



Contents lists available at ScienceDirect

Future Generation Computer Systems

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

Editorial

Special section: Federated resource management in grid and cloud computing systems

Rajkumar Buyya^a, Rajiv Ranjan^{b,*}

^a Cloud Computing and Distributed Systems (CLOUDS) Laboratory, Department of Computer Science and Software Engineering, The University of Melbourne, Australia

^b Service Oriented Computing Research Group, School of Computer Science and Engineering, The University of New South Wales, Sydney, Australia

ARTICLE INFO

Article history:

Received 5 June 2010

Accepted 8 June 2010

Available online 15 June 2010

Welcome to the special issue of Future Generation Computer System (FGCS) journal. This special issue compiles a number of excellent technical contributions that significantly advance the state-of-the-art in federated management [1–3] of Grid and Cloud computing environments. Federated management of administratively distributed grids and clouds offers significant benefits including: (i) improving the ability of resource providers to meet SLA compliance [4] for clients and offer improved service by optimizing the service placement and throughput according to users' QoS needs; (ii) enhancing the peak-load handling and dynamic system expansion capacity of every Grid/Cloud domain through federation [5] without the need for setting up a new software or hardware infrastructure in every location, and (iii) adapting to failures including natural disasters and regular system maintenance more gracefully as providers can transparently migrate their services to other domains in the federation, thus avoiding SLA violations and resulting penalties. Hence, federated management not only ensures business continuity but also augments the reliability of the participating resource providers. Next, we briefly describe the technical contributions, which were selected for publication in this special issue. All of the selected papers underwent a rigorous peer-review process.

A major performance issue in large-scale decentralized distributed systems, such as grids, is how to ensure that jobs finish their execution within estimated completion times in the presence of resource performance fluctuations. Previously, several techniques including advance reservation, rescheduling and migration have been adopted to resolve/relieve this issue; however, there are some non-trivial practical hurdles that need to be addressed. The use of leveraging additional resources from clouds may be an attractive alternative, since resources in clouds are much more reliable than those in grids. In the paper "Rescheduling for Reliable Job Completion with the Support of Clouds", Young Choon Lee and

Albert Y. Zomaya investigate the effectiveness of rescheduling using cloud resources to increase the reliability of job completion. Specifically, schedules are initially generated using grid resources while cloud resources (relatively more expensive) are used only for rescheduling to deal with delays in job completion. A job in their study refers to a bag-of-tasks (BoT) application that consists of a large number of independent tasks; this job model is common in many science and engineering applications. They have devised a novel rescheduling technique, called rescheduling using clouds for reliable completion (RC2) and applied it to three well-known existing heuristics. Their experimental results reveal that RC2 significantly reduces delays in job completion.

In scientific Cloud workflows, large amounts of application data need to be stored in distributed data centers. To effectively store these data, a data manager must intelligently select data centers in which these data will reside. This is, however, not the case for data which must have a fixed location. When one task needs several datasets located in different data centers, the movement of large volumes of data becomes a challenge. In the paper "A Data Placement Strategy in Scientific Cloud Workflows", Dong Yuan, Yun Yang, Xiao Liu, and Jinjun Chen propose a matrix based k -means clustering strategy for data placement in scientific cloud workflows. The strategy contains two algorithms that group the existing datasets in k data centres during the workflow build-time stage, and dynamically clusters newly generated datasets to the most appropriate data centres based on dependencies during the runtime stage. Simulations show that their algorithm can effectively reduce data movement during workflow execution.

Shared logging is an increasingly important problem for accounting and usage tracking in overlapping and federated Grid environments. Large scientific projects such as the Large Hadron Collider have outgrown the capacity of any existing single Grid, making federations of Grids more and more common. In the paper titled "Distributed Usage Logging for Federated Grids", Erik Elmroth and Daniel Henriksson present a non-intrusive solution to the increasingly important problem of shared logging for overlapping and federated Grid environments. The solution addresses three usage scenarios of hierarchical Grids, mutual cross-Grid resource

* Corresponding author.

E-mail address: rranjans@gmail.com (R. Ranjan).

utilization, and federated Cloud computing infrastructures. The approach is evaluated by extending the existing SweGrid Accounting System (SGAS) with a light-weight component that makes the system applicable to a wide range of usage scenarios. The proposed architecture is characterized by its simplicity, flexibility, and generality, and the new key component by its non-intrusiveness, flexibility, and ability to manage high load. They also present requirements derived from three usage scenarios, and also include an in-depth description of the architecture and design, as well as the implementation and performance evaluation of a new component written for use with SGAS. They conclude from a performance evaluation that the sharing of usage data is not likely to be a limiting performance factor even in large-scale Grid scenarios.

