



Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Editorial

Guest editors' introduction: Special issue on Cluster, Grid, and Cloud Computing

Rapid advances in architectures, networks, and systems and middleware technologies are leading to new concepts in and platforms for computing, ranging from Clusters and Grids to Clouds and Datacenters. IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid) is a series of very successful conferences, with the overarching goal of bringing together international researchers, developers, and users to provide an international forum to present leading research activities and results on a broad range of topics related to these concepts and platforms, and their applications.

In this issue:

This special issue is a consolidation of the top papers selected from CCGrid 2012, the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing held in Ottawa, Canada during May 13–16, 2012. Authors of the top 13 papers from CCGrid were invited to extend and submit to this special issue. All extended papers were assessed by 3–4 peer reviewers as per the standard of the FGCS journal and 6 papers were accepted, based on the first round of reviews, to ensure that this special issue is brought out in a timely manner. A summary of accepted papers and their contributions is briefly noted below:

In “Self-healing of workflow activity incidents on distributed computing infrastructures”, da Silva et al. present an approach to classify operational incidents in scientific gateways used for distributed computing infrastructure, and use this information to heal the underlying issue when possible. The authors study the long-tail effect issue, and propose a new algorithm to control task replication. The healing process is parameterized on real application traces acquired in production on the European Grid Infrastructure. The paper also presents experimental results obtained on the Virtual Imaging Platform, and demonstrates that the proposed method speeds up execution up to a factor of 4, consumes up to 26% lesser resource time than a control execution.

In “Efficient Distributed Monitoring with Active Collaborative Prediction”, Feng et al. argue that the problem of probe selection for fault prediction based on end-to-end probing is essentially a collaborative prediction (CP) problem. Using an extensive experimental dataset from the EGI grid, the authors demonstrate that the combination of applying the Maximum Margin Matrix Factorization approach to CP and Active Learning shows excellent performance, reducing the number of probes typically by 80% to 90%. The authors also compare their approach with other collaborative prediction strategies and show that Active Probing is most efficient at dealing with the various sources of data variability.

In “Scheduling Concurrent Applications on a Cluster of CPU–GPU Nodes”, Ravi et al. study the problem of optimizing the overall throughput of a set of applications deployed on a cluster of heterogeneous nodes comprised of multi-core CPUs and many-core GPUs.

The authors formulate the issue of scheduling on such systems in two different ways. In the first formulation, jobs can be executed on either the GPU or the CPU of a single node. In the second, jobs can be executed on the CPU, GPU, or both, of any number of nodes in the system. Through experimentation, the authors demonstrate that the schemes proposed for first formulation outperform a blind round-robin scheduler and approximate the performances of an ideal scheduler that involves an impractical exhaustive exploration of all possible schedules. They also demonstrate that the scheme proposed for the second formulation outperforms the best of existing schemes for heterogeneous clusters, TORQUE and MCT, by up to 42%. Furthermore, they also evaluate the robustness of our proposed scheduling policies under inaccurate inputs to account for real execution scenarios and show that, with up to 20% of inaccuracy in the input, the degradation in performance is marginal (less than 7%) on the average.

In “User Transparent Data and Task Parallel Multimedia Computing with Pyxis-DT”, Kessel et al. propose an easy-to-use programming model that brings the benefits of HPC to multimedia content analysis (MMCA). Specifically, they argue that existing user transparent parallelization tools generally use a data parallel approach in which data structures (e.g., video frames) are scattered among the available nodes in a compute cluster. For certain MMCA applications a data parallel approach induces intensive communication, however, which significantly decreases performance. In these situations, applying alternative approaches can be beneficial. To help with such situations, the authors propose Pyxis-DT: a user transparent parallel programming model for MMCA applications that employs both data and task parallelism. Through experimentation, the authors demonstrate that for realistic MMCA applications the concurrent use of data and task parallelism can significantly improve performance compared to using either approach in isolation. Extensions for GPU clusters are also presented.

In “Using Model Checking to Analyze the System Behavior of the LHC Production Grid”, Remenska et al. present an analysis of DIRAC (Distributed Infrastructure with Remote Agent Control) with mCRL2, process algebra with data. The authors describe their efforts to reverse engineer two critical and related DIRAC subsystems, and subsequently model their behavior with the mCRL2 toolset. The authors describe how this enabled them to easily locate race conditions and livelocks which were confirmed to occur in the real system. The authors further formalize and verify several behavioral properties of the two modeled subsystems.

