



Editorial

A note on new trends in data-aware scheduling and resource provisioning in modern HPC systems



The Big Data era [1,2] poses new challenges as well as significant opportunities for High-Performance Computing (HPC) systems such as how to efficiently turn massively large data into valuable information and meaningful knowledge? It is clear that computationally optimized new data-driven HPC techniques are required for processing Big Data in rapidly-increasing number of applications, such as Life Sciences, Particle Physics and Socio-economical systems.

The realm of HPC systems lies in sharing of the “multi-core” hardware resources among the software applications. Key characteristics of HPC systems include high processor density, high speed Input/Output (IO), and high-density cooling techniques. In the Pre Grid computing era (before 2000), the HPC was always exclusively referred to as “supercomputing”. In grid-based HPC era, the Globus project in conjunction with cluster job scheduling systems such as Portable Batch System (PBS), and Platform LSF (recently acquired by IBM) has dominated the middleware research efforts and powered numerous computational grids across the globe. PBS and Platform LSF implemented scheduling techniques and cluster resource provisioning mechanism for allocation of available compute resources to HPC applications.

However, with the emergence of modern HPC systems (e.g., Amazon EC2 Cluster CPU Instances, Univa Grid Engine, IBM HPC Cloud, etc.) powered by cloud computing and virtualisation technologies, the job scheduling techniques implemented by traditional HPC schedulers (e.g. Platform LSF, PBS) are facing serious limitations. The main reason for this state of affairs is that modern HPC systems *cannot* support on-demand scalability, strong performance guarantees, and improved fault-tolerance, which the traditional HPC scheduling techniques are unable to cater for or take advantage of. In reality, the traditional HPC schedulers were not designed for the cloud computing and virtualization era.

Hence, this special issue solicits papers related to topics including techniques for optimizing the performance of traditional HPC applications on new “multi-core” cloud systems, novel extensions to traditional HPC schedulers for Big Data application scheduling, dynamic resource provisioning for HPC applications on the cloud, techniques for optimizing HPC application specific performance and energy constraints, workflow scheduling techniques, and so on.

The Call for this Special Issue received a number of submissions. After a two-phase review process we accepted eight papers of very high quality. This includes one research survey paper that gives an interesting overview of the current big data problems and solutions. The other papers are related to scheduling strategies on Cloud-related platforms with individual scheduling goals.

The paper “Remote Sensing Big Data Computing: Challenges and Opportunities” [3] gives a comprehensive overview on the big data and data-intensive problems, including the analysis of remote sensing (RS) big data, challenges, current techniques, and existing works for processing big data. The complexity of RS data is clearly analyzed and the research issues are fully elaborated with efficient managing of the massive RS data, the intensive irregular data access patterns, loading and transmission, optimal scheduling for data-dependent tasks, as well as efficient and productive programming. The paper dives into the challenges and opportunities for the state-of-the-art parallel systems.

The paper “Resource-Aware Hybrid Scheduling Algorithm in Heterogeneous Distributed Computing” [4] presents a novel scheduling technique for tasks over distributed computing architectures. The main idea of the scheduling algorithm is to apply resource-aware techniques by applying hierarchical clustering of the available resources into allocation groups for both batch jobs and workflows, which are then scheduled with one of the classical algorithms based on the requirements. The technique fits well for multimedia workloads that are characterized by both compute intensive filtering tasks and communication intensive tasks, as well as for big data applications.

The paper “Artificial Neural Network Support to Monitoring of the Evolutionary Driven Security Aware Scheduling in Computational Distributed Environments” [5] proposes an artificial neural network based model of the learning monitoring Grid system for supporting the security awareness of evolutionary driven Grid schedulers. The approach considers both the efficiency and the safety of task execution processes, and also improves the reliability of computing resources.

The paper “Cloud-aware data intensive workflow scheduling on volunteer computing systems” [6] proposes a partitioning and data-centric approach to schedule workflows on both volunteer and Cloud resources with the goal to meet the deadline in the given user SLA. With this approach a workflow is partitioned into sub-workflows to minimize data dependencies and these sub-workflows are then distributed on volunteer resources according to the proximity of resources and the load balancing policy. Based on the prediction of the execution time of the sub-workflows their re-scheduling onto public Clouds may be organized to improve the system performance by increasing the percentage of workflows that meet the time deadline.

The paper “CloudFlow: Data-aware Programming Model for Cloud Workflow Applications on Modern HPC System” [7] describes a new model for the efficient use of shared-data on

MapReduce frameworks and to conduct both task-level and job-level scheduling. The purpose of developing this model is to address complex data flow in multiple Cloud applications and to improve the overall performance of MapReduce jobs via proper scheduling. The novel CloudFlow Programming Model supports user-defined Map/Reduce Functions, which allows the user to orchestrate their required data-flow logics. Both theoretical proof and experimental evaluations demonstrate the effectiveness of the proposed scheduling framework.

