Network and Parallel Computing

Ching-Hsien Hsu · Xiaoming Li · Xuanhua Shi

Received: 21 March 2014 / Accepted: 21 March 2014 / Published online: 20 April 2014
© Springer Science+Business Media New York 2014

1 Introduction

It's our pleasure to bring you this special issue of the International Journal of Parallel Programming on the network and parallel computing. Prior to the publication to this special issue, all papers were presented in the 10th IFIP International Conference on Network and Parallel Computing (NPC' 13) held from September 19 to September 21 2013 in Guiyang, China. Since 2003, NPC has become a valuable venue for engineers and scientists to present their excellent ideas and experiences in system?fields of distributed and parallel computing. The NPC 2013 continues this tradition and in particular extends areas of interest to Big data computing and Parallel and multicore issues and opportunities. Huazhong University of Science and Technology, China organized this year's NPC conference.

The twelve papers included in this special issue were selected after two rounds of selections from in total 109 submissions. In the first round of reviewing, each submission received at least 3 reviews by the guest editors in consultation with the NPC 2013 Program Committee. The PC also sought additional external reviews for

C.-H. Hsu (✉)

Department of Computer Science and Information Engineering, Chung Hua University,
Hsinchu, Taiwan
e-mail: chh@chu.edu.tw

X. Li

Department of Electrical and Computer Engineering, University of Delaware, Newark, DE, USA
e-mail: xli@udel.edu

X. Shi

School of Computer Science and Technology, Huazhong University of Science and Technology,
Wuhan, China
e-mail: xhshi@hust.edu.cn

contentious papers. Authors whose papers were selected in the first round were then asked to revise their papers based on the reviewer feedback and present an extended version for the second round of reviewing. The overall selection rate is 11 %. The guest editors regret that many high quality papers can not be included in this special issue, and are confident the papers presented here demonstrate idea novelty and sound technical depth. The twelve papers cover a variety of topics ranging from parallel computing theories 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 also stimulate future research in merging cloud and communication services.

2 Heterogeneous Architectures and Computing

The paper entitled “YuruBackup: A Space-Efficient and Highly Scalable Incremental Backup System in the Cloud” [1] by Quanqing Xu, Liang Zhao, Mingzhong Xiao, Anna Liu, and Yafei Dai presents YuruBackup, a space-efficient and highly scalable incremental backup system in the Cloud. YuruBackup enables fine-grained data de-duplication with hierarchical partitioning to improve space efficiency to reduce bandwidth of both backup and restore processes, and storage costs. On the other hand, YuruBackup explores a highly scalable architecture for fingerprint servers that allow adding one or more fingerprint servers dynamically to cope with increasing clients.

The paper entitled “Automatic composition of heterogeneous models based on Semantic Web services” [2] by Hui Huang, Ligang He, Xueguang Chen, Minghui Yu, and Zhiwu Wang studies the problem of heterogeneous model composition by employing the techniques based on semantic web services and Artificial Intelligence (AI) planning. In their paper, the heterogeneous model composition problem is converted to the problem of planning in nondeterministic domains under partial observability. An automatic composition method is presented to generate the composite model based on the planning as model checking technique. The experiment results show the feasibility and capability of their approach in dealing with the complex problems involving heterogeneous models.

The paper entitled “Accelerating Smith-Waterman Alignment of Species-based Protein Sequences on GPU” [3] by Xiaowen Feng, Hai Jin, Ran Zheng, Lei Zhu, and Weiqi Dai develops an efficient implementation of the Smith-Waterman algorithm, which is an optimal method of finding the local sequence alignment. The algorithm requires a large amount of computation and memory space, and is also constrained by the memory access speed of the Graphics Processing Units (GPUs) global memory when accelerating by using GPUs. Their new implementation improves performance by optimizing the organization of database, increasing GPU threads for every database sequence, and reducing the number of memory accesses to alleviate memory bandwidth bottleneck. Experimental results show performance improved by about 32 % on average when compared with CUDASW++2.0 and DOPA with Ssearch trace for 100 shortlisted sequences on NVIDIA GTX295. It also outperforms CUDASW++2.0 with Ssearch trace for 100 shortlisted sequences by about 52 % on NVIDIA GTX460.

The paper entitled “Power Efficiency for Hardware/Software Partitioning with Time and Area Constraints on MPSoC” [4] by Edwin Sha, Li Wang, Qingfeng Zhuge, Jun Zhang, and Jing Liu proposes two algorithms for hardware/software partitioning problem on MPSoC, to minimize power consumption with time and area constraints. The first Tree_Partitioning algorithm generates optimal partitioning results for tree-structured control-flow graphs using dynamic programming. For the general partitioning problem, they propose the DAG_Partitioning algorithm to produce near optimal solution efficiently for directed-acyclic graphs (DAGs).

