

POLITECNICO DI TORINO
Repository ISTITUZIONALE

EuFRATE: European FPGA Radiation-hardened Architecture for Telecommunications

Original

EuFRATE: European FPGA Radiation-hardened Architecture for Telecommunications / Bozzoli, Ludovica; Catanese, Antonio; Fazzoletto, Emilio; Scarpa, Eugenio; Goehringer, Diana; Pertuz, Sergio; Kalms, Lester; Wulf, Cornelia; Charaf, Najdet; Sterpone, Luca; Azimi, Sarah; Rizzieri, Daniele; La Greca, Salvatore Gabriele; Merodio Cordinachs, David; King, Stephen. - ELETTRONICO. - (2023), pp. 1-6. (Intervento presentato al convegno Design, Automation & Test in Europe Conference & Exhibition (DATE) - DATE 2023 tenutosi a Antwerp (Belgium) nel 17-19 Aprile 2023)
[10.23919/DATE56975.2023.10137035].

Availability:

This version is available at: 11583/2974542 since: 2023-06-12T08:52:14Z

Publisher:

IEEE

Published

DOI:10.23919/DATE56975.2023.10137035

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

EuFRATE: European FPGA Radiation-hardened Architecture for Telecommunications

Ludovica Bozzoli*, Antonino Catanese*, Emilio Fazzoletto*, Eugenio Scarpa*
Diana Goehringer[†], Sergio A. Pertuz[†], Lester Kalms[†], Cornelia Wulf[†], Najdet Charaf[†]
Luca Sterpone[‡], Sarah Azimi[‡], Daniele Rizzieri[‡], Salvatore Gabriele La Greca[‡]
David Merodio Codinachs[§], Stephen King[§]

*Argotec Group

[†]Technische Universität Dresden

[‡]Politecnico di Torino

[§]European Space Agency

Abstract—The EuFRATE project aims to research, develop and test radiation-hardening methods for telecommunication payloads deployed for Geostationary-Earth Orbit (GEO) using Commercial-Off-The-Shelf Field Programmable Gate Arrays (FPGAs). This project is conducted by Argotec Group (Italy) with the collaboration of two partners: Politecnico di Torino (Italy) and Technische Universität Dresden (Germany). The idea of the project focuses on high-performance telecommunication algorithms and the design and implementation strategies for connecting an FPGA device into a robust and efficient cluster of multi-FPGA systems. The radiation-hardening techniques currently under development are addressing both device and cluster levels, with redundant datapaths on multiple devices, comparing the results and isolating fatal errors. This paper introduces the current state of the project’s hardware design description, the composition of the FPGA cluster node, the proposed cluster topology, and the radiation hardening techniques. Intermediate stage experimental results of the FPGA communication layer performance and fault detection techniques are presented. Finally, a wide summary of the project’s impact on the scientific community is provided.¹

Index Terms—FPGA cluster, parallel processing, radiation hardening, FPGA-FPGA communication

I. INTRODUCTION

Field Programmable Gate Array (FPGA) devices play a key role in the modern architecture of satellite On-Board Computers (OBC) and Data Handlers (DH). Their intrinsic flexibility allows them to implement important functionalities for system communication interfaces, Digital Signal Processing (DSP), and data storage. This growing appreciation is mainly due to their reprogrammable capabilities and hardware-accelerating features. These devices support on-field reconfiguration i.e., a reprogramming of the implemented electronic circuit, either in the case of different tasks to perform or when nominal work parameters change. Furthermore, when multiple FPGAs are used in a unique computing architecture, they are capable to increase the computational capacity of satellite payloads by leveraging distributed computing and performances. However, FPGAs like any other Electrical, Electronic, and Electro-mechanical (EEE) components are prone to radiation-induced

errors. Considering the adoption of SRAM-based FPGAs in mission-critical applications, the main problem is their sensitivity to Single Event Effects (SEEs) and in particular Single Event Upsets (SEUs) and Multiple Event Upsets (MBUs).

