Des Autom Embed Syst (2014) 18:119–120
DOI 10.1007/s10617-014-9136-7

EDITORIAL

Editorial

Davide Bertozzi · Giorgos Dimitrakopoulos · Sören Sonntag

Published online: 26 April 2014

The new computing systems environment is characterized by the convergence of the commodity, embedded, and supercomputer markets. At the same time, evolution is constrained by a design complexity crisis and by power consumption limitations, where performance gains cannot rely on technology scaling. The technology wave is directed at the commoditization of the technology base that was previously associated only with specific platforms. Key building blocks have been identified across domains that rely mostly on many-core processors and heterogeneous computation and extend into data center, microservers, wireless, automotive, and space applications.

The connecting piece of all heterogeneous hardware components is the interconnect architecture. That might be an on-chip interconnect for modern embedded systems or an off-chip interconnect for powerful HPC solutions. With their increasing number of processing nodes these interconnects pose significant challenges to the designer. On the HPC/datacenter side, systems are expected to scale well beyond today's 1 million nodes, serving thousands of simultaneous tenants—decoupled in performance—and tens of millions of users. These emerging fabrics are likely to be increasingly embedded within the CPU complex, together with support for networking virtualization and also for a growing multitude of new and legacy protocols.

At the same time, modern SoC platforms face increasingly complex interconnect-related problems at the chip level that go beyond physical integration and include the need for scalable performance, design-time flexibility and workload adaptivity. Modern SoCs for wireless and automotive platforms contain multiple heterogeneous processing nodes, memory and I/O

D. Bertozzi (✉)
University of Ferrara, Ferrara, Italy
e-mail: davide.bertozzi@unife.it

G. Dimitrakopoulos
Democritus University of Thrace, Xanthi, Greece
e-mail: dimitrak@ee.duth.gr

S. Sonntag
Intel Corp., Munich, Germany
e-mail: soeren.sonntag@intel.com

controllers and a vast set of other feature-rich domain-specific components. Design teams face a growing pressure to make end-user products more efficient, more reliable, and less expensive, all in the face of faster product life cycles. On-chip communication infrastructure and Network-on-Chip solutions are expected to have an increasingly large share of the SoC market with the requirement of constant efficiency increase and support for more sophisticated architectures.

This special issue brings together the best five papers of the 7th International Workshop on Interconnection Network Architecture: On-Chip, Multi-Chip (INA-OCMC) held in conjunction with the 8th HiPEAC Conference in Berlin, January 2013. The papers represent extended and peer-reviewed versions of the original work focusing on many aspects of network design and covering both on- and off-chip networks. In particular, these papers include the following topics:

- *Centaur: A Hybrid Network-on-Chip Architecture Utilizing Micro-Network Fusion* deals with the design of a fused NoC architecture composed of a ultra-low-latency ring-based micro-network optimized for control messages and a conventional throughput optimized NoC that handles the rest of the traffic. The proposed architecture increases performance with a negligible area-power overhead versus traditional designs.
- *RIDER: Ring Deflection Router with Buffers* presents a novel router design that embraces both deflection-based switching using an internal ring structure and minimal buffering to achieve low-cost NoC router architectures.
- *Automated Design Space Exploration for FPGA-based Heterogeneous Interconnects* presents an automated NoC synthesis methodology targeting FPGA technologies. The developed interconnect synthesis method can produce optimized and customized topologies using crossbars and shared busses driven by the application traffic requirements and other implementation constraints.
- *Randomizing task placement and route selection do not randomize traffic (enough)* develops a theoretical model that predicts the performance of a generic dragonfly network under uniform traffic and delivers performance-optimal dragonflies at a minimum cost. Then, based on the theoretical predictions that have been validated via simulations, quantifies the effectiveness of randomizing task placement of these topologies.
- *Tandem Queue Weighted Fair Smooth Scheduling* introduces two light-weight, fair schedulers that accommodate an arbitrarily large number of requestors and are suitable for ultra high-speed links (in the order of 100 Gbps). The presented techniques improve short-term fairness and are able to deliver very smooth service when flow weights are approximately equal.
- This special issue is concluded with a future outlook paper regarding Network-on-Chip technology addressing upcoming research challenges related to architecture, technology and physical integration.

As guest editors, we thank the authors and presenters for their contributions, and the Editor-in-Chief for enabling this special issue.

Davide Bertozzi
Giorgos Dimitrakopoulos
Sören Sonntag