

# Network and Parallel Computing

Ching-Hsien Hsu · Valentina Salapura

Received: 17 November 2014 / Accepted: 17 November 2014 / Published online: 13 December 2014  
© Springer Science+Business Media New York 2014

## 1 Introduction

It is our great 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, all papers were presented at the 11th IFIP International Conference on Network and Parallel Computing (NPC 2014) held from September 18 to September 20, 2014 in Yilan, Taiwan. Since 2003, NPC has become a valuable venue for scientists and engineers to present their excellent ideas and experiences in the field of distributed and parallel computing. The NPC 2014 conference continues this tradition and in particular extends the areas of interest to big data computing and parallel and multicore issues and opportunities. Chung Hua University, Taiwan organized this year's NPC conference.

The eleven papers included in this special issue were selected after two rounds of reviews from a total of 196 submissions. In the first review round, each submission received at least three reviews from the NPC 2014 Program Committee and additional external reviewers. We selected candidate papers for publication in this special issue from the very top ranked papers of the 196 NPC submissions. Authors of papers selected in the first round were then asked to revise their papers based on the reviewer feedback and submit a revised and extended manuscript for a second review round. Of the original 196 submissions, only the highest ranked eleven papers were selected

---

C.-H. Hsu (✉)  
Department of Computer Science and Information Engineering,  
Chung Hua University, Hsinchu, Taiwan  
e-mail: chh@chu.edu.tw

V. Salapura  
IBM T.J. Watson Research Center, Yorktown Heights, NY, USA  
e-mail: salapura@us.ibm.com

for inclusion in this special issue, representing a highly selective acceptance rate of 5.6%. The guest editors regret that we could not include more of the many high quality papers submitted to NPC this year.

The eleven papers appearing in this special issue cover a broad range of topics ranging in network and parallel computing from parallel computing theory to large-scale parallel computing applications, including virtualization, HPC, context-aware technologies, resource management, big data, semantic web services, GPU technologies, green computing, scheduling and load balancing, etc. All of these papers not only provide novel ideas and state-of-the-art techniques in the field, but will undoubtedly also stimulate future research in the converging cloud computing and communication technologies.

## 2 Parallel Architectures and Networking

The paper “Efficient Buffer Management for Tree Indexes On Solid State Drives” by Chengcheng Yang, Peiquan Jin, Lihua Yue and Puyuan Yang presents a new buffering scheme for tree indexes on SSDs. It not only considers the asymmetric read/write speeds of SSDs by first evicting cold clean pages, but also exploits the internal parallelism of SSDs by clustering dirty pages and then writing them to a new address with one coarse-grained I/O request. The reported experimental results demonstrate the benefits of the proposed techniques.

The paper “Using Packet Processing Object Modules Interchangeably as Stand-Alone Programs or Multi-App Components” by Ralph Duncan, Peder Jungck, Kenneth Ross, Dwight Mulcahy and Minh Nguyen describes a novel approach for deploying packet programs in object module format and deploying them as either a stand-alone program or as one of a collection of cooperating modules. The proposed framework provides a way to place multiple application modules in a single program image, which speeds up execution by eliminating the need to pass packets and metadata from application to application. The reported results show that working with packet processing applications in object and executable form is promising.

With the advance of network technology and the explosive growth of Internet applications, games have developed from single-player games to Massively Multiplayer Online Games. The paper “Operating System Enhancement for Supporting Massively Multiplayer Online Games in a Server Cluster” by Mei-Ling Chiang, Bo-Wen Yu, Chi-Shian Shia and Jiun-Jiun Hwang develops a new mechanism to transparently hand off a client’s live game connection between servers. The proposed mechanisms have been implemented in the Linux kernel. The reported experiments using a real multiplayer online game show that the proposed mechanisms can effectively reduce the need to hand off the on-going game connection and access the shared game database.

In multi-program environment, cache contention among processors can significantly degrade system performance. The paper “A Phase Behavior Aware Dynamic Cache Partitioning Scheme for CMPs” by Xiaofei Liao, Rentong Guo, Danping Yu, Hai Jin and Li Lin presents a novel dynamic cache partitioning mechanism based on the phase behavior of programs. The method proposed in this work is validated with measured results for applications from SPEC CPU2006 benchmark suite. They report

that their method can achieve a speedup up to 1.214 for co-scheduled applications compared to the performance of a shared cache scheme.