To solve this problem, radiation-hardened (rad-hard) FPGAs have been proposed which are costly for new space areas including CubeSats and nanosatellites. Therefore, utilization of Commercial-Off-The-Shelf (COTS) FPGA devices is becoming more common even though no hardening techniques are applied to their silicon. Therefore, suitable mitigation techniques are required to ensure the system’s reliability, operations’ continuity, and consistency of processed data. Moreover, multiple FPGAs or FPGA clusters are becoming highly popular due to the distribution of functionalities among different FPGAs. Therefore, radiation hardening techniques can be implemented at both the board level and cluster level.

The main scientific contribution of the EuFRATE project is the novel realization of a multi COTS FPGA devices computing architecture for telecommunications applications with resilient and radiation-hardened architecture able to outperform digital telecommunication payload algorithms of at least one order of magnitude.

The paper is representing the intermediate stage of the EuFRATE project, aiming to provide an overview on the technical status of the project in terms of achievements, challenges addressed, and open bits until the end of the project. The paper is organized as follows. Section II describes the state of the art of the high performance computing architecture based on FPGAs and adopted in the space environment. Section III presents the contributions of partners to the project. Section IV describes the proposed the novel FPGA cluster architecture, the idea behind the fault tolerance mechanism and the application use cases. Section V presents the current status of the project, while Section VI presents the preliminary results achieved. Finally, Section VII presents the summary and the future activities of the project.

II. STATE OF THE ART

The rising computational capacity required in various application fields can be achieved by connecting multiple FPGAs to a large system [1], [2]. One challenge of such systems

¹The EuFRATE project is financially supported by the European Space Agency (ESA) in response to the call ESA AO/1-10240/20/UK/ND within the ARTES 4.0 CORE competitiveness generic programme line Component A: Advanced Technology, Activity Reference 5C.416

is the scalability, as multiple compute devices need to be distributed over a large infrastructure [3]. An increasingly common solution is the use of clusters consisting of efficient compute devices interconnected by high performance networks [4]. Nevertheless, large FPGA systems, such as the CHIPit from Synopsys, can be very costly. To reduce these costs, a cluster of COTS FPGA boards [5] can be used. Examples such as the Zedwulf [6] and the ZCluster [2] demonstrate the growing need for such low-cost FPGA clusters. Both approaches combine several low-cost commercial FPGA boards, such as the ZedBoard containing an ARM-FPGA-based SoC. Using an FPGA cluster does not only reduce the costs, but it also increases the availability of the system due to the enhanced number of available computing FPGAs. Another advantage is expandability, as new FPGAs can be added without having to replace the entire system when running out of resources. This improves both the longevity and robustness versus permanent faults such as Total Ionization Dose (TID) and fault tolerance of the system due to the removal of the system single point of failure.

Even though FPGAs and in particular SRAM-based FPGAs have been widely used in aerospace applications the large volume of the configuration memory of SRAM cells is extremely sensitive to radiation-induced effects in particular Single Event Upsets (SEUs) and Multiple Event Upsets (MEUs). In order to perform an accurate radiation analysis, many radiation test experiments have been performed calculating the FPGA sensitivity cross-section and application error rate [7], [8]. However, performing a radiation test is costly. Therefore, several methodologies have been proposed such as fault injection and simulation environments to evaluate the sensitivity of the SRAM-based FPGAs to soft errors [9]. To mitigate the effect of SEUs and MEUs, several techniques have been applied to the design implemented on SRAM-based FPGAs [10]. Among them, Triple Modular Redundancy (TMR) is the traditional method even though it is costly in terms of area, speed and power consumption, which is not negligible especially considering mission critical applications. Therefore, several partial or selective TMR techniques have been proposed [11]. Another mitigating method is scrubbing that refers to rewriting the content of the configuration memory in order to correct the SEUs and MEUs in the memory content. If the system is not implemented in such a way that a previous state can be recovered, then the device will be put in a reset state and completely restarted after rewriting the configuration memory. This leads to an operational downtime which could happen to be unsustainable in a mission-critical scenario such as space missions. Nevertheless, as the main advantage, this method permits to detect and correct faults avoiding their dangerous accumulation. A further technique is Dynamic Function eXchange (DFX), which allows to reprogram certain parts of the FPGA configuration memory during operation. In this project, we apply a combination of the mentioned techniques, such as TMR, scrubbing and DFX, in order to exploit the advantages of each technique both at the node and the cluster level.