As the utilization of Cloud platforms grows, users are realizing that the implicit promise of clouds (leveraging them from the tasks related with infrastructure management) is not fulfilled. A reason for this is that current clouds offer interfaces too close to that infrastructure, while users demand functionalities that automate the management of their services as a whole unit. To overcome this limitation, in the paper “From Infrastructure Delivery to Service Management in Clouds”, Luis Rodero-Merino et al. propose a new abstraction layer closer to the lifecycle of services that allows for their automatic deployment and escalation depending on the service status (not only on the infrastructure). This abstraction layer can sit on top of different cloud providers, hence mitigating the potential lock-in problem and allowing the transparent federation of clouds for the execution of services.

Francesco Palmieri and Silvio Pardi in their paper titled “Towards a Federated Metropolitan Area Grid Environment: the SCoPE Network-aware Infrastructure” discuss opportunities and challenges in the Grid resource management infrastructure and network control plane design, critical to the provisioning of network-assisted extensible Grid services on the metropolitan scale. Such services can empower a real high performance distributed computing system built on dark fibre transport networks, administered within a single domain and offering plenty of cheap bandwidth to e-science applications. This approach makes the transport infrastructure the main enabling factor of a novel Grid vision: the “Metropolitan Area Grid” (MAG) aiming at unifying many geographically distributed federated computational and storage resources into a common “virtual site” abstraction, so that they can cooperate as if they were in the same Server Farm and Local Area Network. Simply stated, the MAG concept aims to make applications running on our metro Grid infrastructure aware of their complete computational and networking environment and capabilities, and able to make dynamic, adaptive and optimized use of heterogeneous network infrastructures connecting various high-end resources. As a proof of concept, they realized within the SCoPE High Performance Computing environment the prototype of a basic MAG architecture by implementing a novel centralized network resource management service supporting a flexible Grid-application interface and several effective network resource reservation facilities.

An efficient resource discovery mechanism is one of the fundamental requirements to resource management and scheduling of applications in federated grid networks. Resource discovery involves searching for the appropriate resource types that match the application requirements. Various kinds of solutions have been proposed lately. However, some of them have serious limitations in regard to scalability, fault-tolerance and network congestion. To this end, Gregor Pipan in his paper titled “Use of the TRIPOD Overlay Network for Resource Discovery”, presents a fully decentralized, efficient and highly scalable resource discovery approach that is applicable to large heterogeneous and highly dynamic distributed systems. The approach is based on a hybrid overlay network, named TRIPOD, which enables the efficient search for resources in the aforementioned highly distributed, dynamic and

largely heterogeneous systems. The key advantages of the proposed approach solution are its efficient proximity searching, ability to search over highly dynamic resource properties, the in-built fault tolerance and robustness and, finally, its very low and manageable network overhead.

Virtual Organization Clusters (VOCs) are another novel concept for overlaying dedicated private cluster systems on existing grid infrastructures. VOCs provide customized, homogeneous execution environments on a per-Virtual Organization basis, without the cost of physical cluster construction or the overhead of per-job containers. Administrative access and overlay network capabilities are granted to Virtual Organizations (VOs) that choose to implement VOC technology, while the system remains completely transparent to end users and non-participating VOs. Unlike existing systems that require explicit leases, VOCs are autonomically self-provisioned and self-managed according to configurable usage policies. In the paper “Virtual Organization Clusters: Self-Provisioned Clouds on the Grid”, Michael A. Murphy and Sebastien Goasguen presented a technology-agnostic formal model that describes the properties of VOCs and a prototype implementation of a physical cluster with hosted VOCs, based on the Kernel-based Virtual Machine (KVM) hypervisor. Test results demonstrate the feasibility of VOCs for use with high-throughput grid computing jobs. With the addition of a “watchdog” daemon for monitoring scheduler queues and adjusting VOC size, the results also demonstrate that cloud computing environments can be self-provisioned in response to changing workload conditions.

We hope that you will find the articles of this special issue to be informative and useful.