### 3 Parallel Programming and Runtime Systems

The paper “GPU Accelerated Finding of Channels and Tunnels for a Protein Molecule” by Byungjoo Kim, Jung Eun Lee, Young J. Kim and Ku-Jin Kim presents a novel method for finding all cavities, channels, and tunnels in a protein molecule in interactive time. A real-time performance is achieved for proximity queries on a protein molecule. The proposed system also provides a method for approximating a convex hull of a molecule in a discrete space. The reported experimental results in comparison with previous methods confirm the time efficiency of the proposed system.

The paper “A credit-based load-balance-aware CTA scheduling optimization scheme in GPGPU” by Yulong Yu, Xubin He, HeGuo, Yuxin Wang and Xin Chen investigates the load balance problem in the current GPGPU schedulers. The authors propose a credit-based load-balance-aware scheduling optimization scheme. The reported speedup over the baseline scheduler achieves up to 26.6% in reported benchmarks.

Server sprawl is a problem faced by data centers, which causes unnecessary waste of hardware resources, collateral costs of space, power and cooling systems and administration. The paper “SOC: Satisfaction-oriented Virtual Machine Consolidation in Enterprise Data Centers” by Xi Li, Anthony Ventresque, John Murphy and James Thorburn proposes a satisfaction-oriented virtual machine consolidation mechanism while taking IT managers’ preferences into consideration. They report that their proposed technique is effective in minimizing the number of hosts required to support a certain set of VMs, while maximizing the satisfaction degree of both managers from the provider and requester side.

### 4 Distributed Systems and Applications

Most essential services in large scale networked distributed systems (ad hoc, wireless or sensor) involve maintaining a global predicate over the entire network. The paper “A Linear Time Self-Stabilizing Algorithm for Minimal Weakly Connected Dominating Sets” by Yihua Ding, James Z. Wang and Pradip K. Srimani propose a new self-stabilizing algorithm for minimal weakly connected dominating sets. They report a significant improvement over the best possible solution in the existing literature.

Online social networks have become people’s daily platform for communication and sense of community. The retweet function in Twitter and the forward function in Sina Weibo are typical forwarding-based behaviors. The paper “Detection of Forwarding-based Malicious URLs in Online Social Networks” by Jian Cao, Qiang Li, Yuede Ji, Yukun He and Dong Guo focuses on forwarding-based behaviors to train a detection model for malicious URLs. These behaviors are naturally produced by the forwarding operation. They report that their evaluation shows that the proposed method has high accuracy and low false positive rates.

The paper “Effectiveness of Statistical Features for Early Stage Internet Traffic Identification” by Lizhi Peng, Bo Yang, Yuehui Chen and Zhenxiang Chen studies

the effectiveness of the statistical features for traffic identification, as accurate traffic classification is crucial for network management including deploying QoS-aware mechanisms, bandwidth budget management, intrusion detection, etc.. The reported experimental results demonstrate that for early stage traffic identification, most statistical features are as effective as the payload size.

The paper “A Text Clustering Approach of Chinese News Based on Neural Network Language Model” by Zhaoxin Fan, Shuoying Chen, Li Zha and Jiadong Yang applies a neural network language model to calculate the similarity of Chinese news based on text and word clustering. They report work has been successfully used in the Sohu news recommendation system and has great significance for research in text clustering and data mining of Chinese text.

## 5 Conclusions

All of the above papers address either original research in network and parallel computing, cloud computing and big data, or propose novel application models in the various parallel and distributed computing fields. They will also inspire further research and technology improvements in application of parallel computing and cloud services. This special issue serves as a resource for education, information, and reference to professors, researchers and graduate students in network and parallel computing, cloud computing and management, and novel application models for distributed systems.

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 inviting us to edit this Special Issue as Guest Editors. In addition, we are deeply indebted to program committee of NPC 2014 for their dedication and expert work in reviewing these papers, and the NPC Steering Committee for their guidance. Last but not least, we are grateful to all authors for their contributions, without which this special issue could not have been produced. We hope that you will enjoy this special issue as much as we enjoyed preparing it.