POSTER: A Novel Robust Core for Detecting Node Failures in FPGA Clusters

Giorgio Cora, Corrado De Sio, Sarah Azimi, Luca Sterpone
Department of Control and Computer Engineering (DAUIN)
Politecnico di Torino, Turin, Italy
giorgio.cora@polito.it, corrado.desio@polito.it, sarah.azimi@polito.it, luca.sterpone@polito.it

ABSTRACT

Field Programmable Gate Arrays (FPGAs) are gaining popularity in different fields, including space applications, where high computational capabilities are required; for this reason, FPGAs are often used as nodes in clusters. When considering mission-critical systems, reliability must be ensured, even in radiation environments such as space. Thus, it is necessary to define a way of monitoring the entire system, ensuring the correct behavior of each node. This work introduces the Beacon Controller, a module to be implemented on the FPGA elements of a cluster for real-time monitoring of the computational elements of the node.

CCS CONCEPTS

• Hardware • Robustness • Safety critical systems

KEYWORDS

FPGA clusters, reliability, radiation hardening

ACM Reference format:

Giorgio Cora, Corrado De Sio, Sarah Azimi and Luca Sterpone. 2024. POSTER: A Novel Robust Core for Detecting Node Failures in FPGA Clusters. In 21th ACM International Conference on Computing Frontiers (CF'24), May 7-9, 2024, Ischia, Italy. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/1234567890

1 Introduction

With the increasing demand for high computational capabilities, efficient and simple results can be achieved by connecting multiple Field Programmable Gate Arrays (FPGAs) to form a larger system forming FPGA Cluster [1]-[2]. In fact, FPGA, renowned for their flexibility and high performance, enables the creation of efficient systems at low-cost, particularly when leveraging Commercial Off-the-Self (COTS) devices.

These kinds of systems, thanks to their properties, are often adopted into safety-critical environments, including space ones.

When operating in these fields, however, achieving good performances is not the only requirement since reliability also must be ensured. FPGAs are prone to errors that can lead to modifications in the implemented design; common examples are Single Event Upset (SEU) and Multiple Bit Upset (MBU), caused by charged energy particles that hit the configuration memory of these devices and modify its content. Moreover, when multiple FPGAs are connected to form a cluster, the failure of a single node, if not correctly treated, can lead to disastrous consequences, possibly causing the entire system to fail.

While many studies have been conducted on the reliability of SRAM-based FPGAs [3], radiation hardened devices [4], and cluster nodes [5], no suitable tools have been proposed for monitoring the health of nodes within an FPGA-based cluster.

In this work, we present the Beacon Controller, a custom IP which can monitor all the main processing unit within the module in which is inserted, as well as the other modules in the cluster, where other beacon controllers are instantiated. Outside of the status and health signals, it is provided with interfaces that allow it to trigger full or partial reconfiguration of the system, if the cluster supports this feature.

2 The Beacon Controller Structure

The Beacon Controller is provided with both hardware and software features, such that it can be managed in an easier way. Fig. 1a represents how the beacon is introduced in an FPGA cluster.

2.1 Hardware Structure

In Fig. 1b the internal structure of the Beacon Controller is highlighted. The device has been developed adopting the TMR technique, in such a way to design a robust component with a reduced probability of failure. A watchdog timer has been implemented to validate the status of the resources under control. They are requested to send periodic signals to the Beacon Controller to indicate their operativity.

Having the capability of monitoring both the resources within the node as well as the other FPGAs of the cluster, the

CF '24, May 7-9, 2024, Ischia, Italy.

ACM ISBN 978-1-4503-0000-0/18/06...$15.00
https://doi.org/10.1145/3649153.3653001

BC is internally subdivided into 2 units: the Local Controller and the Global Controller.

The Local Controller is responsible for analyzing the status of the local resources within the node and if the watchdog timer expires, reconfiguration of the failing module is triggered, while the other nodes of the cluster are notified through the Global Controller. It can manage up to 30 local hardware resources simultaneously, providing a balance between the number of monitored elements within the node and the complexity and area occupation of the device itself.

