

VU Research Portal

HPS-HDS

Pop, Florin; Iosup, Alexandru; Prodan, Radu

published in

Future Generation Computer Systems
2018

DOI (link to publisher)

[10.1016/j.future.2017.09.012](https://doi.org/10.1016/j.future.2017.09.012)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Pop, F., Iosup, A., & Prodan, R. (2018). HPS-HDS: High Performance Scheduling for Heterogeneous Distributed Systems. *Future Generation Computer Systems*, 78(Part 1), 242-244.
<https://doi.org/10.1016/j.future.2017.09.012>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl



Editorial

HPS-HDS: High Performance Scheduling for Heterogeneous Distributed Systems



Florin Pop^{a,b,*}, Alexandru Iosup^{c,d}, Radu Prodan^e

^a University Politehnica of Bucharest, Romania

^b National Institute for Research and Development in Informatics, Bucharest, Romania

^c Delft University of Technology, The Netherlands

^d Vrije Universiteit Amsterdam, The Netherlands

^e University of Innsbruck, Austria

ARTICLE INFO

Keywords:

Scheduling algorithms

Resource management

Fault tolerance

Big data

Heterogeneous distributed systems

ABSTRACT

Heterogeneous Distributed Systems (HDS) are often characterized by a variety of resources that may or may not be coupled with specific platforms or environments. Such type of systems are Cluster Computing, Grid Computing, Peer-to-Peer Computing, Cloud Computing and Ubiquitous Computing all involving elements of heterogeneity, having a large variety of tools and software to manage them. As computing and data storage needs grow exponentially in HDS, increasing the size of data centers brings important diseconomies of scale. In this context, major solutions for scalability, mobility, reliability, fault tolerance and security are required to achieve high performance. More, HDS are highly dynamic in its structure, because the user requests must be respected as an agreement rule (SLA) and ensure QoS, so new algorithm for events and tasks scheduling and new methods for resource management should be designed to increase the performance of such systems. In this special issues, the accepted papers address the advance on scheduling algorithms, energy-aware models, self-organizing resource management, data-aware service allocation, Big Data management and processing, performance analysis and optimization.

© 2017 Published by Elsevier B.V.

1. Introduction

High-performance scheduling has been traditionally approached as an NP-complete optimisation problem having execution time as the single objective of interest. The increased complexity of the currently emerging heterogeneous distributed systems aggregating parallel computers, clusters, peer-to-peer systems, or virtualized Cloud operating either in centralized data centers or distributed towards the edge of the network, bring along additional objectives of equal importance, such as energy consumption, operational costs, reliability, security, uncertainty, elasticity, just to mention few. The optimization challenge is that these objective are often in contradiction with each other, requiring new algorithms with increased complexity that converge to acceptable trade-off solutions.

The problem of task scheduling and resource allocation under SLA constraints is one of the main challenges the context of heterogeneous distributed systems, because we face with the need to satisfy the contract between service providers and clients considering that the total profit in the system depends on how the system meets different demands.

Additionally, we have a trade-off between the scope of cloud providers that is represented by the need to maximize profit and the client satisfaction. An SLA that is in place can include various parameters such as computational power, memory and/or hard disk space, network bandwidth, availability of the resources, data security, consumed power and so on. In order to meet the the agreements with clients, service providers offer and maintain a large number of resources than it is the case, or end up breaking the SLA's because they cannot meet the agreed parameters.

Thus optimizing the resource allocation process is a must in order to maximize the provider's profit and to honor the SLA. Honoring a SLA contract requires building an accurate resource allocation model which assesses the amount (and characteristics) of the needed resources. Building such a model became difficult because the amount of resources needed in order to meet such an SLA can vary significantly over time.

2. This special issue

This special issue called for original papers on latest research and innovations, solutions and developments on High Performance Scheduling for HDS. Authors were solicited to submit complete unpublished papers oriented on the following, but not limited to: foundational models for task scheduling and resource

* Corresponding editor.