III. DESCRIPTION OF THE PARTNERS

The realization of the project is carried out in a partnership among Argotec group (Italy), Politecnico di Torino (Italy) and Technische Universität Dresden (Germany). The project has been funded by the European Space Agency (ESA) in the framework of the ARTES 4.0 CORE competitiveness programme line.

A. Argotec group

Argotec is the leader in the development of small satellites for Deep Space missions, whose Avionics department has got a productive background in the exploitation of FPGA devices in their missions. From the recent achievements of Argotec, LICIA CubeSat can be mentioned as part of the NASA DART mission to image the asteroid after the incident.

B. Politecnico di Torino (PoliTO)

PoliTO is a public university in Italy. The “Aerospace and Reconfigurable Computing” lab of the Department of Control and Computer Engineering is involved in radiation effects analysis and mitigation of FPGAs since 2006. In the EuFRATE project, PoliTO is responsible for proposing and implementing the radiation hardening techniques at both single board and cluster level.

C. Technical University of Dresden (TUD)

TUD is a public university in Germany. The chair for “Adaptive Dynamic Systems” of the Faculty of Computer Science is specialized in the development of FPGA-based systems, including single and multiple devices. In the EuFRATE project, TUD is responsible for implementing the FPGA cluster architecture.

IV. EUFRATE PRINCIPAL CONCEPTS

FPGA-cluster computation is rapidly expanding and High-Performance Computing (HPC) centres in the EU are expecting to reach exascale performance within 5 years. The EuFRATE architecture will advance state-of-the-art techniques by implementing a SEE-resilient partial reconfiguration performed inside of each FPGA device, permitting savings of resources with respect to an external scrubbing solution and increasing the overall system availability.

A. Architecture

The EuFRATE conceptual architecture is shown in Fig. 1. All the devices shall communicate to each other by means of two links: a high-throughput protocol, such as the Xilinx Aurora 64b/66b able to support up to 25.8Gbit/s and a robust and single-error immune rad tolerant beacon bus. Aurora shall be employed for high bandwidth data transfer with minimum latency. Additional elements are the configuration port that interfaces with an external multi-boot Non-Volatile Memory (NVM), such as a NOR-based Flash memory, containing all the bitstreams of each device composing the cluster node, and a reconfiguration arbiter. The reconfiguration arbiter will be implemented with self-repair and reconfiguration capabilities.

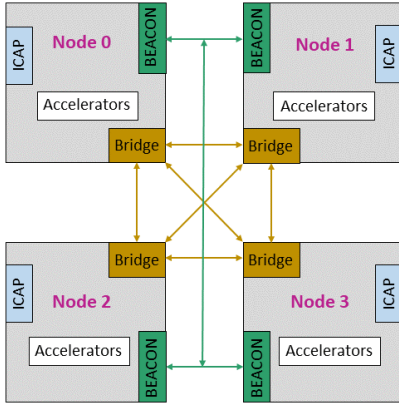


Fig. 1. The main architecture of EuFRATE cluster prototype based on-line fault-tolerant and self-repair reconfigurable architecture with four nodes.

B. Fault Tolerance

The fault tolerance have crucial importance in the EuFRATE which has been tackled with three different approaches: The detection of an error, its subsequent correction, and the proper orchestration of the cluster during these steps which leads to a complete recovery of the system. Furthermore, we exploited the EuFRATE cluster structure of computational units by adding an inter-node communication architecture able to support fault detection and correction. Each node will be SEU resilient and it will implement place and route mitigation techniques. Furthermore, it will benefit from the system redundancy. In case of failure of a node detected during the nominal activity, the spare node is activated to substitute the failing node. The failing node is powered-off and restarted as a new spare element for the next failure occurrence.

C. Use Cases

FPGA-based clusters are widely used in many different application fields, such as the Internet of Things, digital signal processing and prototyping, asymmetric encryption algorithms [12], JPEG encoder [13] and digital wireless channel emulator (DWCE) [14]. In satellite systems for communication, efficient multi-channel support is essential. Therefore, as a testbed for the EuFRATE project, one of its main computing blocks, LDPC (Low Density Parity Check) decoding carried out in a single instruction multiple data fashion, was chosen.

