

Editorial

**Bernabé Dorronsoro · Dzmitry Kliazovich ·
Pascal Bouvry · Albert Y. Zomaya**

Published online: 15 March 2016
© Springer Science+Business Media Dordrecht 2016

Cloud computing is nowadays very popular and extended among Internet users. It provides a wide variety of services, from high performance computing to storage or web services. Behind the cloud, there are huge datacenters that are usually composed of a large variety of computing, storage, and networking devices, among others, that boost their power consumption. This important issue does not only concern the increasing cost of operating these centers, but it also implies an important environmental impact.

Optimizing the energy consumption in datacenters is therefore of high importance not only for companies, but for governments too. There are several

ways that may lead to a reduction of the energy consumption in datacenters. An efficient utilization of the resources is obviously one of them, but not the only one. Research on novel more efficient cooling systems, new computing and network architectures, and storage systems is needed for the next generation green datacenters. In addition, the design of highly efficient software and services is a must for the green cloud.

This special issue focuses on novel solutions that contribute to the reduction of energy consumption in the cloud, from a number of different perspectives.

The work “Online Bi-Objective Scheduling for IaaS Clouds Ensuring Quality of Service”, by Tchernykh et al., addresses the online resource management problem in clouds by minimizing the energy consumption and maximizing the revenue of the provider, while respecting the Quality of Service of the system. The proposed online schedulers are designed in two levels: in the first level it is decided whether a job is accepted or not, and in the second level it is allocated into the servers, if accepted in level one. The paper offers a detailed analysis of the performance of a number of different schedulers on real HPC workloads.

Kliazovich et al. propose in their work a novel approach for modeling the resource allocation in clouds, considering communication tasks explicitly, together with computation ones. The proposed

B. Dorronsoro (✉)
School of Engineering, University of Cádiz, Cádiz, Spain
e-mail: bernabe.dorronsoro@uca.es

D. Kliazovich · P. Bouvry
Computer Science and Communications Research Unit,
University of Luxembourg, Luxembourg City, Luxembourg

D. Kliazovich
e-mail: dzmitry.kliazovich@uni.lu

P. Bouvry
e-mail: pascal.bouvry@uni.lu

A. Y. Zomaya
School of Information Technologies, The University
of Sydney, New South Wales, Australia
e-mail: albert.zomaya@sydney.edu.au

communication-aware model allows making separate resource allocation decisions for communications and computations, unlike the other existing approaches. Therefore, schedulers can take into account not only computing but also network resources for information transmission. The paper demonstrates how taking into account the two kinds of resources at the same time allows the algorithms finding more accurate and efficient schedules.

A flexible virtual machines scheduler for task-based cloud applications is presented in “A Pliant-based Virtual Machine Scheduling Solution to Improve the Energy Efficiency of IaaS Clouds”. The new method provides energy-efficient and accurate solutions for IaaS cloud management. The results show, based on simulations guided by real traces, how significant savings in energy consumption are achieved when using the proposed method, with respect to a number of heuristic schedulers.

The paper by Tang et al., with title “An Energy-Efficient Task Scheduling Algorithm in DVFS-enabled Cloud Environment”, presents a novel scheduler for parallel applications with deadlines. The algorithm works in three steps: (1) finding accurate solutions with the well-known HEFT heuristic, (2) compacting the solution by relocating the tasks assigned to underutilized processors, therefore freeing resources, and (3) using dynamic voltage/frequency scaling to fill all gaps (respecting deadlines) present in the schedule. Results show how the proposed solution can significantly outperform other existing approaches.

The work “Topology-Aware Virtual Machine Placement in Data Centers” proposes a Virtual Machine (VM) allocation mechanism that takes into account the network topology so that the VMs are allocated in the smallest possible area of the datacenter, while preventing the formation of bottlenecks. The energy consumption of both servers and switches is considered in the allocation process of VMs, and those that remain idle are switched off.

The work by Borgetto et al. focuses on the design of an architecture to extend the classical cloud model to integrate personal computers with the cloud. Users can have access to a number of services (defined as remote applications), to which they subscribe, that are either locally installed or remotely executed in the cloud or even in other user computer, depending on the capacities of the local computer and the criterion

of the system. This is done in a transparent way to the user.

The paper “Green Cloud Meta-Scheduling: A Flexible and Automatic Approach”, proposes an energy-efficient approach to schedule VMs into the existing servers in a datacenter in order to minimize their energy consumption, while ensuring that the service level agreement (SLA) is respected. For that, the framework is able to follow different policies in the whole allocation process (e.g., choosing VMs to deploy, where to deploy them, or assigning them the tasks to compute) according to a number of objectives, and automatically choosing the most convenient one in every situation. The approach is validated versus others from the literature through simulations.

The work by Boryło et al. suggests the need of cooperation between cloud providers and the network operators interconnecting them in order to minimize the CO₂ emissions related to their service provisioning. Different models of cooperation are considered. Experiments on two real platforms validate the approach proposed in the work.

The paper by Marszałkowski et al. studies the influence of a scheduler’s problem size in the power consumption in parallel systems with hierarchical memory architecture. The motivation for this study is that the different technologies used for the memory at each level of the hierarchy provide different characteristics in terms of speed, size, and power consumption, among others. The paper concludes that, the setting of the parallel platform has a major influence on makespan and energy results, and improving the performance of computations involves a careful parameters tuning.

Shi et al. target the minimization of the influence that VMs have on the other VMs that are hosted in the same server, due to the limited resources of the server. This performance isolation of VMs plays an important role to guarantee cloud computing SLAs. The paper presents a novel architecture that adaptively adjusts the resource allocation of virtualized web servers, which present a highly dynamic workload, in order to minimize the resources contention in the servers to satisfy the SLA. In addition, the CPU frequency is controlled too in order to minimize the server power consumption.

The guest editors would like to express their deepest gratitude to the Editor in Chief of the Journal of Grid Computing, Prof. Peter Kacsuk, for giving us the opportunity to organize this special issue and for his great support. We are also very grateful to the Springer team and to the anonymous reviewers for

their precious help. Finally we would like to thank the authors of the selected papers for their high quality contributions.

The Guest editors
Dorronsoro, Kliazovich, Bouvry & Zomaya