

Special Issue on Network and Parallel Computing

Vijayalakshmi Srinivasan¹ · Yunquan Zhang²

Published online: 16 October 2015
© Springer Science+Business Media New York 2015

It is our pleasure to bring you this special issue of the *International Journal of Parallel Programming* on network and parallel computing. Prior to the publication of this special issue, these journal articles were presented at the 12th IFIP International Conference on Network and Parallel Computing (NPC' 2015) held from September 17 to September 19, 2015 in New York, USA. NPC continues to be a top-tier venue for researchers to present their work in the areas of distributed and parallel computing.

NPC 2015 attracted submissions in the areas of parallel and multicore systems and applications, and also in network scalability and security. We had a strong selection committee, and averaged four reviews per paper. The committee was very selective in choosing the papers, and only 13 journal articles were accepted for presentation at the conference. These 13 top-quality journal articles are included in this special issue, covering a variety of topics ranging from parallel computing applications, including networks, virtualization and scalability to multicore and accelerator architectures. We have placed the papers into four categories.

1 Networks

HyperFatTree: A Large-scale Tree-based Network with Low-radix Switches, Yong Su, Zheng Cao, Zhiguo Fan, Zhan Wan, Xiaoli Liu, En Shao, Xuejun An and Ninghui

✉ Vijayalakshmi Srinivasan
viji@us.ibm.com
Yunquan Zhang
zyq@mail.rdcps.ac.cn

¹ IBM Thomas J. Watson Research Center, Yorktown Heights, NY, USA

² Laboratory of Parallel Software and Computational Science, Chinese Academy of Sciences, Beijing, China

Sun. This paper proposes a HyperFatTree structure which uses low radix switches to construct scalable large interconnect for HPC systems. The main idea of HyperFatTree is to have full bisection bandwidth sub-network, and complete graph topology to interconnect sub-networks.

An Opportunistic Network Coding Routing for Opportunistic Networks, Jiansheng Yao, Chunguang Ma, Gang Du and Qi Yuan. This paper proposes opportunistic network coding (ONC) for OppNets by exploiting communication between neighboring nodes and between meeting nodes. The paper first develops the ONC model under some assumptions (which are hard to justify in real settings) and then develops an ONC routing protocol based on the model.

Parallel Algorithms for Generating Random Networks with Given Degree Sequences, Maksudul Alam and Maleq Khan. This paper provides novel parallel algorithms for generating random networks with given degree sequences, based on the Chung-Lu model. The algorithm uses the CL graph model and is implemented for a distributed memory system.

2 Applications: Scalability, Virtualization and Security

CovertInspector: Identification of Shared Memory Covert Timing Channel in Multi-tenanted Cloud, Sheng Wang, Weizhong Qiang, Hai Jin and Jinfeng Yuan. This paper presents a novel scheme to identify Shared Memory Covert Timing Channel (SMCTC) by intercepting high-resolution timer instructions and guest VM page faults. The system can fully detect malicious covert channel processes and prevent this attack while incurring acceptable performance overhead.

Vshadow: Promoting Physical Servers into Virtualization World, Song Wu, Yongchang Li, Xinhou Wang, Hai Jin and Hanhua Chen. This paper proposes a new Physical servers To Virtual machines (P2V) conversions system called Vshadow, which includes a local physical-to-virtual conversion, a VM image replication, and a VM live migration.

DMR: A Deterministic MapReduce for Multicore Systems, Yu Zhang and Huifang Cao. This paper proposes a deterministic MapReduce library, called deterministic-MapReduce (DMR) which is based on deterministic message passing multi-threaded model (DetMP).

3 Multicores and Accelerators

RECU: Rochester Elastic Cache Utility—Unequal Cache Sharing is Good Economics, Chencheng Ye, Jacob Brock, Chen Ding and Hai Jin. This paper studies the trade-off between fairness and overall performance for computing power renting and proposes two elastic cache partition strategies.

An Optimal Page-Level Power Management Strategy in PCM-DRAM Hybrid Memory, Jinbao Zhang, Xiaofei Liao, Hai Jin, Dong Liu, Li Lin and Kao Zhao. This paper proposes a new power management method for hybrid memory which adaptively places PCM and DRAM pages according to the pages' local and global access information.

Instruction Fusion for Multiscalar and Many-Core Processors, Yaojie Lu and Sotirios Ziavras. This paper proposes an instruction fusion technique for multiscalar or many-core processors to alleviate the utilisation wall impact.

Two-level Task Scheduling for Irregular Applications in GPU Platform, Jing Li, Lei Liu, Yuan Wu, Xiaobing Feng and Chengyong Wu. This paper proposes a two-level scheduling to distribute irregular tasks which also considers resource sharing for better cache utilization.

4 Parallelism

Hierarchical Read-Write Optimizations for Scientific Applications with Multi-variable Structured Datasets, Preeti Malakar and Venkatram Vishwanath. This paper proposes HieRO, which leverages the fast interconnect and performs hierarchical optimizations for reads and writes in scientific applications with structured datasets.

Performance Evaluation and Enhancement of Process-Based Parallel Loop Executions, Xingjing Lu, Long Chen and Zhiyuan Li. This paper proposes the PPLE parallel execution model that will enable parallelism for C/C++ non-numerical applications. It naturally privatizes pointer-dereferenced variables, and it can be used to efficiently exploit DOACROSS parallelism.

Determinism at Standard Library Level in Transactional Memory Based Applications, Vesna Smiljkovic, Osman Unsal, Adrian Cristal and Mateo Valero. This paper presents DeTrans-lib, the first standard C library that provides deterministic execution of TM-based applications at application and standard-library level. DeTrans-lib avoids deadlocks by performing transaction serialization in deterministic order.

5 Conclusions

All of the above papers represent original research in network and parallel computing systems and propose either new microarchitecture solutions or application models to exploit such systems. We hope that this special issue serves as a landmark source for education, information, and reference to professors, researchers and graduate students working in the area of networks and parallel computing.

The guest editors would like to express their gratitude to Prof. Alex Nicolau, the Editor-in-Chief of the International Journal on Parallel Programming, for giving us a chance to host this Special Issue. In addition, our sincere thanks to the program committee of NPC 2015 whose dedication and hard work resulted in such a high-quality conference. Last but not least, we are grateful to the authors for choosing to publish their research at this venue without which this special issue would not have been possible. We hope that this special issue will be a strong addition to the area of next generation Network and Parallel Computing.