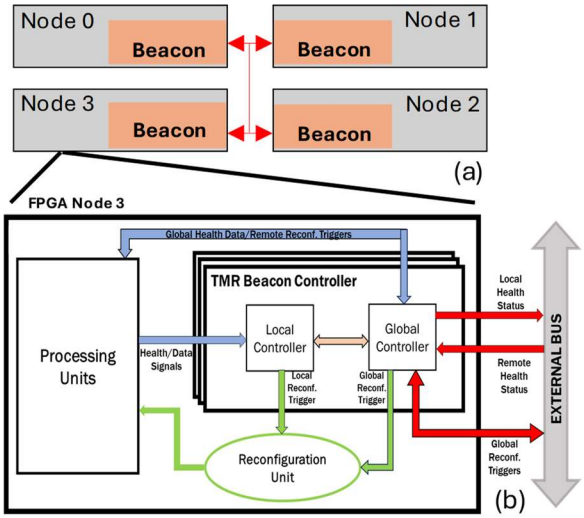

Figure 1: (a) FPGA Cluster Implementing the Beacon Controller (b) Beacon Controller Internal Structure

The Global Controller oversees the status of the entire system: it manages all the information regarding the health of each node and, eventually, requests a remote reconfiguration. It can manage up to 3 external nodes, considering the smallest possible cluster size being of 4 elements, but can be easily parametrized such that more external resources can be monitored.

2.2 Software Drivers

The Beacon Controller can be managed by accessing its internal registers through specific software drivers.

Before being operative, the Local Controller must be started and the timeout value for the watchdog must be set. This can be achieved by setting up the Control and Data Registers respectively. To prevent unwanted writing operations, these registers are normally locked and should be unlocked only during the initialization phase.

The information related to the health status of the monitored modules is stored into specific data structures, which are updated constantly.

Three additional drivers are available to retrieve information from the Global Controller, keep the node updated on the status of the other devices and, eventually, trigger the reconfiguration of a remote node.

3 Experimental Results

Table I. FPGA Resources Utilization of Target modules

FPGA Resource	Resource Utilization
LUTs	3185 (2.35%)
FF	3282 (2.04%)
BRAM	0 (0%)
Muxes	32 (0.64%)

In this section, the results from a fault injection campaign, targeting SEUs in configuration memory, are analyzed. The Beacon Controller has been implemented on a Xilinx Kintex Ultrascale KCU105 and the injections were performed on a single board using the SEM-IP tool provided by Vivado.

In Table I, the resource utilization of the TMR beacon Controller with respect to the available resources of the FPGA is reported.

To verify the effectiveness of the Beacon Controller, more than 800000 injections were performed over the FPGA area where it was implemented. The proposed custom IP has been implemented next to a Xilinx Microblaze processor and the results of the fault injection campaign have been reported in Table II. The processing unit, hardened through TMR techniques, initializes the Beacon Controller, and then enters its nominal execution phase. The outcome of the injection process was classified according to the following categories:

- *Good*: Beacon operations executed correctly.
- *Halt*: Beacon stopped working.
- *TMR Masked*: failure in one of the replicas.
- *Unmasked:* wrong values in Beacon Registers.

Table II. Beacon Controller Intervention Rate

Error Type	Good	Halt	TMR Masked	Unmasked
Error Rate	788814 (96.00%)	10638 (1.29%)	13215 (1.61%)	9028 (1.10%)

The Beacon Controller achieved extremely good performances and low failures rate, granting high stability while monitoring the required resources, and, at the same time, keeping a low resource utilization.

4 Conclusion

This paper presents a custom IP that can be used to monitor the status of computational resources within every single node in a FPGA cluster to improve system reliability.

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] T. Li et al., "Investigation into SEU Effects and Hardening Strategies in SRAM Based FPGA," in 2017 17th European Conference on Radiation and Its Effects on Components and Systems (RADECS), Oct. 2017, pp. 1–5. doi: 10.1109/RADECS.2017.8696177.

[4] E. Vacca, S. Azimi, L. Sterpone, "Failure rate analysis of radiation tolerant design techniques on SRAM-based FPGA", Microelectronics Reliability, Volume 138, 2022, 114778, ISSN 0026-2714, DOI: 10.1016/j.microrel.2022.114778.

[5] L. Sterpone, M. Porrmann, 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.