E-mail addresses: florin.pop@cs.pub.ro, florin.pop@ici.ro (F. Pop),

A.iosup@tudelft.nl, A.iosup@vu.nl (A. Iosup), radu@dps.uibk.ac.at (R. Prodan).

management in HDS, distributed scheduling algorithms, adaptive and machine learning based scheduling algorithms, dynamic resource provisioning, load-balancing and co-allocation, data-aware scheduling, self-* techniques for high performance scheduling, scheduling in big data platforms, content distribution systems for large data, data-intensive computing applications, scheduling for map-reduce and Hadoop, high-throughput computing applications, cloud workload profiling and deployment control, workflow scheduling and scalability analysis, scheduling for many-task computing, cloud resource virtualization and composition, task offloading and scheduling techniques for mobile cloud computing, resource management for high performance cloud computing, scheduling for green computing, quality management and service level agreement (SLA), reliability and fault tolerance of HDS, performance evaluation.

The special issues follows the previous published work like ARMCO (advanced topics in resource management for ubiquitous cloud computing: an adaptive approach) [1], and MidHDC (advanced topics on middleware services for heterogeneous distributed computing) part 1 and part 2 [2,3].

In the first paper, a competitive analysis of fundamental scheduling algorithms on a fault-prone machine and the impact of resource augmentation is presented by Antonio Fernandez Anta et al. [4]. The main conclusion highlighted by the authors shows that all deterministic and work-conserving algorithms require speedup to be competitive.

Fredy Juarez et al., in the second paper, titled “Dynamic Energy-Aware Scheduling for Parallel Task-Based Application in Cloud Computing” [5], describe a model for for estimating the application energy consumption and propose an energy-aware run-time scheduler for task-based applications using a multi-heuristic resource allocation algorithm to get solutions in polynomial time.

In the third paper, Rodrigo da Rosa Righi et al. are proposing MigPF, a migration model for bulk-synchronous parallel programs on heterogeneous environments [6]. Based on the simple ideas from the TCP protocol, MigPF adapts the interval between migration calls. As a new contribution, the authors propose a new algorithm that self-organizes process migration.

A novel sampling method for massive data as the input of Spark framework is introduced by Zhuo Tang et al. in the 4th paper [7]. Based on this, the authors describe and algorithm for filling the settled number of buckets with roughly equal number of tuples.

Michał P. Karpowicz et al. highlights in the 5th paper the design and implementation of energy-aware application-specific CPU frequency governors for the heterogeneous distributed computing systems [8]. The authors propose a benchmarking methodology to identify models of CPU workload dynamics.

The 6th paper, authored by Wenbin Jiang et al. [9], proposes a fine-grained parallelism for intra prediction based on GPU. The authors present a new regular prediction formula for parallelism and two optimized encoding orders are adopted for multi-levels parallelism.

The 7th paper addresses the problem of Data Centres (DCs) integration into the Smart Grid scenario by proposing a technique for scheduling and optimizing their operation allowing them to participate in Smart Demand Response programs [10]. Tudor Cioara et al.

A bio-inspired coral-reefs optimization paradigm to model cloud elasticity and a game theory-based approach to identify the best cloud resource reallocation schema Fuzzy linguistic SLA formalization is considered by Massimo Ficco et al. in the 8th paper [11].

The 9th paper, authored by Lizhe Wang et al. [12], considers a Cloud-enabled HPC platform for large-scale remote sensing dataset applications. The authors propose in this paper, Hilbert-R⁺, a tree based data indexing for optimal remote sensing big data indexing.

They consider collaborative large-scale remote sensing workflow processing across data centers with VMs and bare-metal provisioning.

Konstantinos Douzis et al. propose a modular and generic IoT management system in the Cloud [13]. The authors developed a sensor data collection service, which separates IoT devices and cloud system with NoSQL data storage support. It supports vendor agnostic and on the fly data collection from IoT devices and fast incoming traffic (POST data around 1.7 s per event).

A resource provisioning framework for bioinformatics applications in multi-Cloud environments is presented by Izzet F. Senturk et al. in the 11th paper [14]. The proposed framework simplifies extending local resources to a multi-cloud environment. The authors show that the considered scheduling algorithm decreases the workflow execution time for a given budget.

The last paper, authored by Rong Jiang et al. [15], propose a high performance and privacy-preserving query scheme over encrypted multidimensional big metering data. In the proposed solution, encrypted metering data are stored in the server of an untrusted heterogeneous distributed system environment.

