CONCURRENCY AND COMPUTATION: PRACTICE AND EXPERIENCE
Concurrency Computat.: Pract. Exper. 2012; **24**:1393–1396
Published online 30 November 2011 in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/cpe.1910

EDITORIAL

Middleware for Clouds and e-Science

1. BACKGROUND

This special issue focuses on middleware for Clouds and e-Science applications, and is based on extended, thoroughly revised papers from the 8th International Workshop on Middleware for Grid, Clouds and e-Science (MGC 2010) and the Workshop on Challenges in e-Science (CIS 2010). The authors were invited to provide extended versions of their original papers taking into account comments and suggestions raised during the peer review process and comments from the audience during the workshops.

Relevant contributions have been provided by Beloglazov and Buyya [1], Veiga *et al.* [2], Klôh *et al*. [3], Mathias and Baude [4], de Alencar *et al.* [5], Porto *et al*. [6], Ogasawara *et al.* [7], and Oliveira *et al.* [8]. These contributions focus on:

- Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in Cloud data centers;
- A checkpointing-enabled and resource-aware Java Virtual Machine (VM) for efficient and reliable e-science applications in grid environments;
- A bicriteria scheduling process with Class of Services support on Grids and Clouds;
- A component-based middleware for hybrid Grid/Cloud computing platforms;
- A peer-to-peer scheduling mechanism for workflows in grid computing;
- A metaphoric trajectory data warehouse for Olympic follow-up;
- A provenance management system for scientific workflows; and
- An adaptive parallel execution strategy for Cloud-based scientific workflows.

The eighth edition of the MGC workshop series encouraged researchers to submit and present original work related to middleware in Grids, Clouds and e-Science. Within this overall scope, MGC 2010 emphasized issues of scalability in terms of Grids and Clouds, mechanisms for computation and data scaling and adaptive and semantic mechanisms for scaling capabilities across disciplines and over time. The workshop was held in Bangalore India in conjunction with the ACM/IFIP/USENIX 11th International Middleware Conference.

The Challenges in eScience Workshop is an initiative of the National Laboratory of Scientific Computing with support from the Brazilian Society for the Progress of Science. The event gathered scientists and practitioners to discuss the latest achievements in supporting eScience applications. The availability of more powerful instruments and the deployment of high-speed networks are empowering researchers with unprecedented volume of data that must be cleaned, analysed and interpreted. Achieving this is, nevertheless, far from trivial. A complex computing arsenal composed of Grids and Clouds equipped with multicore machines is driving the development of software to support scientists in the difficult task of understanding the data deluge. In this context, the Challenges in eScience Workshop has invited world leaders in eScience to discuss solutions and point to directions the community should follow to fulfil the potential of in-silico research. The event was held in Petrópolis, RJ, Brazil, in conjunction with the IEEE 22nd International Symposium on Computer Architecture.

2. SPECIAL ISSUE PAPERS

Beloglazov and Buyya [1] formally define the single VM migration and dynamic VM consolidation problems. They conduct competitive analysis and prove competitive ratios of optimal online deterministic algorithms for the defined problems. Then, the authors propose novel adaptive heuristics for dynamic consolidation of VMs, which significantly reduce energy consumption, while ensuring a high level of adherence to the service level agreements.

Veiga *et al.* [2] suggest that increasingly more Java applications are gaining relevance in fields related with e-Science (with Grid and Cloud computing), having long (or very long) execution times and manipulating large amounts of data/information. Significant examples include chemistry, computational biology and bioinformatics, with many available Java-based APIs. When the execution of such applications is terminated because of a failure, all work already performed is simply lost, and applications must restart from scratch, wasting resources and time, and delaying completion with no deadline guarantees. These issues are addressed by incorporating comprehensive and efficient mechanisms for checkpointing and migration in a Java VM (Jikes RVM), rendering applications more robust and flexible, and able to move to other nodes, without any intervention from the programmer.

Klôh *et al.* [3] present a workflows scheduling model that is based on a hybrid bicriteria scheduling algorithm based on a Class of Services approach. The proposed model is based on the analysis of several related works on resource scheduling, trying to incorporate the most relevant aspects considered in these works, thus covering some shortcomings pointed out by them. We are particularly interested in the transfer of intermediate data — that build up a workflow — with reduced cost and identifying the reliability level of the resources. We also work-out different bicriteria scheduling approaches, with the criterion priority order chosen by the user. This is implemented by optimizing the first criterion and based on the outcome tries to optimize the next criterion, in accordance with a variation limit. The proposed model aims to allow scheduling with: reduced response time to the user, better use of resources, reducing the make-span and choosing the best resource using historical data from user applications. To validate this model, a set of tests were performed comparing the performance of the proposed model with the Join the Shortest Queue scheduling algorithm. The analysis of the obtained results showed an improvement in the overall quality made available to the users — as they meet the priority criteria — and also a performance gain.

