic particles (heavy ions, protons, neutrons,…) or by cross-talk noise, were generally masked and thus disregarded in older technology. However, as operating voltages become lower and clocking frequency increases, the design of integrated circuits with technology node below 0.25um requires of particular attention to soft-error effects. Indeed, NoC routers are not the exception since they have a huge number of combinational logic elements. In fact, the probability of the occurrence of a soft-error is higher in 3D-NoC routers which are more complex than 2D-NoC routers because of the increasing number of connecting ports. Therefore, the study of 3D-NoC routers reliability towards soft-errors becomes mandatory for future large-scale integration of dependable system-on-chips.

One of the most critical sections in the NoC's control path is the Route Computation Unit (RCU), as it is the one responsible for selecting the next output port (i.e. direction) that a packet must take at every hop. If the RCU fails, packets may be forwarded to wrong outbound ports (i.e., misrouting) eventually leading to deadlocks (cyclic dependencies between packets) or packet loss.

In this project, some techniques and methods are proposed to enhance the reliability of 3D-NoCs by detecting and correcting errors provoked by transient faults in the RCU. The primary characteristic of our method is the ability to reliably and quickly detect misrouting based on a combination of a fault tolerance technique called *double-sampling*, and a complementary illegal turn detection method based on existing signals. These two methods are combined into a specific hardware unit called *fault detection circuit*, which can make decisions to recover from faults. The second characteristic is the capability to correct misrouting either by the reuse of the route computation samples provided by double-sampling or by directly rerouting the in-transit packet.

3D-NoC Architecture Overview

In this project was consider a partially vertically connected 3D-NoC mesh architecture, wherein the routers include, in addition to the usual five ports (East, West, South, North, Local), either an Up port, a Down port, or both (i.e., 5, 6 and 7 port configuration). The Up and Down ports of the router are connected vertically using Through-Silicon Via (TSV), as shown in Figure 1.

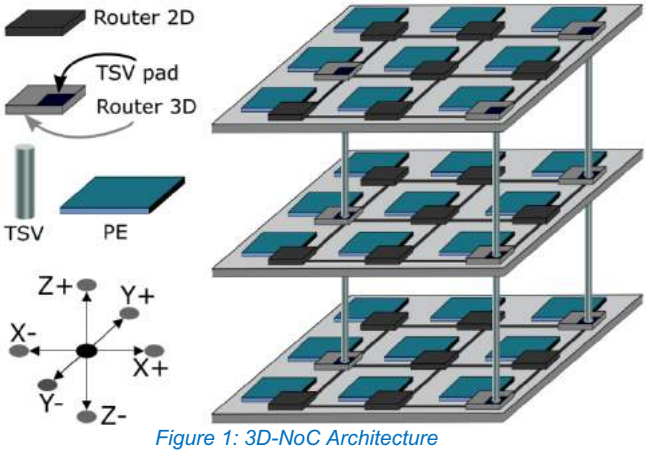

Figure 1: 3D-NoC Architecture

Since we consider RCU is responsible for selecting the output port and the Virtual Network (VN), a transient fault in the RCU can leave an in-transit packet with an incorrect route, implying a wrong output port direction, an erroneous VN or both. In this case, to detect such faults is required a mechanism capable of analyzing the output port and the VN selected by the RCU.

Techniques and Methods to Tolerate Soft-Error

Figure 2 illustrates the basic operation of the double sampling mechanism used in this project. We propose to use the rising edge of the clock as latching event of the regular flip-flop and the falling edge of the clock as latching event of the redundant sampling element to get the Output Port (OP) and Virtual Network (VN) computed by the RCU ([OP,VN]). The RCU is performed in one clock cycle, and the double-sampling is running in parallel with RCU, thus resulting in no additional delay.

The Sample_1 and Sample_2 are obtained in different instants, as shown in Figure 2. Those two samples are compared by the additional Fault Detection Circuit (FDC) after route computation. If the result of comparison between these two

samples is different, then a fault has affected one of the samples. Therefore, both samples need to be analyzed with the purpose of determining which one is faulty. On the other hand, if the result is equal, either there was not a fault, or a same fault affecting both samples. In this case, we only select Sample 1 to validate if there have been errors or not.