The paper “ParSA: High-throughput Scientific-data Analysis Framework with Distributed File System” [8] presents a software framework (ParSA) for high performance analysis on scientific data. ParSA introduces application-aware data distribution in HDFS and locality-aware task scheduling to accelerate well-known analysis methods for climate data. Experimental results demonstrate the high efficiency and scalability of the framework.

The paper “Workload Balancing and Adaptive Resource Management for the Swift Storage System on Cloud” [9] proposes a framework of workload balancing and resource management for Swift with a number of algorithms to achieve a workload aware balancing solution. Based on an experiment that demonstrates the weakness of the default load balancing mechanisms of Swift, several algorithms are proposed to optimize the overloaded and under-loaded nodes in the cluster of Swift. The proposed mechanism is lightweight and requires no source code change in the guest OS and storage applications.

The paper “A high performance framework for modeling and simulation of large-scale complex systems” [10] proposes an efficient computational component model framework, a vectorial component model (VCM), to fully exploit the parallelism of multi-input component models. A two-level composite parallel scheme (CMP) is designed for VCM, which well demonstrates how to use VCM to tackle research challenges in anti-missile of naval vessel parallel simulations.

We thank all reviewers for providing valuable review comments and for their substantial contributions to the final revised versions appearing in this special issue. We would like also to thank Prof. Peter Sloot, Editor-in-Chief, for his great help and productive advice in preparing this special issue.

References

- [1] L. Wang, R. Ranjan, J. Kolodziej, A. Zomaya, L. Alem, Software tools and techniques for big data computing in healthcare clouds, *Future Gener. Comput. Syst. J.* 43–44 (2015) 38–39.
- [2] Ching-H Hsu, Special section: Intelligent big data processing, *Future Gener. Comput. Syst. J.* 36 (2014) 16–18.
- [3] Yan Ma, et al., Remote sensing big data computing: Challenges and opportunities, *Future Gener. Comput. Syst.* 51 (2015) 47–60.
- [4] Mihaela-Andreea Vasile, et al., Resource-aware hybrid scheduling algorithm in heterogeneous distributed computing, *Future Gener. Comput. Syst.* 51 (2015) 61–71.
- [5] Grzonka Daniel, et al., Artificial Neural Network support to monitoring of the evolutionary driven security aware scheduling in computational distributed environments, *Future Gener. Comput. Syst.* 51 (2015) 72–86.
- [6] Toktam Ghafarian, et al., Cloud-aware data intensive workflow scheduling on volunteer computing systems, *Future Gener. Comput. Syst.* 51 (2015) 87–97.
- [7] Fan Zhang, et al., CloudFlow: A data-aware programming model for cloud workflow applications on modern HPC systems, *Future Gener. Comput. Syst.* 51 (2015) 98–110.
- [8] Tao Zhang, et al., ParSA: High-throughput scientific data analysis framework with distributed file system, *Future Gener. Comput. Syst.* 51 (2015) 111–119.
- [9] Zhenhua Wang, et al., Workload balancing and adaptive resource management for the swift storage system on cloud, *Future Gener. Comput. Syst.* 51 (2015) 120–131.
- [10] Feng Zhu, et al., A high performance framework for modeling and simulation of large-scale complex systems, *Future Gener. Comput. Syst.* 51 (2015) 132–141.

Guest Editors

Dr. Jie Tao

*Steinbuch Centre for Computing, Karlsruhe Institute of Technology,
Karlsruhe, Germany*

E-mail address: jie.tao@kit.edu.

Prof. Joanna Kolodziej

*Institute of Computer Science for Research and Development,
Cracow University of Technology, Cracow, Poland*

E-mail address: jokoldziej@pk.edu.pl.

Dr. Rajiv Ranjan*

Dr. Prem Prakash Jayaraman

*Commonwealth Scientific and Industrial Research Organisation,
Canberra, Australia*

E-mail addresses: raj.ranjan@csiro.au (R. Ranjan),

Prem.jayaraman@csiro.au (P.P. Jayaraman).

Prof. Rajkumar Buyya

*CEO, Manjrasoft Pty Ltd, Melbourne, Australia
Director, Cloud Computing and Distributed Systems Laboratory,*

Department of Computer Science and Software Engineering,

The University of Melbourne, Australia

E-mail address: rbuyya@unimelb.edu.au.

16 April 2015

Available online 8 May 2015

* Corresponding editor.