Editorial

Introduction: Special Section on Architecture of Future Many Core Systems

1. Introduction

In the past decades, high-performance computing has played a key role to advance the emerging technologies from social networks, search engines, and cell phones to humanoid robotics. The ever-increasing growth of high-performance computing has continuously brought more opportunities for system designers to revisit both enterprise and end-user applications. The significant advancement of high-performance computing is owed to many-core architecture which has been the main focus of numerous academic and industrial research projects in the last years and it is expected to continue as a dominant research topic in the next decade.

Despite significant studies in many-core architectures, further advancement is prohibited by several important issues such as power wall, reduced yield, memory wall, and design complexity. In particular, *Future Many Core Systems* (FMCSs) require more specialized architectures in order to further integrate the advantages of more than ten billions transistors available on-chip. In addition to such design challenges, with the paradigm shift in daily applications used in Cloud and datacenters from E-commerce, On-Line Transactional Processing (OLTP), gaming, online storage, and Genomics to Big Data applications such as clustering and graph analytics, FMCSs should be designed such that the new applications can benefit from emerging multi-core systems.

To foster the research momentum on many-core architectures, Elsevier *Microprocessors and Microsystems* has scheduled a special issue on FMCS addressing both micro-architectural and system-level issues. For this special issue, we have received 14 submissions from nine countries, out of which seven submissions have been selected in the first round of review process. Finally, after detailed review of the revised submissions, four papers have been accepted for publication in this issue. More than sixty reviewers, carefully selected among distinguished researchers from both industry and academia, have helped us to evaluate manuscripts submitted for publication in this special issue.

2. Contributions

The first paper is entitled "*ElCore: Dynamic Elastic Resource Management and Discovery for Future Large-Scale Manycore Enabled Distributed Systems*" by Javad Zarrin, Rui L. Aguiar, and João Paulo Barraca, addresses the problem of resource management for the upcoming large-scale many-core computing environments by exploring different resource allocation schemes. This work presents an elastic, dynamic, and fully decentralized resource management architecture which can be exploited in such distributed environments. The proposed architecture comprises a set of modules which can dynamically be instantiated on or allocated to an arbitrary node on demand. Experimental results demonstrate that the proposed resource management scheme is highly scalable and accurate accuracy for resource allocation and provides a remarkable enhancement in the entire system performance.

The second paper, by Giuseppe Tuveri, Paolo Meloni, Francesca Palumbo, Giovanni Pietro Seu, Igor Loi, Francesco Conti and Luigi Raffo, entitled "*On-the-fly Adaptivity for Process Networks over Shared-Memory Platforms*", proposes an adaptive, run-time and dynamic process migration technique, targeting power and throughput efficiency of shared-memory platforms. The proposed methodology is able to improve the applications' parallelism degree and resource utilization with the available resources or the application workload variation. Experiments, carried out by a shared-memory platform prototyped on an FPGA, indicate the efficiency of the proposed approach in terms of speed-up enhancement, power saving, and performance overhead.

A thread parallelization methodology for accelerating the Quorum Planted Motif Search algorithm in multi-core and many-core systems is presented in the third contribution entitled "*Parallel Implementation of Quorum Planted (l,d) Motif Search on Multicore/Many-core Platforms*" by Fazeleh Sadat Kazemian, Mahmood Fazlali and Ali Katanforoush. The proposed methodology benefits from dynamic scheduling of threads and loop parallelization provided by OpenMP library, which result in efficient load balancing of tasks. Experimental results demonstrate that the proposed approach is able to speed up the qPMS9 algorithm and outperforms prior parallel and sequential algorithms by decreasing tasks execution time.

The last paper is entitled "*ASHA: An Adaptive Shared-Memory Sharing Architetcure for Multi-Programmed GPUs*" by Hamed Abbasitabar, Mohammad Hossein Samavatian and Hamid Sarbazi-Azad. This work targets to enhance utilization of GPU resources and improves the spatial multi-programming by exploiting the proposed adaptive shared-memory sharing architecture. Experimental results show that the proposed architecture is able to speed up programs which can be performed simultaneously and alleviate performance degradations which is a major overhead in the spatial multi-programming without ability of resource sharing.

3. Concluding Remarks

Many-core systems play an important role in future computing systems used in different applications. Processing the huge volume of the data generated every day in different areas, require new

architectural supports. Such architectures employ numerous processing cores which should be managed and utilized efficiently. This special issue is meant to address some of the important issues related to the design and use of such many-core systems.

We would like to thank all the authors who have submitted papers to this special issue. We would also like to offer our sincere thanks to the reviewers, without whom this issue would not have been possible and who provided very constructive comments and suggestions to the authors. We would also like to express our gratitude to the Editor-in-Chief, Prof. Jozwiak, for his support and help in making this special issue a reality. We hope you will enjoy the papers in this special issue.