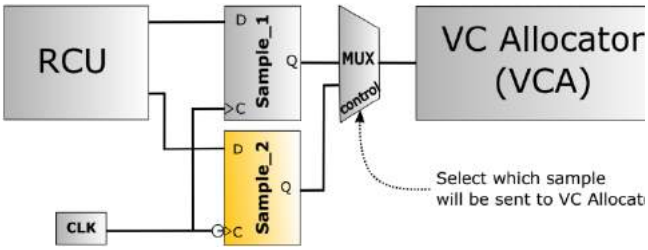

Figure 2: A Double-Sampling implementation

The detection of an incorrect route vector ([OP, VN]) can be performed directly by the VC Allocator (VCA) or by an extra Fault Detection Circuit (FDC), which runs in parallel with the VCA, as shown in Figure 3. Because the VCA includes a table indicating the possible input-output connections, it can detect illegal output port and VN requests with minor modifications.

For faults that do not involve illegal turns, more sophisticated approaches are needed. This is the case when the output port and VN are both valid, but the packet needs to make illegal turns at later hops to reach its destination. To deal with these faults that cannot be detected by the VCA, an extra FDC is included in the NoC router, as shown in Figure 3. To this end, the FDC checks three conditions. First, it compares the Zd and Zc to know if the packets which need to go in the Up direction were wrongly selected to go Down, or vice versa. Additionally, it checks if the packets are directed to a local port, in this case, the FDC compares the position of the current router with the destination of the packets. Second, the FDC checks if the direction adopted by the packet leaves to an illegal turn in the upcoming packet path. Third, it checks the two samples provided by the double-sampling and selects which of them must be finally used by the VCA.

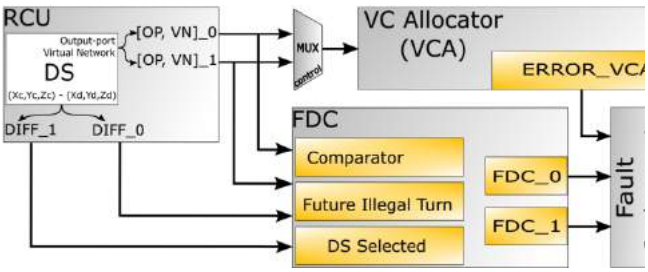

Figure 3: Fault detection Circuit – FDC

The VCA and the FDC work together to detect transient faults in the RCU. The faults reported by the FDC and VCA are analyzed and judged as follows: If both samples (Samples_1 and Sample_2 from double-sampling) are different, this method will check both samples trying to find which one is correct to be used, and the packet will continue its path using this result. If booth samples are faulty the packet will be rerouted.

Repeating the route computation process is not simple since one RCU is shared among all VCs in the same input port. This means that each port is able to do one route computation per clock cycle. To avoid replicating the RCU for each VC, we propose a simple rerouting mechanism that uses the existing RCU. The idea is to ensure that a limited number of new packets enter an input port in which one or more packets have requested rerouting. The solution consists in preventing packets from entirely leaving the input port, ensuring the buffers do not get available to receive new packets until all packets have been rerouted.

Evaluation and Validation

The performance of the proposed RCU architecture is estimated through the average packet latency under three different traffic patterns (Uniform, Bit-complement). The resulting performance comparison is shown in Figure 4 when there are faults in the ROUT3D-FDR and no faults in the ROUT3D baseline. The degradation in performance is noticed because to recover from misrouting the VC, the same input port needs to be blocked to avoid the reception of new headers while executing the rerouting routine. In other hand, using our techniques the reliability of the 3D-NoC increase in the presence of faults, which means that more packets can arrive his destination. Additionally, the area and power overheads were estimated using an operating frequency of 1GHz, a power supply of 1V, and a commercial ST FD-SOI 28nm library. The measured area and power overheads including double-sampling, rerouting, and fault detection circuit were 4.1 % and 6.8 % respectively.

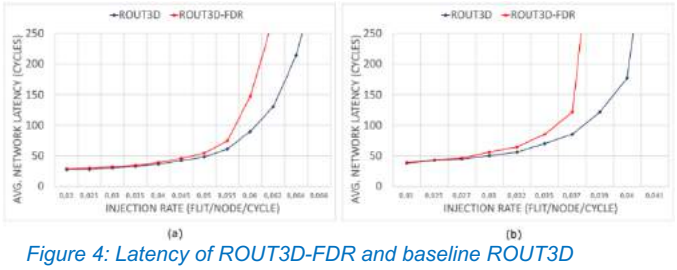

Figure 4: Latency of ROUT3D-FDR and baseline ROUT3D under (a) uniform traffic and (b) bit complement traffic

2. References

[1] A. Charif, N.E. Zergainoh, A. Coelho, M. Nicolaidis: "Rout3d: A lightweight adaptive routing algorithm for tolerating faulty vertical links in 3d-nocs," in 22th IEEE European Test Symposium (ETS'17). ACM IEEE, pp. 1–6, 2017