Acknowledgments

The guest editors would like to thank all authors for their valuable contributions and efficient cooperation in the preparation of this special issue. We appreciate the efforts all referees for to provide timely, constructive, thoughtful and knowledgeable reviews and for their substantial contributions to the final revised versions appearing in this special issue. We express our gratitude and thank to them for their hard work. We address our personal warm regards to Prof. Peter Sloot, the Editor-in-Chief of Future Generation Computing Systems Journal for his support and advises, as well as to the editorial and managerial journal team from Elsevier for their assistance and excellent cooperative collaboration to produce this valuable scientific work.

References

- [1] F. Pop, M. Potop-Butucaru, Armco: Advanced topics in resource management for ubiquitous cloud computing: An adaptive approach, *Future Gener. Comput. Syst.* 54 (2016) 79–81. <http://dx.doi.org/10.1016/j.future.2015.07.016>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X15002484>.
- [2] F. Pop, X. Zhu, L.T. Yang, Midhdc: Advanced topics on middleware services for heterogeneous distributed computing. part 1, *Future Gener. Comput. Syst.* 56 (2016) 734–735. <http://dx.doi.org/10.1016/j.future.2015.11.020>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X15003738>.
- [3] F. Pop, X. Zhu, L.T. Yang, Midhdc: Advanced topics on middleware services for heterogeneous distributed computing. Part 2, *Future Gener. Comput. Syst.* 74 (2017) 86–89. <http://dx.doi.org/10.1016/j.future.2017.05.010>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X1730972X>.
- [4] A.F. Anta, C. Georgiou, D.R. Kowalski, E. Zavou, Competitive analysis of fundamental scheduling algorithms on a fault-prone machine and the impact of resource augmentation, *Future Gener. Comput. Syst.* 78 (2018) 245–256. <http://dx.doi.org/10.1016/j.future.2016.05.042>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301674>.
- [5] F. Juarez, J. Ejarque, R.M. Badia, Dynamic energy-aware scheduling for parallel task-based application in cloud computing, *Future Gener. Comput. Syst.* 78 (2018) 257–271. <http://dx.doi.org/10.1016/j.future.2016.06.029>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X1630214X>.
- [6] R. da Rosa Righi, R. de Quadros Gomes, V.F. Rodrigues, C.A. da Costa, A.M. Alberti, L.L. Pilla, P.O.A. Navaux, Migpf: Towards on self-organizing process rescheduling of bulk-synchronous parallel applications, *Future Gener. Comput. Syst.* 78 (2018) 272–286. <http://dx.doi.org/10.1016/j.future.2016.05.004>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301145>.
- [7] Z. Tang, X. Zhang, K. Li, K. Li, An intermediate data placement algorithm for load balancing in spark computing environment, *Future Gener. Comput. Syst.* 78 (2018) 287–301. <http://dx.doi.org/10.1016/j.future.2016.06.027>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16302126>.
- [8] M.P. Karpowicz, P. Arabas, E. Niewiadomska-Szynkiewicz, Design and implementation of energy-aware application-specific {CPU} frequency governors for the heterogeneous distributed computing systems, *Future Gener. Comput. Syst.* 78 (2018) 302–315. <http://dx.doi.org/10.1016/j.future.2016.05.011>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301212>.

- [9] W. Jiang, M. Long, L.T. Yang, X. Liu, H. Jin, A.L. Yuille, Y. Chi, Fpip: A novel fine-grained parallel partition based intra-frame prediction on heterogeneous many-core systems, *Future Gener. Comput. Syst.* 78 (2018) 316–329. <http://dx.doi.org/10.1016/j.future.2016.05.009>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301194>.
- [10] T. Cioara, I. Anghel, M. Bertocchini, I. Salomie, D. Arnone, M. Mammina, T.-H. Velivassaki, M. Antal, Optimized flexibility management enacting data centres participation in smart demand response programs, *Future Gener. Comput. Syst.* 78 (2018) 330–342. <http://dx.doi.org/10.1016/j.future.2016.05.010>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301200>.
- [11] M. Ficco, C. Esposito, F. Palmieri, A. Castiglione, A coral-reefs and game theory-based approach for optimizing elastic cloud resource allocation, *Future Gener. Comput. Syst.* 78 (2018) 343–352. <http://dx.doi.org/10.1016/j.future.2016.05.025>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301364>.
- [12] L. Wang, Y. Ma, J. Yan, V. Chang, A.Y. Zomaya, pipscloud: High performance cloud computing for remote sensing big data management and processing, *Future Gener. Comput. Syst.* 78 (2018) 353–368. <http://dx.doi.org/10.1016/j.future.2016.06.009>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301923>.
- [13] K. Douzis, S. Sotiriadis, E.G. Petrakis, C. Amza, Modular and generic iot management on the cloud, *Future Gener. Comput. Syst.* 78 (2018) 369–378. <http://dx.doi.org/10.1016/j.future.2016.05.041>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301662>.
- [14] I.F. Senturk, P. Balakrishnan, A. Abu-Doleh, K. Kaya, Q. Malluhi, Ümit V. Çatalyrek, A resource provisioning framework for bioinformatics applications in multi-cloud environments, *Future Gener. Comput. Syst.* 78 (2018) 379–391. <http://dx.doi.org/10.1016/j.future.2016.06.008>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301911>.
- [15] R. Jiang, R. Lu, K.-K.R. Choo, Achieving high performance and privacy-preserving query over encrypted multidimensional big metering data, *Future Gener. Comput. Syst.* 78 (2018) 392–401. <http://dx.doi.org/10.1016/j.future.2016.05.005>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X16301157>.