Guest-Editors

Hossein Asadi
Sharif University of Technology, Iran

Paolo Ienne
Ecole Polytechnique Fédérale de Lausanne (EFPL), Switzerland

Hamid Sarbazi-Azad
Sharif University of Technology and Institute for Research in Fundamental Sciences (IPM), Iran

Hossein Asadi received the B.Sc. and M.Sc. degrees in computer engineering from the Sharif University of Technology (SUT), Tehran, Iran, in 2000 and 2002, respectively, and the Ph.D. degree in electrical and computer engineering from Northeastern University, Boston, MA, USA, in 2007. He was with EMC Corporation, Hopkinton, MA, USA, as a Research Scientist and Senior Hardware Engineer, from 2006 to 2009. From 2002 to 2003, he was a member of the Dependable Systems Laboratory, SUT, where he researched hardware verification techniques. From 2001 to 2002, he was a member of the Sharif Rescue Robots Group. He has been with the Department of Computer Engineering, SUT, since 2009, where he is currently a tenured Associate Professor. He is the Founder and Director of the Data Storage Systems Laboratory at SUT. He has authored and co-authored more than fifty technical papers in reputed journals and conference proceedings. His current research interests include data storage systems and networks, solid-state drives, operating system support for I/O and memory management, and reconfigurable and dependable computing.

Dr. Asadi was a recipient of the Technical Award for the Best Robot Design from the International RoboCup Rescue Competition, organized by AAAI and RoboCup, a recipient of Best Paper Award at the 15th CSI Internation Symposium on Computer Architecture and Digital Systems (CADS) in 2010, and the Distinguished Lecturer Award from SUT in 2010, one of the most prestigious awards in the university. He is also recipient of *"Extraordinary Ability in Science"* visa from US Citizenship and Immigration Services in 2008. He has served as the publication chair of several conferences during the past three years. He has also served as a Guest Editor of IEEE Transactions on Computers and as a Program Co-Chair of the 18th International Symposium on Computer Architecture & Digital Systems (CADS2015). He will also serve as a Program Chair of 22th National CSI Computer Conference (CSICC2017). He is a member of ACM and a senior member of IEEE.

Paolo Ienne has been a Professor at the EPFL since 2000 and heads the Processor Architecture Laboratory (LAP). Prior to that, from 1990 to 1991, he was an undergraduate researcher with Brunel University, Uxbridge, U.K. From 1992 to 1996, he was a Research Assistant at the Microcomputing Laboratory (LAMI) and at the MANTRA Center for Neuro-Mimetic Systems of the EPFL. In December 1996, he joined the Semiconductors Group of Siemens AG, Munich, Germany (which later became Infineon Technologies AG). After working on datapath generation tools, he became Head of the embedded memory unit in the Design Libraries division. His research interests include various aspects of computer and processor architecture, electronic design automation, computer arithmetic, FPGAs and reconfigurable computing, and multiprocessor systems-on-chip.

Dr. Ienne was a recipient of Best Paper Award at the 40th Design Automation Conference (DAC) in 2003, at the International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES) in 2007, at the 19th International Conference on Field-Programmable Logic and Applications (FPL) in 2009, and at the 20th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA) in 2012. In 2008, he was General Co-Chair of the 6th IEEE Symposium on Application Specific Processors (SASP) and Guest Editor of a Special Section on Application Specific Processors which appeared in October 2008 on the IEEE Transactions on Very Large Scale Integration Systems. In 2010, he was the Program Sub-committee Chair of the Design Automation Conference (DAC) on High-Level and Logic Synthesis. From 2010 to 2012, he was a Topic Co-Chair of Design Automation and Test in Europe (DATE) for Architectural and High-Level Synthesis topic. In 2011, he was a Program Co-Chair of the 20th IEEE Symposium on Computer Arithmetic (ARITH) and a Program Co-Chair of the 22nd IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP). In 2012, Ienne was a Guest Editor of the Special Section on Computer Arithmetic of the IEEE Transactions on Computers. In 2014, he was Chair of the Program Committee of the 23rd International Workshop on Logic & Synthesis (IWLS) and a Guest Editor of the IEEE Micro Special Issue on Reconfigurable Computing. He has been a member of some fifty program committees of international workshops and conferences in the areas of design automation, computer architecture, embedded systems, compilers, FPGAs, and asynchronous design. He has been an associate editor of ACM Transactions on Design Automation of Electronic Systems (TODAES), since 2011, and of ACM Computing Surveys (CSUR), since 2014.

Hamid Sarbazi-Azad received his BSc in electrical and computer engineering from Shahid-Beheshti University, Tehran, Iran, in 1992, his MSc in computer engineering from Sharif University of Technology, Tehran, Iran, in 1994, and his PhD in computing science from the University of Glasgow, Glasgow, UK, in 2002. He is currently a professor in the department of computer engineering at Sharif University of Technology, and heads the School of Computer Science, Institute for Research in Fundamental Sciences (IPM), Tehran, Iran. His research interests include high-performance computer/memory architectures, NoCs and SoCs, parallel and distributed systems, performance modelling/evaluation, and storage systems, on which he has published more than 300 refereed conference and journal papers. He received Khwarizmi International Award in 2006, TWAS Young Scientist Award in engineering sciences in 2007, and Sharif University Distinguished Researcher awards in years 2004, 2007, 2008, 2010 and 2013.

He is a member of managing board of Computer Society of Iran (CSI), and has served as the editor-in-chief for the CSI Journal on Computer Science and Engineering, an associate editor for IEEE Transactions on Computers and ACM Computing Surveys, and editorial board member for Elsevier's Computers & Electrical Engineering, International Journal of Computers & their Applications, and Journal of Parallel and Distributed Computing and Networks. He has also served as a guest-editor for several special issues on high-performance computing architectures and networks in related journals.