High-Performance Reconfigurable Computing Granularity

Mário P. Véstias
INESC-ID/ISEL/IPL, Portugal

INTRODUCTION

There are several examples where High-Performance Computing (HPC) systems are being upgraded with reconfigurable technology integrated in the computing architecture to improve the performance. System designers and engineers of HPC systems are aware of the potential speed up that can be achieved by integrating Field Programmable Gate Arrays (FPGAs) in their multiprocessing systems. Besides performance, FPGAs 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. High-performance reconfigurable computing can be fine or coarse-grained. Makimoto's wave forecasts that this decade will be dominated by coarse-grained reconfigurable computing. Many cases exist where coarse-grained reconfigurable hardware achieves better performance and lower power consumption compared to fine-grained logic at the cost of reduced hardware flexibility. This trade-off must be taken into consideration when designing HPC architectures. In this article we will provide a description of fine and coarse-grained reconfigurable hardware for high performance computing. We are mainly concerned with the granularity aspect in high performance reconfigurable 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

DOI: 10.4018/978-1-4666-5888-2.ch348

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 architectural design concern of HPRC systems is the granularity of the reconfigurable platform. A fine granularity level permits a high flexibility in the design of the hardware accelerator, but at the cost of high overhead associated with the reconfigurable logic and interconnections. On the other side, a coarse granularity reduces the flexibility but potentially achieves more efficient hardware designs since the reconfigurability overhead is lower. In the following sections several architectures will be analyzed considering the granularity of the reconfigurable fabric.

HIGH-PERFORMANCE RECONFIGURABLE COMPUTING

HPRC architectures can be classified according to the granularity of the reconfigurable system. 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 fine-grained architectures are usually more adequate than coarse-grained architectures achieving faster solutions. On the other side, coarse-grained architectures are faster and more efficient when running word-level applications.

Formally, we define fine-grained HPRC as those platforms that use fine-grained reconfigurable logic and coarse-grained HPRC as those systems that try to improve the performance and power consumption of reconfigurable systems by increasing the granularity at which operations are computed. In the following, examples of fine and coarse-grained architectures are described in the context of high performance computing.

FINE-GRAINED HPRC

Fine-grained high-performance reconfigurable computers consist of basically on multiple FPGAs for hardware acceleration combined with general-purpose systems. Well-known fine-grained HPRC includes Cray XD1 (Cray Inc., 2006), SRC MAPstation (SRC Computers), Starbridge Hypercomputer (Starbridge Hypercomputers), SGI RASC RC100 blade (Silicon Graphics), Novo-G (George, A., Lam, H., Stitt, G., 2011), Maxwell (Baxter, 2007), Convey HC-1 (Brewer, 2010), (Bakos, 2010) and Convey HC-2 (Convey Computer, 2011).

Cray XD1 consists of multiple general-purpose processors combined with multiple FPGAs in a set of chassis. Each Cray XD1 chassis consist of up to six compute blades and a blade has two AMD Opteron 200 processors and one Xilinx Virtex-II Pro FPGA. Therefore, a single chassis consists of twelve processors with up to 8 GBytes of memory for each processor and six FPGAs (see Figure 1). Each processor has 1 or 2 RapidArray links connected to a non-blocking RapidArray fabric switch. The switch provides up to 96 GB/s of total bandwidth.

The application must be manually partitioned between hardware running on the FPGA and software running on the Opteron processors.