V. DEVELOPMENT STATUS OF THE PROJECT

The EuFRATE architecture is currently at an intermediate development stage. The hardware architecture of the cluster on the FPGA level is composed of:

- Blazes and Blaze On-Node (BaByLoN) Processing System, consisting of the general-purpose domain of the node including: A hardened microprocessor (TOWER: Triple OverWatch kERnel), running the main reliability and cluster management services; and a plain microprocessor (GarDeN: General and Real-time DomaiN) to support the application specific computations and main node data-flow management.




Fig. 2. Block diagram of the NA and its connectivity with the router and the internal network and to create messages.

- TMR Beacon Controller, consisting of a dedicated monitoring IP core able to periodically monitor the health status of the node and the tile. It is composed of two sub-controllers: the local controller, in charge of monitoring the node; the global controller, in charge of monitoring the status received from the other nodes.
- Router + Network Adapter (NA), consisting of dedicated IP cores in charge of managing the exchange and flow of data across nodes and within the node itself, according to the requests of the BaByLoN processing system.
- DDR4 memory controller enables the access to the on-chip external memory from the GarDeN and the NA.
- Aurora interface for the data link. At least three of them are required to form a tile with a full-mesh topology, while four should be present for the nodes that interface with the outside and to form a network of four tiles.

A. Architectural Description

The internal interconnection network of a node is connected to a router via a NA to enable the communication with other nodes inside or outside of a tile. The NA is a crucial component of the EuFRATE system. Fig. 2 describes the internal modules of the NA and its links to the internal interconnection network, composed by the router and the GarDeN MicroBlaze processor. To enable the exchange of messages between nodes, an additional communication layer for the application is required. The NA contains an RX unit to receive packages, a TX unit to create and transmit packages, a DMA unit to connect to the internal network, and a controller to control the DMA and add the application layer.

In detail, the DMA component is envisioned to enable the stream of data from the router (i.e., from outside the node) to inner components (i.e., node memory) through the NA. In Fig. 2, the red lines denote AXI4-stream interfaces (lowest resource consumption), the purple lines AXI4-full interfaces (highest resource consumption) and the dashed blue lines mark AXI4-lite interfaces.

On each node there is an internal connection to connect the internal components with each other. The router is connected to all three aurora cores via AXI4-stream protocol, which is a much lighter and resource efficient protocol. Thus, the




Fig. 3. General scheme of radiation-induced fault detection and mitigation applied to EuFRATE.

routing information is controlled via the NAs by the processing domain in charge of managing the dataflow. The router can have a fourth port, which can be used if more than one tile exists, creating a network of tiles.

The most critical cluster management services are executed on the TOWER processor. The TOWER collects all information about the state of the cluster that it receives from the Beacon Health Controller and the GaRDeN software. It answers to telecommands from the spacecraft and provides information about the current state of the system. While the TOWER software is executed bare-metal on the TMR processor, the operating system FreeRTOS is exploited on the GaRDeN processor. Several software tasks share the processor and enable to monitor the execution of hardware tasks on hardware accelerators and concurrently communicate with the TOWER processor on the own FPGA respectively with the GaRDeN processors on the other nodes. The GaRDeN software is also aware of the state of execution of hardware tasks on neighboring nodes. In case of a failure, the GaRDeN software checks the availability of hardware accelerators on the own FPGA as well as on neighboring nodes. It redirects the data that failed to be processed and regulates its execution with the goal to reach a balanced load distribution within the FPGA cluster.

B. Error Detection Architecture

The fault tolerance strategies adopted to EuFRATE covers several different parts of the architecture and act accordingly on each element while focusing on limiting the performance and resource overheads. These strategies consist of both detection and mitigation of radiation-induced faults, implemented at node and cluster level. Figure 3 represents the general scheme of the developed fault detection and mitigation strategies. In order to perform the detection of radiation-induced errors, an ad-hoc watchdog unit named Beacon Health Controller (BHC) is developed. The BHC is composed of two sub-systems, the local controller dedicating to internal monitoring of each node and Global Controller for external monitoring of the node through the other nodes of the cluster. In order to perform the detection of the SEU occurrence in the single node, a hardware-oriented approach has been designed and implemented which consists of a robust ad-hoc watchdog unit,

Beacon Health Controller, monitoring the correct functionality of several modules inside the single FPGA node and detect the occurrence of SEU in each module. Please notice that we applied the TMR Technique to the BHC module in order to design a radiation-robust watch dog monitoring unit.

