Special Section on High-Performance Computing for Embedded Multi-Core Systems. JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY 26(3): 373–374 Mar. 2011. DOI 10.1007/s11390-011-1140-9

Preface

Multi-core processor brings a major technical innovation in computing hardware. The leap from single-core to multi-core technology has permanently altered the concept about computing. Embedded devices with multi-core technology will rapidly spread throughout the world, and the shift from single-core to multi-core posts a lot of challenges. This special section aims to address some of these challenges by including three invited papers and six regular papers.

The nine papers cover a variety of subjects in high-performance computing for embedded multi-core systems, and represent the state-of-the-art techniques in the field:

In the first paper entitled "Unified UDispatch: A User Dispatching Tool for Multicore Systems", the authors study how to dispatch threads to processor cores so as to improve the system performance in multi-core systems. The authors also provide an extension, called UDispatch+, to dispatch threads without any modification of application source codes. The UDispatch and UDispatch+ are integrated and wrapped for more portability, and introduced as a tool called Unified UDispatch (UUD) with more detailed experiments and description. It can dispatch the application threads to specific cores at the discretion of users with up to 171.8% performance improvement on a 4-core machine.

The second paper is entitled "Data Transmission with the Battery Utilization Maximization". In this paper, the authors explore strategies to maximize the amount of data transmitted by a 3G module under a given battery capacity. In particular, the authors present algorithms under different workload configurations with and without timing constraint considerations. Experiments have been conducted to verify the validity of the strategies and provide insights in energy-efficient data transmission.

In the third paper entitled "Leakage-Aware Modulo Scheduling for Embedded VLIW Processors", a leakage-aware modulo scheduling algorithm is proposed to achieve leakage energy saving for applications with loops on VLIW architectures. The proposed algorithm is designed to maximize the idleness of function units integrated with the dual threshold domino logic, and reduce the number of transitions between the active and sleep modes. The technique has been implemented into the Trimaran compiler, and the experiments have been conducted using a set of embedded benchmarks from DSPstone and Mibench on the cycle-accurate VLIW simulator of Trimaran. The results show that the technique achieves significant leakage energy saving compared with a previously published DAG-based leakage-aware scheduling algorithm.

In the fourth paper entitled "Energy Efficient Block-Partitioned Multicore Processors for Parallel Applications", based on the voltage island and dynamic voltage and frequency scaling (DVFS) techniques, the authors investigate the energy efficiency of block-partitioned multi-core processors, where cores are grouped into blocks with the cores on one block sharing a DVFS-enabled power supply. The authors study both symmetric and asymmetric block configurations, and develop a system-level power model. The energy efficiency of different block configurations is further evaluated through extensive simulations with both synthetic as well as a real life application.

The fifth paper entitled "A Resource-Efficient Communication Architecture for Chip Multiprocessors on FPGAs" presents the novel design and implementation of a resource-efficient communication network for multiprocessors on FPGAs. The technique can reduce not only the required number of routers for a given number of PEs (processing elements) by introducing a new PE-router topology, but also the resource requirement of each router. The communication network relies on the NEWS channels to transfer packets in a pipelined fashion following the path determined by the routing network. The experimental results on various Xilinx FPGAs show good performance in the typical range of network load for multiprocessor applications.

The sixth paper is entitled "VERTAF/Multi-Core: A SysML-Based Application Framework for Multi-Core Embedded Software Development". In this paper, the authors present a new VERTAF/Multi-Core framework and show how software code can be automatically generated from SysML models of multi-core embedded systems. The authors illustrate how model-driven design based on SysML can be seamlessly integrated with Intel's threading building blocks (TBB) and the quantum framework (QF) middleware. A digital video recording system is used to illustrate the benefits of the framework. The experiments show how SysML/QF/TBB help in making multi-core embedded system programming model-driven, easy, and efficient.

In the seventh paper entitled "Configuration Reusing in On-Line Task Scheduling for Reconfigurable Computing Systems", the authors present a technique called reusing-based scheduling (RBS), for on-line scheduling and placement in which configuration reusing is considered as a main characteristic in order to reduce reconfiguration overhead and decrease total execution time of the tasks. Several experiments have been conducted on the proposed algorithm. The results show considerable improvement in overall execution time of the tasks and run-time complexity of the algorithm.