[2] A. Charif, N.E. Zergainoh, M. Nicolaidis: "Addressing transient routing errors in fault-tolerant networks-on-chips," in 2016 21th IEEE European Test Symposium (ETS), May 2016, pp. 1–6

[3] L. Huang, X. Zhang, M. Ebrahimi, G. Li: "Tolerating transient illegal turn faults in nocs," Microprocess. Microsyst., vol. 43, no. C, pp. 104–115, Jun. 2016

[4] X. Zhang, M. Ebrahimi, L. Huang, G. Li, "Fault-resilient routing unit in nocs," in 2015 28th IEEE System-on-Chip Conference (SOCC), pp. 164–169, Sept 2015

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, R. Velazco, G. Hubert, H. Mecha, 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 memory cells. The soft-error rate (SER) determination is still a challenge to evaluate the technology sensitivity and to extrapolate the trends for future IC's generations.

A common procedure in electronics consists in reducing the power supply value to the lowest possible value that can retain information. Although this value depends on the technology miniaturization, it means that the power supply falls from 3.3 V down to, e.g., 0.6 V. This is done in order to achieve two purposes: First of all, a decrease in the power consumption and, secondly, to minimize electrical stress on the device and prolong its effective life time.

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 the GHADIR research group in UCM (*Universidad Complutense de Madrid)* since 2014. Other organisms and companies involved in this research line are ONERA (*Office National d'Etudes* et *de Recherches Aérospatiales*), in charge of simulations and predictions, and Cypress Semiconductor, who has provided proprietary information about the DUTs (Devices Under Test) layout. Experiments usually consist in a tunable power supply far enough from the irradiation chamber and an accurate measuring system. Cypress 180 nm and 65 nm SRAMs as well as Renesas' 130 nm to 110 nm SRAMs have been tested at the GENEPI-2 facility.

GENEPI-2 (*GEnérateur de NEutrons Pulsé Intense*) is a particle accelerator located at *Laboratoire de Physique Subatomique et de Cosmologie* (LPSC, Grenoble - France) dedicated to high-energy 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 beam intensity is continuously measured on target and the 14 MeV neutron production is monitored on line by a silicon detector.

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

Since neutrons propagate in the whole experimental room, the DUT test platform (hardware platform providing the suitable environment to the DUT and the readout electronics) also sits in the neutron flux. Therefore a shield of borated polyethylene was assembled and used to protect it (see Figure 2).

Figure 2: The DUT facing the neutron source

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 recent facility upgrade.

2. Recent outcomes

As previously stated, SRAMs from Cypress Semiconductor and Renesas were tested. Cypress Semiconductors' devices were different versions of the 2Mx8bit CY62167.

- CY62167D: 130 nm
- CY62167E: 90 nm
- CY62167G: 65 nm

Besides, the 110-nm RMLV1616A memory, by Renesas, was tested with an identical set-up.

Results obtained for the Cypress Semiconductor Memories

Experimental results clearly indicate that the soft-error cross sections increase as the power supply value decreases [1-2]. Thus, single bit upsets are about 10 times more likely at ~0.5 V than at the nominal value, 3.3 V. This behavior is also observed in multiple bit upsets with multiplicities ranging from 2 to 5. Events with larger multiplicities are very common for the 90-nm SRAM, but much less likely in the other ones (see Figure 3). Also, as shown in Figure 3, the soft-error cross sections show a decrease as the technology shrinks. In general, devices at 65 nm are more radiation tolerant than equivalent ones at 130 nm.

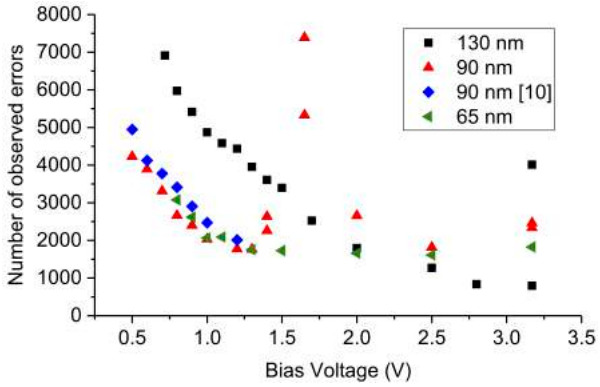

Figure 3: Number of errors observed for different voltages

Bitflips were classified in different kinds of errors using proprietary information provided by Cypress Semiconductor. Another important detail is that the 65-nm device has Error Correction Codes (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 allow estimating the expected number of errors due to radiation in several natural environments such as sea level or high-altitude flights.