The TMR BHC is divided into two different parts: the local and global controllers. The local controller is responsible for the local monitoring of selected modules of the node which consist of Babylon, GaRDeN and TOWER, as well as hardware accelerators, which are expected to send their health signal, a single toggling bit, to the local health controller. As soon as the local health controller does not receive the expected health signal, it triggers the detection of an error.

The global controller is in charge of monitoring the health status of other nodes inside the cluster as well as sending the health signal to the other nodes based on the local node condition. To make a robust communication between nodes inside the cluster, a mitigation-oriented bus has been designed, called Robust Bus, composed of discrete point-to-point connections between the cluster nodes. The Global Controller is responsible for monitoring the Robust Bus, periodically checking the health status signals from the other nodes to detect if any node is having functional issues or is not producing any legal health status at all.

C. Error Mitigation Architecture

After the detection through the beacon health controller unit, the correction of the faults is performed applying two different techniques, scrubbing and DFX.

The first technique, scrubbing, is exploiting the Xilinx Soft Error Mitigation (SEM) IP core which performs SEU correction for configuration memory. The SEM IP is in charge of constantly performing readback of the configuration memory of the node, detecting single and multiple upsets in the configuration memory by checking several ECC codes. Once the SEU or MBU is detected, the SEM IP will correct the error in the configuration memory. This technique has many advantages such as low detection and correction latency, however, it comes with some drawbacks such as SEM-uncorrectable errors which cannot be corrected directly by using the SEM IP core.

Therefore, the second complementary correction technique has been developed, exploiting the Xilinx Dynamic Function eXchange (DFX) provided by Xilinx. DFX is the ability to dynamically modify the identified and selected blocks of logic by downloading partial bit files while the remaining logic continues to operate without interruption. While originally the DFX technique was introduced to change the functionality of a module on the fly without the need to fully reconfigure the whole FPGA, we have exploited this technique to recover from the faults identified by the BHC Unit, overwriting the bits configuring the target module. Applying the combined BHC unit and DFX technique, the requirements regarding detection and correction at the node-level are fulfilled.

TABLE I
AREA RESOURCE USAGE COMPARISON

Resource type	GaRDn [#]	DRPM [#]
LUTs	98,239	74,408
DFFs	123,475	93,772
LUT RAM	8,677	4,559
BRAM	134	42




Fig. 4. A picture of the actual development stage of the EuFRATE FPGA cluster with 4 nodes.

VI. PRELIMINARY EXPERIMENTAL RESULTS

The EuFRATE architecture is expected to open new horizons in resilient reconfigurable computing in space and enable high performance telecommunication protocols in space electronic systems. The EuFRATE prototype is illustrated in Fig. 4. The Xilinx Kintex Ultrascale XCKU040 considered as FPGA node FPGA. The actual implementation on this device achieves an improvement of reconfiguration time of around 28% with respect to state-of-the-art dynamically reconfigurable processing systems (DRPM) [15]. The resources usage of the EuFRATE architecture has been compared with the DRPM system implemented on a Virtex-IV SRAM-based FPGA and the data are reported in Table I.

A. Error Tolerance Preliminary Data

We perform the preliminary radiation-based reliability analysis of the EuFRATE system, calculating the SEU rate. We defined a possible mission profile adopting the EuFRATE testbaed, as reported in Table II. Several analyses have been made to elaborate the characteristics of the deployment environment. In particular, Solar Energetic Particles (SEP) and Galactic Cosmic Rays (GCS) is considered as the main threats while the ontribution of trapped particles is considered negligible.