The eighth paper is entitled "Partitioning the Conventional DBT System for Multiprocessors". In this paper, the authors propose a novel multithreaded architecture for DBT (dynamic binary translation) systems through partitioning distinct function modules to adequately utilize multiprocessors. This novel architecture divides the common working routine of a DBT system into three phases: dynamic translation, optimization, and translated code execution, and then ramifies them into different threads to enable them concurrently executed. In this architecture, several efficient methods are presented

to cope with intractable works puzzling most researchers, such as communication mechanism, cache layout, and mutual exclusion between threads. Experimental results using SPECint 2000 indicate that this new architecture can achieve a 10.75% performance improvement compared to previous work.

In the ninth paper titled "Energy Efficiency of a Multi-Core Processor by Tag Reduction", the authors extend the tag reduction from a single-core processor to a multi-core processor to investigate the potential of energy saving for multi-core processors. The authors formulate this problem as an equivalent one which is to find an assignment of the whole instruction pages in the physical memory to a set of cores such that the tag-reduction conflicts for each core can be mostly avoided or reduced. The authors then propose three algorithms to solve this assignment problem. The experimental data have been collected from the real operating system instead of the traditional way using a processor simulator which cannot simulate operating system functions and the full memory hierarchy. The experimental results show that the algorithms outperform the previous work.

Finally, we would like to take this opportunity to thank the authors for their contributions to this special issue, thoughtful reviewers, and the JCST editorial office staff for their excellent work.

Guest Editors:

Min-Yi Guo, Professor, Department of Computer Science and Engineering, Shanghai Jiao Tong University
Shanghai 200240, China guo-my@cs.sjtu.edu.cn

Zi-Li Shao, Associate Professor, Department of Computing, The Hong Kong Polytechnic University
Hung Hom, Kowloon, Hong Kong, China cszlshao@comp.polyu.edu.hk

Edwin Hsing-Mean Sha, Professor, Department of Computer Science, Erik Jonsson School of Eng. & C.S.
The University of Texas at Dallas, Richardson, TX 75083-0688, U.S.A. edsha@utdallas.edu

Min-Yi Guo received the B.Sc. and M.E. degrees in computer science from Nanjing University, China; and the Ph.D. degree in computer science from the University of Tsukuba, Japan. He is currently distinguished professor and head of the Department of Computer Science and Engineering, Shanghai Jiao Tong University (SJTU), China. Before joined SJTU, he had been a professor and department chair of School of Computer Science and Engineering, University of Aizu, Japan. He received the National Science Fund for Distinguished Young Scholars from NSFC in 2007. His present research interests include parallel/distributed computing, compiler optimizations, embedded systems, pervasive computing, and cloud computing. He has more than 230 publications in major journals and international conferences in these areas, including IEEE Trans. Parallel and Distributed Systems, IEEE Trans. Nanobioscience, ACM Trans. Autonomous and Adaptive Systems, J. Parallel and Distributed Computing, INFOCOM, IPDPS, ICS, CASES, ICPP, WWW, PODC, etc. He received 5 best paper awards from international conferences. He is on the editorial board of IEEE Trans. Parallel and Distributed Systems, IEEE Trans. Comput., and J. Comput. Sci. & Tech. He is a senior member of IEEE, member of ACM, IEICE IPSJ, and CCF.

Zi-Li Shao received the B.E. degree in electronic mechanics from the University of Electronic Science and Technology of China, Sichuan, China, in 1995, and the M.S. and Ph.D. degrees from the Department of Computer Science, University of Texas at Dallas, in 2003 and 2005, respectively. He has been an associate professor with the Department of Computing, Hong Kong Polytechnic University, since 2010. His research interests include embedded software and systems, real-time systems, and related industrial applications.

Edwin Hsing-Mean Sha received the Ph.D. degree from the Department of Computer Science, Princeton University, USA in 1992. From August 1992 to August 2000, he was with the Department of Computer Science and Engineering at University of Notre Dame, USA. Since 2000, he has been a tenured full professor in the Department of Computer Science at the University of Texas at Dallas. He has published more than 270 research papers in refereed conferences and journals. He has served as an editor for many journals, and as program committee member and chair for numerous international conferences. He received Teaching Award, Microsoft Trustworthy Computing Curriculum Award, NSF CAREER Award, and NSFC Overseas Distinguished Young Scholar Award. He also receirved the 2011 ACM Best Paper Award from ACM Transactions on Design Automation of Electronic Systems (TODAES). He was awarded as Chiang Jiang Honorary Chair Professorship and China Thousand Talents Program associated with Hunan University, China.