Experiments also showed that events such as micro Latch-ups and SEFIs (Single Event Function Interruptions) may occur in SRAMs at intermediate voltages (1.1-1.7 V) yielding strange clusters of errors on the order of 25 bitflips and MBUs (see Figure 4). Strangely, this phenomenon is not observed neither at nominal voltages (~3.3 V) nor at ultra low values (~0.5 V). These phenomena were not observed in previous experiments probably due to the much lower neutron flux before the GENEPI-2 upgrade.

Another interesting outcome of the experiments is that the large number of events has allowed the refinement of the strategies developed by the research team to classify errors into single or multiple bit upsets. These strategies were based on the search of statistical anomalies in the set of affected addresses [3-4].

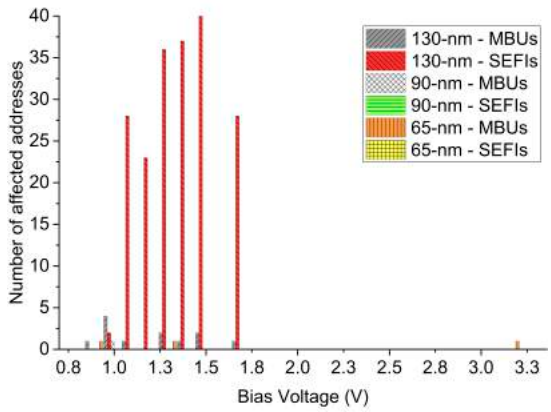

Figure 4: Number of memory addresses affected by MBUs and SEFIs.

Results obtained for the 110-nm A-LPSRAM

An old version of this 2Mx8bit soft-error-free memory built in 150-nm technology was tested in previous campaigns. Recently, a more advanced model was released by Renesas in 110-nm technology so different samples were purchased and tested with nominal and low bias voltages (from 0.5V to 3.3V).

Experimental results, still in study, seem to show that at nominal voltage the memories are extremely tolerant to 14-MeV neutrons, backing up the manufacturer's claim of being soft-error free at typical environments [5-6].

However, strange clusters of errors have been observed in the memories as they were in the previous model in 150-nm CMOS technology. There is a slight difference in the signature of these errors, since in 150-nm devices usually occurred at ultra low voltage values, 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.

3. References

[1] J.A. Clemente et al., "Sensitivity Characterization of a COTS 90-nm SRAM at Ultralow Bias Voltage," in IEEE Transactions on Nuclear Science, vol. 64, no. 8, pp. 2188-2195, Aug. 2017

[2] 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, in press

[3] J.A. Clemente et al., "Statistical Anomalies of Bitflips in SRAMs to Discriminate SBUs From MCUs," in IEEE Transactions on Nuclear Science, vol. 63, no. 4, pp. 2087-2094, Aug. 2016

[4] F.J. Franco et al., "Statistical Deviations From the Theoretical Only-SBU Model to Estimate MCU Rates in SRAMs," in IEEE TNS, vol. 64, no. 8, pp. 2152-2160, Aug. 2017

[5] R. Velazco et al., "Evidence of the Robustness of a COTS Soft-Error Free SRAM to Neutron Radiation," in IEEE Transactions on Nuclear Science, vol. 61, no. 6, pp. 3103-3108, December 2014

[6] J.A. Clemente et al., "Single Events in a COTS Soft-Error Free SRAM at Low Bias Voltage Induced by 15-MeV Neutrons," in IEEE Trans. on Nuclear Science, vol. 63, no. 4, pp. 2072-2079, Aug. 2016

Delay/Disruption Tolerant Networks Performance and Reliability

Keywords: Delay/disruption tolerant networks, satellite constellations, nanosatellites

Members: J.A. Fraire, A. Charif, M. Solinas, N.E. Zergainoh, R. Velazco

Cooperations: UNC, CONAE, JPL, TUD

1. Context and goals

Delay/Disruption Tolerant Networks (DTNs) have recently been considered as an alternative to extend Internet boundaries into space. In particular, recent studies have considered their applicability in Low-Earth Orbit (LEO) satellite constellations. To overcome link disruptions, DTN nodes temporarily store and carry in-transit data until a suitable next-hop link becomes available. To overcome delays, end-to-end feedback messages are no longer assumed continuous or instantaneous. This distinctive characteristic allows DTN to operate in environments where communications can be challenged by latency, bandwidth, data integrity, and stability issues.

