Editorial

Mei Yang* and Yingtao Jiang

Department of Electrical and Computer Engineering,
University of Nevada, Las Vegas,
4505 Maryland Pkwy, Box 454026,
Las Vegas, NV 89154, USA
E-mail: Mei.Yang@unlv.edu E-mail: yingtao@egr.unlv.edu
*Corresponding author

Peng Liu

Department of Information Science and Electronic Engineering,
Zhejiang University,
No. 2 Yugu Rd., Hangzhou, Zhejiang 310027, China
E-mail: liupeng@isee.zju.edu.cn

Maurizio Palesi

Department of Computer Science and Telecommunications Engineering,
6, Viale Andrea Doria, 95125 Catania, Italy
E-mail: mpalesi@diit.unict.it

Biographical notes: M. Yang received her PhD in Computer Science from the University of Texas at Dallas in August 2003. She is currently an Assistant Professor in the Department of Electrical and Computer Engineering at University of Nevada, Las Vegas (UNLV). Before she joined UNLV, she was an Assistant Professor in the Department of Computer Science at Columbus State University from August 2003 to August 2004. Her research interests include computer architectures, wireless sensor networks, and embedded systems.

Y. Jiang received his PhD in Computer Science from the University of Texas at Dallas, Richardson, Texas, in 2001. Immediately after graduation, he joined the Department of Electrical and Computer Engineering, University of Nevada, Las Vegas as an Assistant Professor. He was promoted to Associate Professor rank in 2007 at the same university. His research interests include algorithms, VLSI architectures, and circuit level techniques for the design of DSP, networking, and telecommunications systems; computer architectures; sensors and measurement instrument; nanotechnologies; and biomedical signal processing, instrumentation, and medical informatics.

P. Liu received his BS in Optical Engineering and MS in Optical Engineering from Zhejiang University, Hangzhou, China, in 1992 and 1996, respectively, and PhD in Communication and Electrical Engineering from Zhejiang University in 1999. In 1999, he joined the Faculty of the Information Science and Electronic Engineering department at Zhejiang University, where he was promoted to Associate Professor in 2002. His research interests include embedded processor microarchitecture, multiprocessor system-on-chip architectures, on-chip interconnection networks, and VLSI design. He is currently a Visiting Scholar at University of Rochester working on high-performance computer architectures.

M. Palesi received his MS and PhD in Computer Engineering from the Università di Catania, Catania, Italy, in 1999 and 2003, respectively. Since December 2003, he has held a research contract as an Assistant Professor with the Dipartimento di Ingegneria Informatica e delle Telecomunicazioni, Facoltà di Ingegneria, Università di Catania. His research focuses on platform-based system design, design space exploration, low-power techniques for embedded systems, and network-on-chip architectures. His main research interests include design space exploration methodologies, networks-on-chips, and low-power design. He serves on the editorial board of very large scale integration (VLSI) design as an Associate Editor since May 2007. He has served as a Guest Editor for the *VLSI Design Journal – Special Issue on Networks-on-Chip* in 2008. He serves as the TPC for a number of IEEE/ACM conferences and workshops.

1 Introduction

The ever increasing needs of present and future computation-intensive applications have fuelled multiple research programs worldwide to develop new and innovative computing techniques and architectures that can deliver accelerated computing capacity, superior data throughput, and the ability to aggregate substantial distributed computing power. More noticeably, in recent years, the field of computer architecture has seen a rapidly expanding interest in power-efficient computing at architectural and software level. In this special issue, the guest editors have put together some of the recent new developments and trends in power-efficient, high performance general purpose and application specific computing architectures. The ten selected papers in this special issue span a wide range of topics, including networks-on-chip (NoC) architectures, embedded processors, microprocessors, and parallel and distributed computing systems.