In order to perform a preliminary analysis for identifying the SEU rate, we exploited the CREME69 simulation tool con-

TABLE II
EXPECTED PROFILE OF THE EUFRATE MISSION

Mission Duration	15 years
Spacecraft Shielding	10 mm (Aluminium Equivalent)
Launch	Direct in GEO
Spacecraft Location	Geosynchronous Orbit




Fig. 5. Flux vs.energy of particles and integral flux vs LET of particles for different solar conditions.

sidering the features such as critical charge and sensitive volumes dimension of the device under the study. A preliminary analysis considering both SEP and GCR for different solar conditions (Solar Minimum, Solar Maximum, Worst Week, Worst Day, Peak 5 Minutes Average) has been performed while the result regarding the solar Minimum and Maximum are represented in Figure 5, which shows the expected flux of particles for the described mission profile. Exploiting the CREME96 simulation tool, with respect to the characteristic of the device under test such as Weibull Distribution parameters, the SEE Rate for Maximum and Minimum cosmic ray has been calculated. Table III reports the calculated SEE rate for the device as well as bits per day and per second. Please notice that these results are calculated with respect to the physical features of the device and is not considering the features of the implemented design. To consider the SEE rate with respect to the implemented EuFRATE design, we calculate the SEE rate per bits considering the essential bits of the EuFRATE design. Essential bits are defined as the bits associated with the circuitry of the design and are a subset of the device configuration bits. If an essential bit is upset, it changes the design circuitry but the upset might not affect the function of the design. Considering the EUFRATE node design and referring to information extracted from the Xilinx FPGA design tool chain, 25% (31,970,357 out of 128,038,080) of the configuration memory bits are considered as essential bits.

To assess the impact on the fault tolerance of the EuFRATE system, a first fault injection campaign has been performed

TABLE III
SEE RATE ESTIMATION FOR DIFFERENT SOLAR CONDITIONS

SEE Rate	Solar Minimum	Solar Maximum
SEE Rate device per day	1.37E+00	8.43E-01
SEE Rate device per second	1.59E-05	9.76E-06
SEE Rate bit per day	1.07E-08	6.58E-09
SEE Rate bit per second	1.24E-13	7.62E-14
SEE Rate of Essential Bits per Day	3.43E-01	2.10E-01

TABLE IV
FAULT INJECTION RESULTS COMPARISON

	GaRDeN [%]	DRPM [%]
Success (Fault Masked)	79.15	10.05
Silent Data Corruption (SDC)	1.82	43.52
Exception	1.97	19.86
Undetected System Halt	0.59	26.57
Detected System Halt	16.47	n.a.

specifically targeting the GaRDeN processing system. The fault injection campaign consisted of emulating SEUs happening in the configuration memory of the FPGA device. Please note that the majority of the injected faults would have been properly detected and corrected by the SEM-IP, so the performed fault injection refers to the faults as if they would have been categorized as uncorrectable by the SEM core. To verify the impact of the injected faults, a benchmark application has been executed on the GaRDeN, and the functional output has been compared with a golden reference. Furthermore, ad-hoc exception handling functions have been defined in the FreeRTOS operating system framework, where the application has been executed, to investigate which fraction of the total faults resulted in a processor exception, and which would have been detected and properly handled by the software layer itself. The main purpose of this fault injection campaign has been the preliminary verification of the TMR Beacon Controller detection capabilities, assuming that for every detection by the controller, a corresponding correction action would have followed, leading to the recovery of the system. Table IV reports the results of the fault injection campaigns comparing the fault injection performed on the GaRDeN subsystem on the Xilinx Ultrascale and static region of the DRPM implemented on the Xilinx Virtex-IV devices. As it is possible to notice, even if the EuFRATE architecture is larger than the DRPM system, the sensitivity to injected faults it is drastically lower, since the number of SDC is 23 times lower. Furthermore, the EuFRATE system is able to detect the majority of the critical system errors (Detected System Halt) and to achieve a success rate almost of 80%. As can be seen, the vast majority of the injected faults resulted in masked faults while the second most probable outcome has been the system halt condition. In such a state, the processor stops every operation and enters a freeze state without completing the application execution and failing in producing coherent output data. Most of the halt conditions (95.5%) have been properly detected by the TMR BHC, reducing the failure conditions down to 0.59% of the totally injected faults.

VII. SUMMARY

In this paper, we presented the actual development status of the EuFRATE project, aiming at developing a radiation-hardened telecommunication payload for aerospace electronic systems on FPGAs. The reconfigurable architecture and the radiation-hardening techniques are addressing both device and cluster level and they are currently under completion. However, preliminary fault injection result demonstrated that the proposed EuFRATE systems is able to achieve enhanced performances and a drastic reduction of the critical system faults due to radiation effects.