Mathias and Baude [4] present a GCM/ProActive-based generic and extensible middleware intended to ease the transition from clusters, to Grids and Clouds and/or a mixture of them. Distributed computational resources are everywhere: in company and institution datacenters and clusters, in computational Grids and more recently in the Cloud. Several software solutions have been proposed for gathering these resources together offering an abstraction layer that hides the complexity of such multidomain platforms. This is undeniably a good approach, but many applications require a better knowledge of the running environment, for instance to keep performance (e.g., nonembarrassingly parallel applications) or to cope with regulations and agreements (e.g., business-to-business applications). In this article, they propose a lightweight component-based middleware addressing the requirements of such applications. The key points of this middleware are: a modular infrastructure that can adapt to the running environment and application connectivity requirements, topology-aware message routing and a minimalistic interface allowing an easy integration of applications for the execution in multidomain environments. They also present two use cases in which this middleware helped solve integration issues: the SOA4All Federated Distributed Service Bus, which benefit from their middleware to integrate service buses and to offer service discovery and invocation across multiple service providers and the DiscoGrid Runtime, which benefit from their middleware as an overlay supporting the execution of connected MPI applications in multidomain environments.

de Alencar *et al.* [5] present computational grids as adequate for the execution of scientific workflows because they allow the use of heterogeneous resources and support a variety of applications. Centralized scheduling solutions for grids may lead to performance degradation and to scalability issues. In this work, they present P2PScheMe, a strategy for peer-to-peer execution of workflows based on the invocation of Grid services. The proposal considers information regarding Grid execution environment to allow workflow scheduling adaptation. This adaptation is performed according to user's QoS requirements. In the paper they describe how the P2PScheMe works and provide a comparative analysis against existing solutions.

DOI: 10.1002/cpe

Porto *et al.* [6], present a trajectory data model that has recently been adopted as a semantic enrichment of traditional time-series values. This paper applies this model to the follow-up of Olympic athlete's data. A database schema in support of trajectory analysis is proposed and analytical queries comparing biochemical data among athletes is presented. Additionally, the paper explores a multidimensional implementation for trajectory data analysis and compares this approach against analysis over the relational database model. The benefits of the trajectory approach analysis to sport scientists are discussed.

Ogasawara *et al.* [7] propose a data provenance management approach to keep track of several trials during scientific experiments executions. Gathering provenance data in a structured way keeps track of such executions and allows for querying experiment data descriptions in a higher level than execution logs. Moreover, the approach presents a strategy for managing provenance data in distributed and heterogeneous environments where experiment data are generated by multiple systems. The main idea is to leverage provenance management techniques by integrating different scientific experiment executions from multiple management systems and to allow scientists to have a unique and integrated vision of the whole experiment.

Finally, Oliveira *et al.* [8] introduce an adaptive approach for executing parallel scientific workflows in the Cloud that benefits from the elasticity characteristic of Cloud environments. This approach adapts itself according to the availability of resources during parallel workflow execution. The central idea is to check the available computational power and dynamically tune the workflow activity size to achieve better performance.

ACKNOWLEDGEMENTS

We would like to thank the authors for contributing papers on their research on middleware for Clouds and e-Science for this special issue, and all the reviewers for providing constructive reviews and in helping to shape this special issue. Finally, we would like to thank Prof. Geoffrey Fox for providing us an opportunity to bring this special issue to the research community. Previous issues can be found at [9–15] and [16–23].