The rapid growth of circuit density has brought us into the multi-core CPU era and on the eve of a many-core (16 or more cores per chip) era. To provide functionally correct, power-efficient, and reliable inter-chip communication, NoC has been proposed as a promising interconnection architecture for both chip multiprocessors (CMPs) and multiprocessors system-on-chip (MPSoCs). The first six papers in this issue present the study of the problems falling into three important research topics in NoCs: scheduling and mapping methods, design of network interfaces, and design space exploration. The first paper by P. Ghosh and A. Sen investigates the IP mapping and voltage islanding problem for MPSoCs which targets to minimise the overall power consumption without violating any design constraint. They provide both optimal solution by solving a mixed integer linear programming and heuristic solution based on random greedy selection. The second paper by X. Gu et al. presents a hardware/software approach – a synergetic operating unit (SOU) to effectively manage the object scheduling and synchronisation for the CMP systems with distributed memories. Compared with the software director, the SOU significantly reduces the average delay for object scheduling and synchronisation as well as the code size for RTOS and director. The third paper by Y. Yang et al. proposes a parallel and pipeline processing method for block cipher algorithms using purely software implementation on a fault-tolerant NoC, named NePA, which is enabled with error correction capability. The parallel software implementation is flexible and can run on various NoC platforms.

The fourth paper by B. Xia et al. proposes the network interface (NI) design based on the mutual interface definition, which decouples the resource dependent part and the resource independent part. By building an NI component library, designers can build the NI according to the requirement by selecting the optimum components from the library. The fifth paper by G. Kornaros describes a methodology and a tool chain for design space exploration of temporal encoding schemes, referred as NCXplorer. NCXplorer aims to assist the designer to achieve the best fit regarding to both switching activity and reduction of crosstalk effects. The sixth paper by D. Pani et al. focuses on the problem of efficiently validating complex HDL designs and proposes a multi-parametric cycle-accurate simulation framework, SysCgrid. SysCgrid is designed to provide automatic generation and parallel execution of multi-parametric simulations by exploiting the computational power of a cluster/grid computing infrastructure through message passing interface (MPI).

Software defined radios (SDR), which use a combination of software and hardware to dynamically support multiple communications standards, have been widely recognised as one of the most important new technologies for wireless communication systems. The next paper by C. Jenkins et al. presents instruction set architecture extensions and hardware design to boost up the advanced encryption standard (AES) processing on a multi-threaded SDR platform, the Sandbridge digital signal processor (DSP). The proposed design achieves significant speedup compared with the fastest software implementation of AES and can be used on other SDR embedded processors.

Modern microprocessor designs usually employ large on-chip caches to overcome the memory-wall problem. Based on the dynamic non-uniform cache architecture (D-NUCA) scheme, the paper by A. Bardine et al. proposes the way-adaptable D-NUCA cache design to improve the average access delay and the power efficiency of the memory hierarchy. Evaluation of the proposed scheme over several benchmarks from SPEC CPU2000 and the NAS parallel benchmarks suites shows significant improvements of both performance and power figures w.r.t. conventional D-NUCA. The next paper by W. Zhang focuses on enhancing data reliability against soft errors in caches and proposes to add the replica victim cache to accommodate the replicas that are frequently evicted from the L1 data cache. The experimental results reveal that a replica victim cache improves the reliability of the L1 data cache significantly with small area overhead.

Parallel disk systems have been widely used to support a variety of data-intensive applications running on high-performance computing platforms. The last paper by F. Shen et al. focuses on reliability analysis for parallel disk systems and develops a quantitative reliability model for energy-efficient parallel disk systems with data mirroring using a Markov process. Based on the proposed model, a reliability analysis tool is developed to evaluate the reliability and energy efficiency of parallel disk systems with two power modes.

We hope this issue represents an interesting combination of current research activities in power-efficient and high-performance computing architectures and systems and provides valuable research results. And we would like to thank all the authors who submitted papers to this special issue and the reviewers for their efforts and excellent work. Also, we would like to thank the Editor-in-Chief, Dr. Nadia Nedjah, for her great support and guidance to our editorial work and Ms. Liz Harris for her valuable work in coordinating this special issue.