3 Virtualization and Resource Management

The paper entitled “CCAP: A Cache Contention-Aware Virtual Machine Placement Approach for HPC Cloud” [5] by Hai Jin, Hanfeng Qin, Song Wu, and Xuerong Guo studied the resource consumption characteristics in High Performance Computing (HPC). In particular, they investigate performance interference due to contention for SLLC in the HPC cloud, and propose an enhanced reuse distance analysis technique with an accelerated cyclic compression algorithm to identify application’s cache interference intensity. For a 2-workload system, the execution time is reduced by 12 %, as well as cache miss rate improved by 13 %.

The paper entitled “Efficiently Restoring Virtual Machines” [6] by Bernhard Egger, Erik Gustafsson, Changyeon Jo, and Jeongseok Son introduces a method to efficiently restore VMs from space-optimized checkpoint images in the paper “Efficiently Restoring Virtual Machines”. With the new method, a VM is available to the user before the entire memory contents of the VM have been restored. The new technique is able to reduce the time-to-responsiveness (TTR) for restored VMs to a few seconds and reduces the TTR by 50 % compared to the Xen hypervisor. Compared to the previously fasted restoration of space-optimized checkpoints, the proposed technique achieves a threefold speedup on average.

The paper entitled “A Parallel Job Execution Time Estimation Mechanism Based on User Submission Patterns within Computational Grids” [7] by Feng Liang, Yunzhen Liu, Hai Liu, Shilong Ma, and Bettina Schnor presents and evaluate a novel execution time estimation approach for parallel jobs, the User-Behavior Clustering for Execution Time Estimation (UBCETE), which can give more accurate execution time estimation for parallel jobs through exploring the job similarity and revealing the user submission patterns to help finding the similar jobs. Compared to the state-of-art algorithms, their approach is shown to improve the accuracy of the job execution time estimation up to 5.6 %, meanwhile the time our approach spends on calculation can be reduced up to 3.8 %.

The paper entitled “A Virtualization Based Monitoring System for Mini-intrusive Live Forensics” [8] by Xianming Zhong, Chengcheng Xiang, Miao Yu, Zhengwei Qi, and Haibing Guan proposed VAIL, a novel virtualization based monitoring system for mini-intrusive live forensics, which employs hardware assisted virtualization technique to gather integrated information for the native computer system. Meanwhile, the execution of the target system will not be interrupted and VAIL keeps immune to attacks from the target system. Their experimental results show that VAIL can obtain

comprehensive digital evidences from the target system as designed, including the CPU state, the physical memory content, and the I/O activities. And in average, VAIL only introduces 4.21 % performance overhead to the target system, which proves that VAIL is practical in real commercial environments.

4 Big Data Intelligence and Applications

The paper entitled “OFScheduler: A Dynamic Network Optimizer for MapReduce in Heterogeneous Cluster” [9] by Zhao Li, Yao Shen, Bin Yao, and Minyi Guo proposed a new dynamic network optimizer called OFScheduler for heterogeneous clusters to relieve the network traffic during the execution of MapReduce jobs. The proposed optimizer tags different types of traffic and utilize the Openflow to adjust transfers of flows dynamically. When evaluated in a simulator and an OpenFlow testbed, the proposed optimizer has a significant effect on increasing bandwidth utilization and improving the performance of MapReduce by 24–63 % for most of jobs in a multi-path heterogeneous cluster.

In the paper entitled “An Adaptive Sampling Mechanism for Partitioning in MapReduce” [10], Kenn Slagter, Ching-Hsien Hsu, and Yeh-Ching Chung propose an adaptive sampling mechanism for total order partitioning that can reduce memory consumption whilst partitioning with a trie-based sampling mechanism (ATrie). The performance of the proposed algorithm is compared to a state of the art trie-based partitioning system (ETrie). and show the proposed mechanism is more adaptive and more memory efficient than previous implementations while uses 2.43 times less memory for case insensitive email addresses, and uses 1024 times less memory for birthdates compared to that of a 2-level ETrie.