In “Flubber: Two-level Disk Scheduling in Virtualized Environment”, Jin et al. present the design of Flubber, a two-level scheduling framework that decouples throughput and latency allocation to provide QoS guarantees to virtual machines (VMs) while maintaining high disk utilization. The high-level throughput control regulates the pending requests from the VMs with an adaptive

credit-rate controller, in order to meet the throughput requirements of different VMs and ensure performance isolation. Meanwhile, the low-level latency control, by the virtue of the batch and delay earliest deadline first mechanism (BD-EDF), re-orders all pending requests from VMs based on their deadlines, and batches them to disk devices taking into account the locality of accesses across VMs. The authors also discuss their implementation and extensive evaluation of Flubber on a Xen-based host and demonstrate that Flubber can simultaneously meet the different service requirements of VMs while improving the efficiency of the physical disk.

We hope the articles in this special issue will provide relevant insights into the emerging trends in Cluster, Cloud and Grid computing.



Pavan Balaji holds appointments as a Computer Scientist and Group Lead at the Argonne National Laboratory, as a Research Fellow of the Computation Institute at the University of Chicago, and as an Institute Fellow of the Northwestern-Argonne Institute of Science and Engineering at Northwestern University. His research interests include parallel programming models and runtime systems for communication and I/O, modern system architecture (multi-core, accelerators, complex memory subsystems, and high-speed networks), cloud computing systems, and job scheduling and resource management. He has nearly

100 publications in these areas and has delivered nearly 120 talks and tutorials at various conferences and research institutes. He is a recipient of the U.S. Department of Energy's Early Career Award. He has also received several other awards including the Director's Technical Achievement award at Los Alamos National Laboratory, an Outstanding Researcher award at the Ohio State University, five best paper awards and various others. He serves as the worldwide chairperson for the IEEE Technical Committee on Scalable Computing (TCSC). He has also served as a chair or editor for nearly 50 journals, conferences and workshops, and as a technical program committee member in numerous conferences and workshops. He is a senior member of the IEEE and a professional member of the ACM. More details about Dr. Balaji are available at <http://www.mcs.anl.gov/~balaji>. Contact him at balaji@mcs.anl.gov.



Rajkumar Buyya is Professor of Computer Science and Software Engineering, Future Fellow of the Australian Research Council, and Director of the Cloud Computing and Distributed Systems (CLOUDS) Laboratory at the University of Melbourne, Australia. He is also serving as the founding CEO of Manjrasoft, a spin-off company of the University, commercializing its innovations in Cloud Computing. He has authored over 425 publications and four text books including "Mastering Cloud Computing" published by McGraw Hill and Elsevier/Morgan Kaufmann, 2013 for Indian and international markets respectively. He

also edited several books including "Cloud Computing: Principles and Paradigms" (Wiley Press, USA, Feb 2011). He is one of the highly cited authors in computer science and software engineering worldwide (h-index = 67, g-index = 140, 21700+ citations). Software technologies for Grid and Cloud computing developed under Dr. Buyya's leadership have gained rapid acceptance and are in use at several academic institutions and commercial enterprises in 40 countries around the world. Dr. Buyya has led the establishment and development of key community activities, including serving as foundation Chair of the IEEE Technical Committee on Scalable Computing and five IEEE/ACM conferences. These contributions and international research leadership of Dr. Buyya are recognized through the award of "2009 IEEE Medal for Excellence in Scalable Computing" from the IEEE Computer Society, USA. Manjrasoft's Aneka Cloud technology developed under his leadership has received "2010 Asia Pacific Frost and Sullivan New Product Innovation Award" and "2011 Telstra Innovation Challenge, People's Choice Award". He is currently serving as the foundation Editor-in-Chief (EiC) of IEEE Transactions on Cloud Computing. For further information on Dr. Buyya, please visit his cyberhome: www.buyya.com.

Pavan Balaji*

Argonne National Laboratory, USA

E-mail address: pavanbalaji@gmail.com.

Rajkumar Buyya

University of Melbourne, Australia

15 April 2013

Available online 22 May 2013

* Corresponding editor. Tel.: +1 630 252 3017.