Journal of Systems Architecture 61 (2015) 615–627

Contents lists available at ScienceDirect

Journal of Systems Architecture

journal homepage: www.elsevier.com/locate/sysarc

Techniques for on-demand structural redundancy for massively parallel processor arrays

Vahid Lari [a,*], Jürgen Teich [a], Alexandru Tanase [a], Michael Witterauf [a], Faramarz Khosravi [a], Brett H. Meyer [b]

[a] Department of Computer Science, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Germany
[b] Electrical and Computer Engineering Department, McGill University, Canada

ARTICLE INFO

Article history:
Received 8 October 2014
Revised 28 July 2015
Accepted 15 October 2015
Available online 30 October 2015

Keywords:
Fault tolerance
Coarse-grained reconfigurable Architectures
Reliability analysis

ABSTRACT

In this paper, we present techniques for providing on-demand structural redundancy for Coarse-Grained Reconfigurable Array (CGRAs) and a calculus for determining the gains of reliability when applying these replication techniques from the perspective of safety-critical parallel loop program applications. Here, for protecting massively parallel loop computations against errors like soft errors, well-known replication schemes such as Dual Modular Redundancy (DMR) and Triple Modular Redundancy (TMR) must be applied to each single Processor Element (PE) rather than one based on application requirements for reliability and Soft Error Rates (SERs). Moreover, different voting options and signal replication schemes are investigated. It will be shown that hardware voting may be accomplished at negligible hardware cost, i. e. less than two percent area overhead per PE, for a class of reconfigurable processor arrays called Tightly Coupled Processor Arrays (TCPAs). As a major contribution of this paper, a formal analysis of the reliability achievable by each combination of replication and voting scheme for parallel loop executions on CGRAs in dependence of a given SER and application timing characteristics (schedule) is elaborated. Using this analysis, error detection latencies may be computed and proper decisions which replication scheme to choose at runtime to guarantee a maximal probability of failure on-demand can be derived. Finally, fault-simulation results are provided and compared with the formal analysis of reliability.

1. Introduction

Due to continuous technology scaling, not only supercomputers but also embedded computer systems have emerged integrating multiple processors and hardware accelerators on a single silicon chip, also known as Multi-Processor System-on-Chip (MPSoC). The number of such computing units is about to exceed 1000 processing elements on a single chip [1]. Integrating such an amount of computational resources into a chip raises challenges regarding the power density and power consumption. Therefore, there has been a significant trend toward employing heterogeneity by using *application-specific* hardware components such as CGRAs. One sample target architecture among many CGRAs are the so called Tightly Coupled Processor Arrays (TCPAs) [2,3]. A TCPA is a perfect choice as an accelerator in an MPSoC to speed up computationally expensive loop programs with regular computations stemming from, e. g., digital media and signal processing. A simplified drawing of a 2D TCPA with 24 PEs is depicted in Fig. 1.

But employing such high density designs into an SoC (System-on-a-Chip) comes with its penalties. As semiconductor manufacturing technology scales, transistors have become more vulnerable to disturbance, and consequent data corruption, due to soft errors induced by cosmic radiation, packaging radiation, and thermal neutrons [4–6]. Unmitigated, soft errors may lead to quality of service degradation and failure. In addition, *mixed critical* applications—that comprise a mixture of applications with different criticality levels [7]—such as occurring in automotive and avionics systems expect different requirements in terms of reliability. For example, automotive applications such as Anti-lock Breaking (ABS) (control-oriented processing), collision and pedestrian detection (data-oriented processing), and multimedia applications fall into different levels of criticality with respect to errors in computed data.

This becomes even more challenging when applications demand varying requirements based on their execution scenarios. As an example, in [8] it is shown that Soft Error Rates (SERs) are not constant and may highly vary over time based on orbital positions. It is thus imperative to counter the increasing proneness of modern MPSoCs to

* Corresponding author. Tel.: +4991318567315; fax: +4991318525149.
E-mail address: vahid.lari@cs.fau.de, vahid.lari@gmail.com (V. Lari).

http://dx.doi.org/10.1016/j.sysarc.2015.10.004
1383-7621/

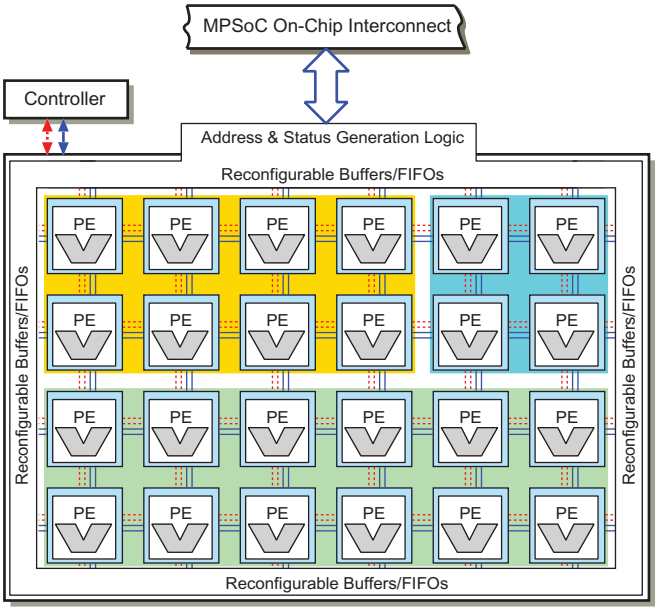

Fig. 1. Tightly Coupled Processor Array (TCPA) used as a loop accelerator in an MPSoC as an example of a coarse-grained reconfigurable array. The different rectangular areas correspond to three applications executing simultaneously on the array.