Florin Pop received his Ph.D. in Computer Science at the University Politehnica of Bucharest in 2008. He received his M.Sc. in Computer Science in 2004 and the Engineering degree in Computer Science in 2003, at the same university. He is Professor within the Computer Science Department and also an active member of Distributed System Laboratory. His research interests are in scheduling and resource management (decentralized techniques, re-scheduling), multi-criteria optimization methods, Grid middleware tools and applications development (satellite image processing and environmental data analysis), prediction methods, self-organizing systems, contextualized services in distributed systems. He is the author or co-author of many publications as books, chapters, papers in international journals and well-established and ranked conferences. He served as guest-editor for International Journal of Web and Grid Services and

he is Managing Editor for International Journal of Grid and Utility Computing. He was awarded with “Magna cum laude” distinction from University Politehnica of Bucharest for his Ph.D. results, one IBM Faculty Award in 2012 or the project “CloudWay – Improving resource utilization for a smart Cloud infrastructure”, two Prizes for Excellence from IBM and Oracle (2008, 2009), Best young researcher in software services Award, FP7 SPRERS Project, Strengthening the Participation of Romania at European R&D in Software Services in 2011 and two Best Paper Awards. He is working in several international and national research projects in the distributed systems field as coordinator and member as well. He is a senior member of the IEEE, and a member of the ACM.



Prof. dr. ir. Alexandru Iosup is currently Tenured Full Professor and University Research Chair with the Computer Systems department at the Vrije Universiteit Amsterdam, the Netherlands, Associate Professor with the Distributed Systems group at TU Delft, and Chair of the SPEC Research Cloud Group. In 2009, he received a Ph.D. in Computer Science from TU Delft. He received the yearly Netherlands Prize for Research in Computer Science (2016), the yearly Netherlands Teacher of the Year (2015), several SPECTacular awards (2012–2017), and numerous research awards and nominations. His research interests are in massivizing computer systems, that is, making computer systems combine desirable properties such as elasticity, performance, and availability, yet maintain their ability to operate efficiently in controlled ecosystems. Topics include cloud computing and big data, with applications in big science, big business, online gaming, and (upcoming) massivized education. His work is funded by a combination of prestigious personal grants (NWO/STW Vidi 2016 and NWO/STW Veni 2011), generous industry gifts and collaborations (Solvinity, Oracle, Intel), and EU and national projects.



Radu Prodan is Associate Professor of computer science at the University of Innsbruck. His research interests in the areas of parallel and distributed systems comprise programming methods, compiler technology, performance analysis, and scheduling for high-performance scientific computing and industrial online gaming applications. Before joining University of Innsbruck, he held previous research positions at ETH Zurich, University of Basel, Swiss Scientific Computing Centre, and University of Vienna. He participated in several national and European projects and is currently coordinating three Austrian projects in the areas of multicore and Cloud computing. Prodan was WP leader in the IST- 034601 edutain@grid and 26185 (SHIWA) European projects, and is currently scientific coordinator of the H2020-ICT-2014-1 644179 project ENTICE. Currently he is co-chair of the “state-of-the-art and continuous learning in UCS” working group of the Network for Sustainable Ultrascale Computing (NESUS) COST Action funded by the EU that brings together partners from over 30 European countries. He is the author of one book, over 100 journal and conference publications, and is the recipient of two IEEE best paper awards.