In spite of the recent advances in the area, the analysis of the fault tolerance of existing DTN solutions remains an open research topic. Indeed, existing DTN solutions had assumed that information on the communication scheduling is highly accurate, disregarding transient and temporary faults, inaccuracies in the orbital prediction, among other unwanted effects. Indeed, studying reliability of DTN is mandatory before seriously considering its applicability in the harsh space environment where radiation effects, vibrations, collisions, out-gassing among others pose significant challenges for satellites.

2. Recent outcomes

In this context, the RIS team has provided an extensive fault injection analysis based on appealing and realistic case studies of delay-tolerant satellite constellations. The state-of-the-art version of Contact Graph Routing (CGR) algorithm was evaluated. The CGR algorithm was implemented in collaboration with the National University of Córdoba (UNC) in DtnSim, a new simulator. DtnSim is also introduced in this document and is already available via open-source licensing[1]. Results obtained from DtnSim in this analysis constitutes the state of the art research in the reliability aspect of DTNs and were published in a journal paper [1]. The main modeling challenges addressed in DtnSim simulator were published in a specific conference [2].

Fault model

Over the last years, the RIS team as well as the semiconductor industry have been particularly concerned by the effects of radiation on integrated circuits and embedded systems in general. The rationale behind this motivation lies not only in the use of these systems in harsh radiation environments but also in the increasing degree of integration of devices embedded in the same chip. Among the possible errors, transient errors occur in the system temporarily and are usually caused by radiation interference, also known as Single Event Effects (SEEs). The random occurrence in time and space of such failure phenomenon, the probability of the error to happen, and the probability of effectively detecting an unwanted behavior, can be modeled by means of an exponential (Poisson) distribution. The exponential distribution model was chosen for this study and can be described by means of the following equation as detailed in [1].

$$F(t) = \int_0^t \lambda e^{-\lambda\tau} = 1 - e^{-\lambda\tau}$$

Adopted exponential distribution model

By means of the configuring lambda with Mean Time to Failure (MTTF) and Mean Time to Repair (MTTR), a fault injection system was designed to study the resulting traffic flow of a DTN under failure conditions. Other effects such as permanent errors as well as orbital predictions inaccuracies are left as future studies.

Simulator Platform

There exist several tools to evaluate DTNs overviewed in [1] and [2], but a new alternative had to be developed. Thus, a simulator called DtnSim was implemented, with support from the National University of Córdoba (UNC). Each node in DtnSim is based on the layered architecture illustrated in Figure 1. This architecture is a simplification of the original DTN architecture that has been adapted for DTN-based satellite systems and is comprised of an application (APP), network (NET), and Medium Access Control (MAC) layer. Physical layer effects such as bit error rate can be modeled within the MAC layer if required.

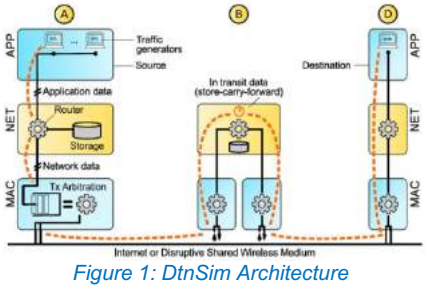

Figure 1: DtnSim Architecture

[1] DtnSim public repository: https://bitbucket.org/lcd-unc-ar/dtnsim

The APP layer is the element that generates and consumes user data. The NET layer is the element in charge of providing delay-tolerant multi-hop transmission (i.e., routes) and was probably the most challenging module of DtnSim to develop. The CGR routing algorithm was implemented as an exchangeable sub-module of this layer. The MAC layer of the DtnSim node is designed to provide a reliable wireless link and to multiplex the shared medium among nodes with wireless interfaces. Also, DtnSim was extended with a fault-injector module based on the exponential model previously described. The complexity and the approaches taken to develop DtnSim were described in a specific conference paper [2].

Simulation Results

In order to assess CGR behavior under transient failures in realistic delay-tolerant satellite constellations, two configurations were proposed: a sun-synchronous along-track and a Walker-delta formation. The former constellation was designed in collaboration with the Argentinian Space Agency (CONAE) and proposed in two related publications [3] [4]. The Walker formation was developed in collaboration with Technical University of Dresden (TUD), deriving in another publication [5]. However, these only analyzed error-free performance of the network. The relevant reliability analysis and comparison was published in [1] and is described below.

