Editorial

Reconfigurable Computing and Hardware/Software Codesign

Toomas P. Plaks,[1] Marco D. Santambrogio,[2] and Donatella Sciuto[2]

[1] *University of Reading, Berkshire RG6 6AH, UK*
[2] *Politecnico di Milano, 20133 Milano, Italy*

Correspondence should be addressed to Toomas P. Plaks, plakstp@aol.com

Received 23 December 2007; Accepted 23 December 2007

Modern consumer appliances as wireless communication and multimedia systems present very strong requirements for the digital parts of these systems: digital design process must provide solutions which possess high performance, flexibility for multifunctional use, and energy efficiency. Over the past decade, the reconfigurable computing platform has been an emerging approach in scientific research and in practical implementations to meet these requirements.

The special issue on "Reconfigurable Computing and Hardware/Software Codesign" addresses the advances in reconfigurable computing architectures, in algorithm implementation methods, and in automatic mapping methods of algorithms onto hardware and processor spaces, indicating the changes in codesign flow due to the introduction of new, reconfigurable hardware platform. Using this platform, the designer faces a new paradigm of computing and programming: the computing system is capable of run-time and autonomous modification of its functionalities following the changing needs of applications.

This new scenario of hardware/software codesign provides a great improvement in the embedded system design and implementation. To cope effectively and timely with the new challenges, the new and more sophisticated dynamic reconfiguration strategies together with codesign methods have to be developed.

In the first paper, "Design flow instantiation for run-time reconfigurable systems: a case study," Y. Qu et al. present a design flow instantiation for run-time reconfigurable systems using a real-life application—part of a WCDMA decoder. The design flow is roughly divided into two parts: system level and implementation. At system level, hardware resource estimation and performance evaluation are applied. At implementation level, technology-dependent tools are used to realize the run-time reconfiguration. The results

show that run-time reconfiguration can save 50% of the area when compared to a functionally equivalent fixed system and achieves 30 times speedup in processing time when compared to a functionally equivalent pure software design.

In "A flexible system level design methodology targeting run-time reconfigurable FPGAs," F. Berthelot et al. present an automatic design generation methodology for heterogeneous architectures. This method automatically generates designs for fixed and partially reconfigurable parts of an FPGA and enables a reconfiguration prefetching technique to minimize reconfiguration latency and buffer-merging techniques to minimize memory requirements of the generated design. This concept has been applied to different wireless access schemes, based on a combination of OFDM and CDMA techniques.

The next paper, "RRES: a novel approach to the partitioning problem for a typical subset of system," by G. B. Knerr et al., integrates some of the most powerful approaches for system partitioning into a consistent design framework for wireless embedded systems, which has led to the development of an entirely new approach for the system partitioning problem. The paper introduces the restricted range exhaustive search algorithm and compares this to popular and well-reputed heuristic techniques based on tabu search, genetic algorithm, and the global criticality/local phase algorithm. This search algorithm proves superior performance for a set of system graphs featuring specific properties found in human-made task graphs, since it exploits their typical characteristics such as locality, sparsity, and their degree of parallelism.

The paper "Software-controlled dynamically swappable hardware design in partially reconfigurable systems," by C. Huang and H. Pao-Ann, considers different wrapper designs for hardware designs such that they can be enhanced

with the capability for dynamic swapping controlled by software. A hardware design with proposed wrappers can be swapped out of the partially reconfigurable logic at run-time in some intermediate state of computation and then swapped in when required to continue from that state. With the capability for dynamic swapping, high-priority hardware tasks can interrupt low-priority tasks in real-time embedded systems so that the utilization of hardware space per unit time is increased.

In "DART: a functional-level reconfigurable architecture for high energy efficiency," S. Pillement et al. deal with functional-level reconfiguration to improve energy efficiency. The paper presents the DART architecture, which supports two modes of reconfiguration: fine-grained and functional level, to achieve the optimized solutions. The compilation framework is built using compilation and high-level synthesis techniques. As a proof of the concept, a 3G mobile communication application has been implemented and the VLSI design of a 0.13 μm CMOS SoC implementing a specialized DART cluster is presented.

The last paper of this issue "Exploiting process locality of reference in RTL simulation acceleration," by A. D. Blumer and C. D. Patterson, addresses the simulation acceleration of digital designs. An analysis of six register transfer level (RTL) code bases shows that only a subset of the simulation processes is executing at any given time. Run-time adaptations are made to ensure that acceleration resources are not wasted on idle processes, and these adaptations may be effected through process migration between software and hardware. Finally, the paper describes an implementation of an embedded, FPGA-based migration system; the empirical data are obtained for use in mathematical and algorithmic modelling of more complex acceleration systems.

<div align="right">

Toomas P. Plaks
Marco D. Santambrogio
Donatella Sciuto

</div>