REFERENCES

1. Beloglazov A, Buyya R. Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in Cloud data centers. *Concurrency and Computation: Practice and Experience* 2012; **24**(13):1397–1420.
2. Veiga L, Simão J, Garrochinho T. A checkpointing-enabled and resource-aware Java VM for efficient and robust e-Science applications in grid environments. *Concurrency and Computation: Practice and Experience* 2012; **24**(13):1421–1442.
3. Klôh H, Schulze B, Pinto R, Mury A. A bi-criteria scheduling process with CoS support on grids and clouds. *Concurrency and Computation: Practice and Experience* 2012; **24**(13):1443–1460.
4. Mathias E, Baude F. A component-based middleware for hybrid grid/cloud computing platforms. *Concurrency and Computation: Practice and Experience* 2012; **24**(13):1461–1477.
5. de Alencar JM, Andrade R, Viana W, Schulze B. P2PScheMe: a P2P scheduling mechanism for workflows in grid computing. *Concurrency and Computation: Practice and Experience* 2012; **24**(13):1478–1496.
6. Porto F, Moura AM, da Silva F, Bassini A, Palazzi D, Poltosi M, de Castro L, Cameron L. A mataphoric trajectory data warehouse for Olympic follow-up. *Concurrency and Computation: Practice and Experience* 2012; **24**(13):1497–1512.
7. Ogasawara E, Marinho A, Murta L, Werner C, Braganholo V, Cruz S, Mattoso M. ProvManager: a provenance management system for scientific workflows. *Concurrency and Computation: Practice and Experience* 2012; **24**(13):1513–1530.
8. Oliveira D, Ogasawara E, Ocana K, Baião F, Mattoso M. An adaptive parallel execution strategy for cloud-based scientific workflows. *Concurrency and Computation: Practice and Experience* 2012; **24**(13):1531–1550.
9. Schulze B, Nandkumar R, Magedanz T. Concurrency and Computation: Practice and Experience–Special Issue on Middleware for Grid Computing. *Wiley Interscience* 2004; **16**(5).
10. Schulze B, Nandkumar R. Concurrency and Computation: Practice and Experience–Special Issue on Middleware for Grid Computing. *Wiley Interscience* 2006; **18**(6).
11. Schulze B, Coulson G, Nandkumar R, Henderson P. Concurrency and Computation: Practice and Experience–Special Issue on Middleware for Grid Computing. *Wiley Interscience* 2007; **19**(14).
12. Schulze B, Nandkumar R, Abramson D, Buyya R. Concurrency and Computation: Practice and Experience–Special Issue on Middleware for Grid Computing. *Wiley Interscience* 2008; **20**(9).
13. Schulze B, Rana O, de Souza JN. Concurrency and Computation: Practice and Experience–Special Issue: Advanced Strategies in Grid Environments-Models and Techniques for Scheduling and Programming. *Wiley Interscience* 2009; **21**(13).

DOI: 10.1002/cpe

14. Schulze B, Fox G. Concurrency and Computation: Practice and Experience–Special Issue: Middleware strategies for clouds and grids in e-Science 2011; **22**(3).
15. Schulze B, Myers J. Concurrency and Computation: Practice and Experience - Special Issue: Advanced Scheduling Strategies and Grid Programming Environments. *Wiley Interscience* 2010; **23**(17).
16. Schulze B, Nandkumar R, Magedanz T. Proceedings of the 1st International Workshop on Middleware for Grid Computing. PUC-Rio, pp. 169-266; Rio de Janeiro, 2003.
17. Schulze B, Nandkumar R. Proceedings of the 2nd International Workshop on Middleware for Grid Computing. *ACM International Conference Proceeding Series* 2004; **76**.
18. Schulze B, Nandkumar R, Henderson P. Proceedings of the 3rd International Workshop on Middleware for Grid Computing. *ACM International Conference Proceeding Series* 2005; **117**.
19. Schulze B, Nandkumar R, Abramson D, Buyya R. Proceedings of the 4th International Workshop on Middleware for Grid Computing. *ACM International Conference Proceeding Series* 2006; **194**.
20. Schulze B, Rana O, Meyer J, Cirne W. Proceedings of the 5th International Workshop on Middleware for Grid Computing. *ACM International Conference Proceeding Series* 2007; **64**.
21. Schulze B, Fox G. Proceedings of the 6th international workshop on Middleware for Grid Computing Leuven, Belgium, Article No. 11, Vol 72.
22. Schulze B, Myers J. Proceedings of the 7th International Workshop on Middleware for Grids, Clouds and e-Science. Champaign, IL, USA, Vol.72, 2009.
23. Schulze B, Buyya R. Proceedings of the 8th International Workshop on Middleware for Grids, Clouds and e-Science. Bangalore, India, 2010.

BRUNO SCHULZE
Distributed Scientific Computing (COMCIDIS), National Laboratory for Scientific Computing (LNCC), Av. Getulio Vargas 333, Petropolis – RJ, Brazil
E-mail: *schulze@lncc.br*

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

FABIO PORTO
Distributed Scientific Computing (COMCIDIS), National Laboratory for Scientific Computing (LNCC), Av. Getulio Vargas 333, Petropolis – RJ, Brazil

DOI: 10.1002/cpe