Both along-track and Walker constellations are based on 16 cross-linked LEO satellites (max. link range of 1000 Km at 500 km height), 25 ground target points, and 6 ground stations. Systems Tool Kit (STK) software was used to propagate these parameters for an analysis period of 24 hs. An intuitive illustration of the node's location on a world map is provided in Figure 2. The left side of the picture plots the ground tracks of the Walker formation while the along-track is on the right.

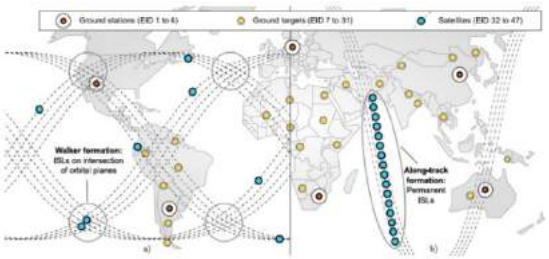

Figure 2: Evaluated Satellite Constellations

Simulation results cannot be plotted here for space reasons, the interested reader is referred to [1] to analyze the curves. In general, the along-track formation exhibited a higher effective fail rate which is consistent with the higher connectivity of a permanently connected constellation. The total bundles received curve on the Walker constellation evidenced a complete traffic delivery. However, the along-track formation is unable to deliver such load. After a thorough analysis of the simulation traces, it was found that a pathological routing behavior impeded CGR to find all feasible routes. These miscalculations are magnified as the failure rate increases. This was further discussed in [1] and derived in some suggestions to improve the existing CGR algorithm. At the time of writing, these modifications are being implemented in collaboration with the Jet Propulsion Laboratory (JPL) in a new release of ION, a DTN software stack, to be released in 2018. The along-track system resulted significantly more stable with the variation of the MTTF. In [1] is provided a detailed analysis of the results.

Future perspectives

A recent cooperation with GomSpace has started to develop reliable planning algorithms.

Future work includes the exploration of further CGR enhancements that could improve its robustness when implemented in fault-prone DTN systems. However, these improvements not only depend on the algorithm but also on the planning stage, which happens on ground.

Indeed, a more integral analysis of reliable communication scheduling and routing is currently on-going in the context of the GOM-X4 mission, with satellites manufactured by GomSpace, a company partially funded by ESA. Two satellites, illustrated in Figure 4, were successfully launched on February 2018 from China, and are currently at early-orbit provisioning stage. They fly in an along-track formation and use inter-satellite links to communicate among them. Both satellites are 6U CubeSats with several payloads which require of particular attention as the power subsystem is not able to provide continuous energy. The communication links are not the exception, thus the utilization of store and forward solutions such as DTN becomes mandatory.

Figure 4: GomSpace CubeSats: GOM-X4A and GOM-X4B

The Saarland University is leading this cooperation with a strong support from J. Fraire, a member of the group that developed the analysis presented in this report. A similar cooperation is on the waiting queue. Indeed, CONAE is still planning to deploy a massive network of small satellites to accomplish their long-term objective of the “segmented architecture” project, being the SARE mission the most relevant technological demonstrator.

3. References

[1] J. Fraire, P. Madoery, S. Burleigh, M. Feldmann, J. Finochietto, A. Charif, N.E. Zergainoh, R. Velazco, “Assessing Contact Graph Routing Reliability in Distributed Satellite Constellations”', Hindawi Journal of Computer Networks and Communications, July 2017

[2] J. Fraire, P. Madoery, F. Raverta, J. Finochietto, R. Velazco, “DtnSim: Bridging the Gap Between Simulation and Implementation of Space-Terrestrial DTNs", 6th Int. Conf. on Space Mission Challenges for Information Tech. (SMC-IT), Sept 2017, Spain

[3] J. Fraire, P. Madoery, J. Finochietto, P. Ferreyra, R. Velazco, “Internetworking Approaches Towards Along-Track Segmented Satellite Architectures”, 2016 IEEE Int. Conf. on Wireless for Space and Extreme Environments (WiSEE), Sept. 2016, Aachen, Germany

[4] C. Barrientos, A. Ferral, L. Cara, J. Fraire, R. Velazco, et al, “A Segmented Architecture Approach to Provide a Continuous Glaciers Monitoring System Based on DTN Comms and Cubesat Platforms”, 1st IAA Latam Symposium on Small Sats, Feb. 2017, Argentina

[5] J. Fraire, M. Feldmann, S. Burleigh, “Benefits and Challenges of Cross-Linked Ring Road Satellite Networks: A Case Study”, 2017 IEEE Int. Conf. on Comms (ICC), May 2017, Paris, France