References

- [1] D. Abramson, R. Buyya, J. Giddy, A computational economy for grid computing and its implementation in the Nimrod-G resource broker, *Future Generation Computer System* 18 (8) (2002) 1061–1074. [http://dx.doi.org/10.1016/S0167-739X\(02\)00085-7](http://dx.doi.org/10.1016/S0167-739X(02)00085-7).
- [2] P. Trunfio, D. Talia, H. Papadakis, P. Fragopoulou, M. Mordacchini, M. Pennanen, K. Popov, V. Vlassov, S. Haridi, Peer-to-Peer resource discovery in Grids: Models and systems, *Future Generation Computer Systems* 23 (7) (2007) 864–878. <http://dx.doi.org/10.1016/j.future.2006.12.003>.
- [3] G. Aloisio, D. Talia, Guest editorial: grid computing: Towards a new computing infrastructure, *Future Generation Computer System* 18 (8) (2002) 5–6.
- [4] A. Di Stefano, G. Morana, D. Zito, A P2P strategy for QoS discovery and SLA negotiation in Grid environment, *Future Generation Computer System* 25 (8) (2009) 862–875. <http://dx.doi.org/10.1016/j.future.2009.03.001>.
- [5] R. Ranjan, A. Harwood, R. Buyya, A case for cooperative and incentive-based federation of distributed clusters, *Future Generation Computer System* 24 (4) (2008) 280–295. <http://dx.doi.org/10.1016/j.future.2007.05.006>.



Rajkumar Buyya is Professor of Computer Science and Software Engineering; 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 Pty Ltd., a spin-off company of the University, commercialising its innovations in Grid and Cloud Computing. He has authored and published over 300 research papers and four text books. The books on emerging topics that Dr. Buyya edited include *High Performance Cluster Computing* (Prentice Hall, USA, 1999), *Content Delivery Networks* (Springer, Germany, 2008) and *Market-Oriented Grid and Utility Computing* (Wiley, USA, 2009). He is one of the highly cited authors in computer science and software engineering worldwide (h-index = 48, g-index = 102, 12000+ citations).

Dr. Buyya has contributed to the creation of high-performance computing and communication system software for Indian PARAM supercomputers. He has pioneered the Economic Paradigm for Service-Oriented Distributed Computing and developed key Grid and Cloud Computing technologies such as Gridbus and Aneka that power the emerging e-Science and e-Business applications. 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 four IEEE conferences (CCGrid, Cluster, Grid, and e-Science). He has presented over 200 invited talks on his vision on IT Futures and advanced computing technologies at international conferences and institutions in Asia, Australia, Europe, North America, and South America. These contributions and the

international research leadership of Dr. Buyya are recognised through the award of “2009 IEEE Medal for Excellence in Scalable Computing” from the IEEE Computer Society, USA. For further information on Dr. Buyya, please visit his cyberhome: www.buyya.com.



Rajiv Ranjan is a Senior Research Associate (Lecturer level B) in the School of Computer Science and Engineering, University of New South Wales (UNSW). This position is funded through a strategic e-Research grant initiative of UNSW. Dr. Ranjan has a Ph.D. in Computer Science and Software Engineering from the University of Melbourne, 2009. He completed a Bachelor of Computer Engineering from North Gujarat University, India, in 2002. Dr. Ranjan is broadly interested in the emerging areas of cloud, grid, and service computing. The main goal of his current research is to advance the fundamental understanding and state of

the art of provisioning and delivery of application services in large, heterogeneous, uncertain, and evolving distributed systems.

Dr. Ranjan has 23 refereed publications, in journals with a high impact factor (according to JCR published by ISI), in proceedings of IEEE's/ACM's premier conferences and in books published by leading publishers (5 journals, 11 conferences, 4 book chapters, and 3 editorials). His h-index is 8, with a total citation count of 240, and his papers are cited an average of 34 times a year (Harzing's Publish or Perish). Dr. Ranjan has often served as Guest Editor for leading distributed systems and software engineering journals including *Future Generation Computer Systems* (Elsevier Press), *Concurrency and Computation: Practice and Experience* (John Wiley & Sons), and *Software: Practice and Experience* (Wiley InterScience). He was the Program Chair for 2010 Australasian Symposium on Parallel and Distributed Computing and 2010 IEEE TCSC Doctoral Symposium. He recently joined the Editorial Board of *International Journal of Information Technology, Communications and Convergence* (Inderscience Publishers). He serves as the editor of IEEE TCSC Newsletter. He has also recently initiated (as co-chair) IEEE TCSC Technical on Cloud Computing.