In the paper “Data Reduction Analysis for Climate Data Sets” [11], Songbin Liu, Xiaomeng Huang, Haohuan Fu, Guangwen Yang, and Zhenya Song study the potential benefit of data reduction for climate data by investigating a total of 46.5 TB climate data sets, including 3 observation data sets (14.1 TB) and 3 climate model output data sets (32.4 TB). Five different data compression algorithms and two types of content deduplication mechanisms are applied to these data sets to study the possible data reduction effectiveness. Further more, the compressibility of different climate component data is also examined. They find that the compression method LCFP can provide the best compression ratio; however, its throughputs, especially the inflate throughputs are much lower than all the others.

The paper entitled “Inaccuracy in Private BitTorrent Measurements” [12] by Hai Jin, Honglei Jiang, Shadi Ibrahim, and Xiaofei Liao investigates the accuracy of previous measurement studies on PT sites, while emphasizing the incentive rules employed and the interplay between these rules and corresponding objective factors. They evaluate the behavior regulation policies of the front-end website and the tracker and examine the semantics of provided data, and design a new crawling methodology and conducted a large-scale measurement study across four representative PT sites over a year. Their study offers fundamental insights into designing an accurate methodology when conducting measurement studies on PT sites.

5 Conclusions

All of the above papers address either original research in network and parallel computing or cloud and big data, or propose novel application models in the various parallel and distributed computing fields. They also trigger further related research and technology improvements in application of parallel computing and cloud services. This special issue serves as a landmark source for education, information, and reference to professors, researchers and graduate students interested in updating their knowledge about or active in network/parallel computing, cloud computing and management, and novel application models for distributed systems.

The guest editors would like to express sincere 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, we are deeply indebted to program committee of NPC 2013 for their dedication and expert work in reviewing and shepherd these papers, the NPC Steering Committee for their guidance. Last but not least, we are grateful to all authors for their contributions and for undertaking two-cycle revisions of their manuscripts, without which this special issue could not have been produced. We hope that this special issue will be a good addition to the area of next generation Network and Parallel Computing.

References

1. Xu, Q., Zhao, L., Xiao, M., Liu, A., Dai, Y.: YuruBackup: a space-efficient and highly scalable incremental backup system in the cloud. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0280-7](https://doi.org/10.1007/s10766-013-0280-7)
2. Huang, H., He, L., Chen, X., Yu, M., Wang, Z.: Automatic composition of heterogeneous models based on semantic Web services. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0279-0](https://doi.org/10.1007/s10766-013-0279-0)
3. Feng, X., Jin, H., Zheng, R., Zhu, L., Dai, W.: Accelerating smith-waterman alignment of species-based protein sequences on GPU. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0284-3](https://doi.org/10.1007/s10766-013-0284-3)
4. Sha, E., Wang, L., Zhuge, Q., Zhang, J., Liu, J.: Power efficiency for hardware/software partitioning with time and area constraints on MPSoC. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0283-4](https://doi.org/10.1007/s10766-013-0283-4)
5. Jin, H., Qin, H., Wu, S., Guo, X.: CCAP: a cache contention—aware virtual machine placement approach for HPC cloud., *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0286-1](https://doi.org/10.1007/s10766-013-0286-1)
6. Egger, B., Gustafsson, E., Jo, C., Son, J.: Efficiently restoring virtual machines. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0295-0](https://doi.org/10.1007/s10766-013-0295-0)
7. Liang, F., Liu, Y., Liu, H., Ma, S., Schnor, B.: A parallel job execution time estimation approach based on user submission patterns within computational grids. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0294-1](https://doi.org/10.1007/s10766-013-0294-1)
8. Zhong, X., Xiang, C., Yu, M., Qi, Z., Guan, H.: A virtualization based monitoring system for mini-intrusive live forensics. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0285-2](https://doi.org/10.1007/s10766-013-0285-2)
9. Li, Z., Shen, Y., Yao, B., Guo, M.: OFScheduler: a dynamic network optimizer for MapReduce in heterogeneous cluster. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0281-6](https://doi.org/10.1007/s10766-013-0281-6)
10. Slagter, K., Hsu, C.-H., Chung, Y.-C.: An adaptive and memory efficient sampling mechanism for partitioning in MapReduce. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0288-z](https://doi.org/10.1007/s10766-013-0288-z)
11. Liu, S., Huang, X., Fu, H., Yang, G., Song, Z.: Data reduction analysis for climate data sets. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0287-0](https://doi.org/10.1007/s10766-013-0287-0)
12. Jin, H., Jiang, H., Ibrahim, S., Liao, X.: Inaccuracy in private BitTorrent measurements. *Int. J. Parallel Program.* doi:[10.1007/s10766-013-0282-5](https://doi.org/10.1007/s10766-013-0282-5)