REFERENCES

- [1] O. Knodel, A. Georgi, P. Lehmann, W. E. Nagel, and R. G. Spallek, "Integration of a highly scalable, multi-fpga-based hardware accelerator in common cluster infrastructures," in *2013 42nd International Conference on Parallel Processing*, 2013, pp. 893–900.
- [2] Z. Lin and P. Chow, "Zcluster: A zynq-based hadoop cluster," in *2013 International Conference on Field-Programmable Technology (FPT)*. IEEE, 2013, pp. 450–453.
- [3] A. M. Caulfield, E. S. Chung, A. Putnam, H. Angepat, J. Fowers, M. Haselman, S. Heil, M. Humphrey, P. Kaur, J.-Y. Kim *et al.*, "A cloud-scale acceleration architecture," in *2016 49th Annual IEEE/ACM international symposium on microarchitecture (MICRO)*. IEEE, 2016, pp. 1–13.
- [4] M. Hernández, A. A. Del Barrio, and G. Botella, "An ultra low-cost cluster based on low-end fpgas," in *Proceedings of the 50th Computer Simulation Conference*, 2018, pp. 1–12.
- [5] N. Mentens, J. Vandorpe, J. Vliegen, A. Braeken, B. D. Silva, A. Touhafi, A. Kern, S. Knappmann, J. Rettkowski, M. S. A. Kadi *et al.*, "Dynamia: Dynamic hardware reconfiguration in industrial applications," in *International Symposium on Applied Reconfigurable Computing*. Springer, 2015, pp. 513–518.
- [6] P. Moorthy and N. Kapre, "Zedwulf: Power-performance tradeoffs of a 32-node zynq soc cluster," in *2015 IEEE 23rd Annual International Symposium on Field-Programmable Custom Computing Machines*. IEEE, 2015, pp. 68–75.
- [7] B. Du, L. Sterpone, S. Azimi *et al.*, "Ultrahigh energy heavy ion test beam on xilinx kintex-7 sram-based fpga," in *IEEE Transactions on Nuclear Science*. IEEE, 2019, pp. 1813–1819.
- [8] S. Azimi, C. D. Sio, A. Portaluri, D. Rizzieri, and L. Sterpone, "A comparative radiation analysis of reconfigurable memory technologies: Finfet versus bulk cmos," in *Elsevier Microelectronics Reliability*. Elsevier, 2022, pp. 1–6.
- [9] L. Sterpone, S. Azimi, and B. Du, "A 3-d simulation-based approach to analyze heavy ions-induced set on digital circuits," in *IEEE Transactions on Nuclear Science*. IEEE, 2020, pp. 2034–2041.
- [10] S. Azimi and L. Sterpone, "Digital design techniques for dependable high performance computing," in *IEEE International Test Conference*. IEEE, 2020.
- [11] L. Sterpone, S. Azimi, and B. Du, "A selective mapper for the mitigation of sets on rad-hard rtg4 flash-based fpgas," in *16th European Conference on Radiation and Its Effects on Components and Systems*. IEEE, 2016.
- [12] X. Bai, L. Jiang, Q. Dai, J. Yang, and J. Tan, "Acceleration of RSA processes based on hybrid ARM-FPGA cluster," in *2017 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, Jul. 2017. [Online]. Available: <https://doi.org/10.1109/iscc.2017.8024607>
- [13] K. Takano, T. Oda, R. Ozaki, A. Uejima, and M. Kohata, "Implementation of distributed processing using a PC-FPGA hybrid system," in *2019 International Conference on Field-Programmable Technology (ICFPT)*. IEEE, Dec. 2019. [Online]. Available: <https://doi.org/10.1109/icfpt47387.2019.00074>
- [14] S. Buscemi and R. Sass, "Design and utilization of an FPGA cluster to implement a digital wireless channel emulator," in *22nd International Conference on Field Programmable Logic and Applications (FPL)*. IEEE, Aug. 2012. [Online]. Available: <https://doi.org/10.1109/fpl.2012.6339253>
- [15] L. Sterpone, M. Porrman, and J. Hagemeyer, "A novel fault tolerant and runtime reconfigurable platform for satellite payload processing," in *IEEE Transactions on Computers*, vol. 62, no. 8, 2013, pp. 1508–1525.