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, 20 – 23 February 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

[11] 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

[12] 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

[13] 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

Software approach to improve the reliability of parallel applications implemented on multi-core and many-core processors

Keywords: Reliability, Parallel Architectures, Multi-core and many-core programming, Redundancy, Partitioning, Fault injection

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

The widespread use of COTS multi-core and many-core processors in computing systems is a consequence of their flexibility, low-power consumption, intrinsic redundancy and high performance. However, there are some dependability issues that have to be overcome. Indeed, the reliability, availability and security of systems based on these devices are reduced due to their increasing complexity, integration scale and sensitivity to natural radiation. In this context, this work provides valuable information about the Single Event Effects (SEE) sensitivity of applications running on multi/many-core processors and proposes a software fault-tolerance technique to improve the reliability of systems implemented on these devices.

The evaluation of the sensitivity was done through software fault-injection campaigns and radiation ground testing. For fault-injection, the principles of the Code Emulating Upset (CEU) approach were used and different variations of the CEU fault-injection method were developed. These variations include a fault-injector based on *fork* principle, a fault-injector on bareboard based on inter-processor interrupts, and a fault-injector for distributed systems based on NoC and inter-processor interrupts [1].

2. Evaluation of the impact of SEEs on parallel applications running on multi-core and many-core processors

The methodology used for evaluation was based on multiple-case studies and their analysis. The different scenarios implemented consider a wide range of system configurations in terms of multiprocessing mode, programming model, memory model, and resources used. For the experimentation, two Commercial Off-The-Shelf (COTS) devices were selected: the Freescale Power PC P2041 quad-core built in 45nm SOI technology, and the KALRAY MPPA-256 many-core processor built in 28nm CMOS technology. The case-studies were evaluated through fault-injection and neutron radiation testing.

The obtained results serve as a useful guideline to developers for choosing the most reliable system configuration according to their requirements. For instance, the evaluation done on the Freescale P2041 multi-core under neutron radiation has proved that despite the Silicon-on-Insulator manufacturing technology and the protection mechanisms implemented on the memory, erroneous results occurred. Additionally, results confirm the sensitivity dependence on software environment. It has been demonstrated that the Asymmetric Multiprocessing Mode is more reliable than the Symmetric.

Multiprocessing mode. Also, one can verify that the intrinsic application characteristics play an important role on the system sensitivity. In addition, it can be observed a considerable vulnerability of the operating systems (OS) especially concerning system crashes and exceptions. This vulnerability has put in evidence that the implementation of redundancy at user level is not enough to overcome dependability issues. Therefore, it is clear the imperative necessity of using the partitioning concept to guarantee both reliability and availability of the system [2-4].

On the other hand, the evaluation under neutron radiation of the implemented scenarios on the KALRAY MPPA-256 many-core processor has proved the effectiveness of the protection mechanisms implemented on its memories. Furthermore, it was also possible to determine the pertinence of enabling the cache memories of this device for improving its performance with minimum reliability consequences. These results also corroborate the SEE sensitivity dependency on software environment. Moreover, through the implementation of parallel applications in the many-core processor, it was possible to confirm that the migration of an application to a distributed system communicating by Networks-on-Chip is a non-trivial task [5-6].

Furthermore, one of the mainstays for migrating from uni-processors to multi/many-core processors in criticality systems is the possibility to take advantage of the inherent redundancy to implement fault-tolerance by replication. As expected, the implementation of the Triple Modular Redundancy on the quad-core has improved significantly the system reliability. However, it has reduced its parallel computing

capability. Nevertheless, redundancy applied to many-core processor is a promising solution.

3. A software approach to improve the system reliability

The approach developed during this work, called called NMR-MPar, is based on N Modular Redundancy and partitioning concepts has been evaluated through the implementation of a case study running on the KALRAY MPPA-256 many-core processor. The approach was applied to two benchmarks, the *cpu-bound* Travel Salesman Problem (TSP) application and the *memory-bound* Matrix Multiplication. Figure 1 illustrates the approach that implements a Quad Modular Redundancy (4-MR) on a clusterized architecture. The system is comprised by two independent Double Modular Redundancy (DMR) configured in a *master-slave* scheme. The *RM1* and *RM2* of the *IO cluster* act as *master* of each module where the Compute Clusters' (CC) cores perform as *slaves*. The first DMR is comprised by 2 multi-core partitions (*P0* and *P1*) to run the distributed application, and 2 mono-core partitions, *P4* which is the master of the DMR and *P5* which is the voter. The second DMR uses 2 multi-core partitions (*P2* and *P3*) to run the application and 2 mono-core partitions, *P6* (*master*) and *P7* (*voter*). Each DMR module sends the results to the other DMR, majority voting to 4 results.before each voter applies.