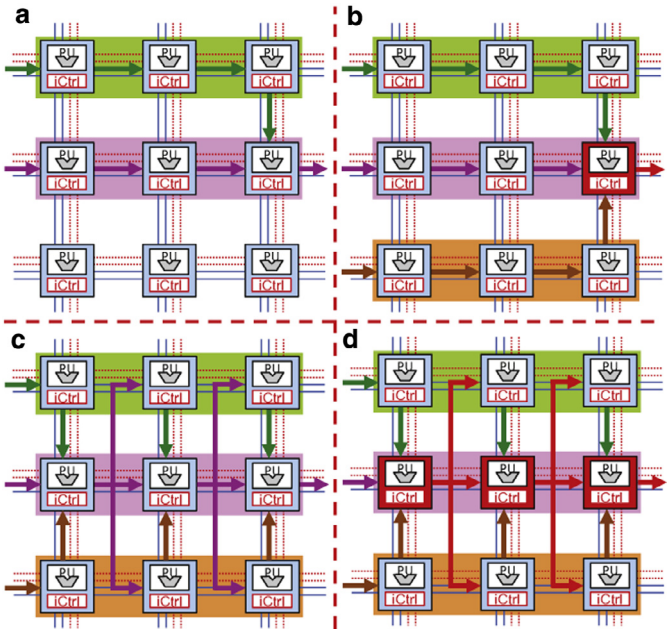

Fig. 2. Redundancy is allocated on demand by claiming identical subarrays to realize (a) DMR, and (b)–(d) TMR. Comparison and voting are performed at the array boundary in (a) and (b) respectively. In (c) and (d), voting is performed in software, respectively in hardware, inside the processor array for immediate detection/correction of potential errors.

errors by applying appropriate fault tolerance measures dynamically. In order to give applications such adaptivity, self-organizing computing paradigms such as invasive computing [9] can be applied, where applications running on an MPSoC and competing for resources are allowed to request an allocation of neighboring processors and later de-allocate them, exploiting and adapting to varying degrees of parallelism and available resources more efficiently.

In this paper, using the principles of invasive computing, we propose an *on-demand* fault tolerance approach, in which individual safety levels may be defined for applications as requirements for execution. Based on such requirements and according to an analysis of reliability that we propose in this paper, the system may adopt structural redundancy schemes such as *dual-modular* (DMR) and *triple-modular* (TMR) on demand and at runtime in order to provide a requested level of protection against soft errors. As a standard metric for safety, we use *Safety Integrity Levels (SILs)* [10] that are explained based on probability of *failure per hour (PFH)* of execution.

Fig. 2 illustrates our approach. For the different options shown in this figure, event upsets occurring in either the SRAM cells such as register file, instruction memory, input/output registers connecting the PE to each other, or in their logical circuits may be trapped or corrected inside the array by DMR (a) or TMR schemes (b)–(d), respectively. In order to achieve this, the compiler replicates loop computations across double (Fig. 2(a)) or triple (Fig. 2(b)–(d)) the number of processors and inserts error detection statements. Here, for voting over the results of replicas, two approaches are proposed, namely, *late voting* at the border of the array or *immediate voting* after each loop iteration. Using these transformations—for applications that demand fault tolerance—multiple target redundancy variants are generated at compile time, one of which may be selected at runtime by the runtime system based on an elaborated reliability calculus.

In summary, the contributions of this paper are three-fold:

- Different redundancy levels, namely *no replication*, *DMR*, and *TMR*, as well as different voting/comparison variants, i. e., voting/comparing the result *lately* at the border of a processor array or *immediately* after each loop iteration, are proposed.
- Mathematical analysis for the gained reliability and the probability of failure in case of each redundancy scheme, derived based on timing characteristics of applications, i. e., timing parameters for error detection latency and execution times on each PE, and observed SER on a system.
- "On-demand" selection of an appropriate redundancy configuration in dependence of reliability requirements for fault-tolerant loop executions on CGRAs.
- Empirical results based on simulations of fault-injections on a real TCPA are provided to confirm the same trends predicted using the mathematical reliability analysis of each redundancy scheme to be applicable to a certain range of observed SER per PE. Whereas in comparison, the mathematical analysis is conservative in its assumptions about error propagations, it may and should be used for runtime decision making on which redundancy scheme to configure in case of systems with high criticality.

The rest of this paper is structured as follows. After a discussion of the state-of-the-art (Section 2), we introduce three different options for structural redundancy in Section 3, and describe the class of applications we consider and the required program transformations to support software replication on replicated hardware claims. Subsequently, in Section 4, we define the criteria of comparison for reliability and derive formulas for their computation. In Section 5, we provide a reliability analysis for each of the proposed on-demand structural redundancy schemes followed by a subsection on fault simulations carried out on real hardware to verify and compare the reliability analysis with observable faults. The paper finally concludes with a summary and suggestions for future work.

2. Related work

Surveys in the literature provide a valuable overview of the general area of fault tolerance. Overviews of the fundamental structures of fault-tolerant computing are available in [11–13], while more recent work [14] covers transient errors and architectures to mitigate them.

CGRA manifest a natural redundancy at the PE level in addition to the instruction level in case of superscalar or Very Long Instruction Word (VLIW) structured PEs. However, compared to standard