

FASTER Run-time Reconfiguration Management

Cătălin Bogdan Ciobanu
Chalmers University of Technology, Sweden
catalin@chalmers.se

Kyprianos D. Papadimitriou
Foundation for Research and Technology -
Hellas, Greece
kppadim@ics.forth.gr

Dionisios N. Pnevmatikatos
Foundation for Research and Technology -
Hellas, Greece
pnevmati@ics.forth.gr

Georgi N. Gaydadjiev
Chalmers University of Technology, Sweden
georgig@chalmers.se

ABSTRACT

The FASTER project Run-Time System Manager offloads programmers from low-level operations by performing task placement, scheduling, and dynamic FPGA reconfiguration. It also manages device fragmentation, configuration caching, pre-fetching and reuse, bitstream compression, and optimizes the system thermal and power footprints. We propose a micro-reconfiguration aware, configuration content agnostic ISA interface and a technology independent Task Configuration Microcode format targeting Maxeler Data Flow computers and Xilinx XUPV5 platforms. We achieve improved resource utilization with negligible performance overhead. Up to 4Gbps for DMA transfers, and up to 3Gbps for FPGA reconfiguration on Xilinx Virtex-5/6 devices is achieved.

Categories and Subject Descriptors

D.3.4 [Programming Languages]: Processors—*Run-time Environments*; D.4.1 [Operating Systems]: Process Management—*Scheduling*; C.1.3 [Processor Architectures]: Other Architecture Styles—*Adaptable Architectures, Data-flow Architectures, Heterogeneous (Hybrid) Systems*

Keywords

FPGA; Run-time System Manager; Partial Reconfiguration

In recent years, processor clock frequencies stagnated due to power and thermal constraints. The additional transistors provided by new technology generations are transformed into supplementary processors or application-specific functional units. Custom accelerators boost performance with low power consumption by specializing the hardware for the workload. By using reconfigurable logic, new hardware units can be dynamically instantiated on-demand. Such hardware adaptability is often preferable to software in performance critical domains, e.g., High Performance Computing (HPC). The design and verification of run-time reconfigurable systems, however, is a time consuming and difficult process.

The Facilitating Analysis and Synthesis Technologies for Effective Reconfiguration (FASTER) project [1] aims to provide a complete methodology allowing easy implementation and verification of systems constituting multiple accelerators in reconfigurable hardware, primarily focusing on partial dynamic reconfiguration, from initial design to dynamic runtime management. It benefits both from region-based reconfiguration, i.e. dynamic alteration of large portions of the device with precompiled circuits, and micro-reconfiguration, which

reconfigures small parts of the device with circuits specialized at runtime based on the input parameter values.

The FASTER Run-Time System Manager (RTSM) software library offloads programmers from low-level scheduling and resource management decisions on partially reconfigurable FPGA-based systems. The RTSM handles task placement and scheduling, converts technology independent representations into FPGA-specific commands, and communicate with the FPGA configuration port. Additional RTSM functionality includes tasks relocation, configuration caching, pre-fetching and reuse, and bitstream compression. In addition, it manages power and energy consumption and monitors temperature for hotspot management and task migration. The RTSM architectural interface consists of the configuration agnostic Instruction Set Architecture (ISA) extension and the technology independent bitstream format. Our customized Molen ISA extension offers support for micro-reconfiguration and for two-level nested reconfiguration.

At design time, the configuration data and the task specific information (e.g., expected execution time and a region/micro-reconfiguration flag) are encapsulated in a configuration agnostic Task Configuration Microcode block (TCM). The FASTER tool-chain compilation phase provides the baseline schedule and relevant task information such as the size of the reconfigurable areas and the estimated reconfiguration times. The RTSM updates selected parameters at run-time, e.g., the current status of the reconfigurable area (e.g., *empty, being reconfigured*) and of the tasks (e.g., *ready for execution, running*), and profiles the task execution times.

We target Maxeler Technologies systems for HPC and Xilinx University Program (XUP) Virtex-5 boards connected via PCIe for the desktop domain. Maxeler provides large Xilinx Virtex-6 Data Flow Engines, each featuring up to 48GB of dedicated high-speed DRAM. The results show that RTSM can improve resource utilization and can be implemented with minimal overhead on overall system performance. Currently, our custom RTSM software running on the host GPP controls both execution and reconfiguration, by loading the bitstreams from the host memory on demand. Our RTSM has shown very good system-level performance with initial measurements reaching DMA transfers of up to 4Gbps and FPGA reconfiguration speeds of up to 3Gbps.

Acknowledgments

This work was supported by the European Commission in the context of FP7 FASTER project (#287804).

1. REFERENCES

- [1] The FASTER Project. URL www.fp7-faster.eu/.