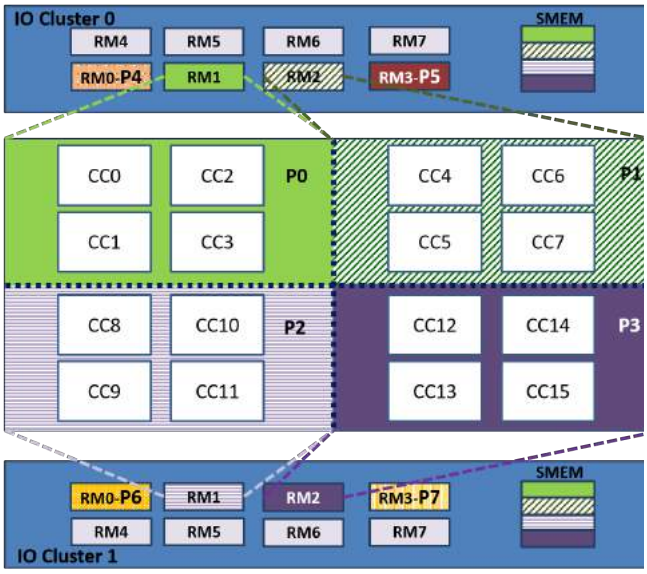

Figure 1: Implementation of NMR-MPar on MPPA

The approach developed during this work, called NMR-MPar, is based on N-Modular Redundancy and partitioning concepts and has been evaluated through the implementation of a case study running on the KALRAY MPPA-256 many-core processor. The approach was applied to two benchmarks, the *CPU-bound* Travel Salesman Problem (TSP) application and the *memory-bound* Matrix Multiplication. The evaluation was done through fault-injection campaigns. Results summarized in figure 2 have proved that the proposed approach is very efficient to improve the reliability on applications running on many-core COTS processors.

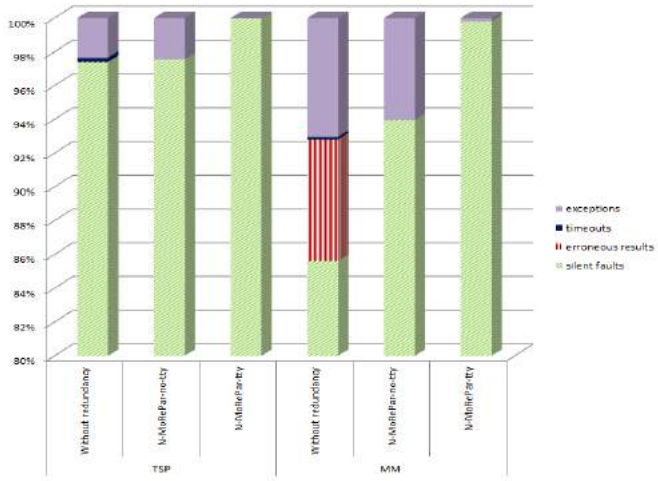

Figure 2: Results of fault-injection campaigns

The obtained results show that only exceptions cannot be completely overcome by this approach. This issue can be surpassed by implementing complementary fault tolerance techniques on application and OS level. Therefore, by combining NMR-MPar with these techniques, it will be possible to use multi-core and many-core COTS processors in harsh radiation environments

4. Future work

This work opens up various possibilities to continue the investigation, such as:

- Evaluate the NMR-MPar approach under neutron and heavy-ions radiation to validate its applicability to avionics and spacecraft applications
- Conjugate this approach with other software implemented fault tolerance techniques at OS and user level to deal with exceptions
- Evaluate NMR-MPar on other system configurations. Benefiting of the large number of cores in the selected many-core processor, it is suitable to evaluate other systems configurations. Also, it can be interesting to implement this approach on other platforms to evaluate its broad applicability on systems based on multi/many-core processors
- Appraise its overhead in terms of energy and power consumption

5. References

[1] V. Vargas, P. Ramos, J.F. Méhaut, R. Velazco: “Swifi Fault injector for heterogeneous many-core processors” Pontificia Universidad Católica del Ecuador, issn: 2528 - 8156 (in press), vol. 106, 2018

[2] 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

[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", In Proc. of Radiation and its Effects on Components and Systems (RADECS), pp. 135-138, September 2015

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

[5] 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, nb1, pp. 483-490, Jan. 2017

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