

Chapter 53

High-Performance Reconfigurable Computing

Mário Pereira Vestias
Instituto Politécnico de Lisboa, Portugal

ABSTRACT

High-performance reconfigurable computing systems integrate reconfigurable technology in the computing architecture to improve performance. Besides performance, reconfigurable hardware devices also achieve lower power consumption compared to general-purpose processors. Better performance and lower power consumption could be achieved using application-specific integrated circuit (ASIC) technology. However, ASICs are not reconfigurable, turning them application specific. Reconfigurable logic becomes a major advantage when hardware flexibility permits to speed up whatever the application with the same hardware module. The first and most common devices utilized for reconfigurable computing are fine-grained FPGAs with a large hardware flexibility. To reduce the performance and area overhead associated with the reconfigurability, coarse-grained reconfigurable solutions has been proposed as a way to achieve better performance and lower power consumption. In this chapter, the authors provide a description of reconfigurable hardware for high-performance computing.

INTRODUCTION

High-Performance Reconfigurable Computing (HPRC) systems integrate reconfigurable technology in the computing architecture to improve performance. System designers and engineers of HPRC systems are aware of the potential speed up that can be achieved by integrating Field Programmable Gate Arrays (FPGAs) or other reconfigurable devices in their multiprocessing systems. Besides performance, reconfigurable hardware devices also achieve lower power consumption compared to General-Purpose Processors (GPP), a major advantage given the high power consumption of existing HPC machines. Better performance and lower power consumption could be achieved using Application Specific Integrated Circuit (ASIC) technology. However, ASICs are not reconfigurable, turning them application specific. Reconfigurable logic becomes a major advantage when hardware flexibility permits to speed up whatever the application with the same hardware module. The first and most common devices uti-

DOI: 10.4018/978-1-5225-7598-6.ch053

lized for reconfigurable computing are fine-grained FPGAs with a large hardware flexibility. To reduce the performance and area overhead associated with the reconfigurability, coarse-grained reconfigurable solutions has been proposed as a way to achieve better performance and lower power consumption. In this chapter we will provide a description of reconfigurable hardware for high performance computing.

BACKGROUND

High-Performance Reconfigurable Computing (HPRC) is a computing paradigm that combines reconfigurable-based processing (e.g., FPGA technology) with general purpose computing systems, whether single general purpose processors or parallel processors. The idea is to introduce hardware flexibility to accelerate computationally intensive tasks and therefore to achieve higher performance computing compared to platforms without reconfigurable hardware. These systems not only potentially improve the performance relative to non-reconfigurable general-purpose computing systems but also reduce the energy consumption. Saves of up to four orders of magnitude in both metrics are reported for some compute intensive applications (El-Ghazawi, 2008). Almost all HPC vendors provide HPRC solutions materializing their beliefs in the capacities of reconfigurable computing as accelerators for high-performance computing.

The first commercial reconfigurable computing platform for high-performance computing was the Algotronix CHS2x4 (Algotronix) consisting of an array of 1024 processors and 8 FPGAs with 1024 programmable cells each. This architecture was followed by many other reconfigurable proposals for HPRC.

A major concern in the design of HPRC systems is how reconfigurable computing is connected to the non-reconfigurable computing side, whether general-purpose computing or dedicated computing (e.g. general purpose processors or ASICs). A few alternatives exist with different expected performances, cost and flexibility. The typical approach is to consider reconfigurable systems, generally FPGAs, mounted in a board that is connected to the main system using some serial bus to operate as a coprocessor. The approach has a relative low cost but has a severe limitation from the serial communication between the host and the coprocessors whose bandwidth determines that for the architecture to be computationally efficient the computation to I/O ratio must be high. To improve the co-processing solution, a few architectures have implemented direct point-to-point connections among the co-processors. This speeds-up the communications between reconfigurable co-processors but keeps the communication bottleneck between the host system and the co-processing system. Some works have refined the communication between the host and the FPGA co-processors through a dedicated network interface. A well-known example of this architecture is Cray XD1 (Cray Inc., 2006). To speed-up even more the communication between the host CPU and the reconfigurable units, all units can be connected to a single communication network, like a shared memory system. In this case, all processing units see each other as part of a unique architecture with access to shared memories (SGI RASC RC100 blade from Silicon Graphics).

Another dimension in the architectural design of HPRC systems is the granularity of the reconfigurable platform. The granularity of a reconfigurable system is the smallest reconfigurable functional unit. In terms of granularity they tend to be defined as fine-grained or coarse-grained. Raw FPGAs are fine-grained since they can be reconfigured at the bit level. Coarse-grain architectures consist of arrays of coarser units, like arithmetic logic units that are reconfigurable at the word level. Coarse-grained architectures are more amenable to design and reconfigure but are less flexible than fine-grained architectures. When the application to run in the reconfigurable architecture can be bit-level optimized then

12 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/high-performance-reconfigurable-computing/214656

Related Content

A Self-Configurable Event Coverage Approach for Wireless Sensor Networks

Ajay Kaushik, Ravi Teja Yakkali, S. Indu, Feroz Ahmed, Daya Gupta, Raunaq Nayar and Sahil Yadav (2019). *International Journal of Mobile Computing and Multimedia Communications* (pp. 1-18).

www.irma-international.org/article/a-self-configurable-event-coverage-approach-for-wireless-sensor-networks/227358

A Mobile Portal Solution for Knowledge Management

Stefan Berger and Ulrich Remus (2009). *Mobile Computing: Concepts, Methodologies, Tools, and Applications* (pp. 2496-2508).

www.irma-international.org/chapter/mobile-portal-solution-knowledge-management/26675

A Cloud Trusting Mechanism Based on Resource Ranking

Ajai K. Daniel (2020). *Handling Priority Inversion in Time-Constrained Distributed Databases* (pp. 130-155).

www.irma-international.org/chapter/a-cloud-trusting-mechanism-based-on-resource-ranking/249428

Joint Source Channel Coding and Diversity Techniques for 3G/4G/LTE-A: A Review of Current Trends and Technologies

Surajit Deka and Kandarpa Kumar Sarma (2016). *Self-Organized Mobile Communication Technologies and Techniques for Network Optimization* (pp. 190-215).

www.irma-international.org/chapter/joint-source-channel-coding-and-diversity-techniques-for-3g4glte-a/151141

Mobile Network Architecture: 3GPP Generations (UMTS, LTE, and Pre-5G)

(2019). *Mobile Network Forensics: Emerging Research and Opportunities* (pp. 130-205).

www.irma-international.org/chapter/mobile-network-architecture/216751