Editorial

Special Issue on Algorithms for the Resource Management of Large Scale Infrastructures

Danilo Ardagna [1,*], **Claudia Canali** [2] **and Riccardo Lancellotti** [2]

[1] Politecnico di Milano, Dipartimento di Elettronica, Informazione e Bioingegneria Via Golgi 42, 20133 Milano, Italy

[2] Dipartimento di Ingegneria Enzo Ferrari, Università di Modena e Reggio Emilia, Via P. Vivarelli 10/1, 41125 Modena, Italy; claudia.canali@unimore.it (C.C.); riccardo.lancellotti@unimore.it (R.L.)

[*] Correspondence: danilo.ardagna@polimi.it

Abstract: Modern distributed systems are becoming increasingly complex as virtualization is being applied at both the levels of computing and networking. Consequently, the resource management of this infrastructure requires innovative and efficient solutions. This issue is further exacerbated by the unpredictable workload of modern applications and the need to limit the global energy consumption. The purpose of this special issue is to present recent advances and emerging solutions to address the challenge of resource management in the context of modern large-scale infrastructures. We believe that the four papers that we selected present an up-to-date view of the emerging trends, and the papers propose innovative solutions to support efficient and self-managing systems that are able to adapt, manage, and cope with changes derived from continually changing workload and application deployment settings, without the need for human supervision.

Keywords: cloud and edge/fog computing; resource management; green computing

---

1. Introduction

The growing demand for fast, powerful, scalable, and reliable computing and communication infrastructures has driven the evolution of the computation paradigm from in-house solutions towards shared infrastructures, such as the cloud computing paradigm (which in the long term can provide a reduction in the total cost of ownership by up to 66%), building up to the most innovative distributed infrastructures being proposed for edge/fog computing. The result is the advent of large distributed infrastructures composed of multiple data centers, where virtualization is applied at both the computing and networking levels, to the point where we can describe the infrastructure as being composed of Software-Defined Data Centers.

These infrastructures are becoming increasingly popular, especially to support modern applications which are characterized by huge volumes of data (in the order of peta or exabytes—the so called big data), possibly collected from thousands of sensors (in the Internet of Things scenario) and may be deployed on tens of thousands of Virtual Machines (VMs), possibly interacting through a virtualized network spread over a geographic area. For example, more than 60% of the Apache Spark installations are deployed on the cloud. However, the management of these large and complex infrastructures represents a major problem. On one hand, there is the need to maximize resource utilization and to minimize the energy consumption of the infrastructure, to reduce costs and for environmental reasons. On the other hand, there is the need to guarantee the resources required by the applications hosted on the infrastructure in terms of computing power, storage, and communication. These requirements may be expressed directly or may be directed via a Service Level Agreement at the level of the application behavior.

The challenge of infrastructure management in this scenario is clear. Even the monitoring of such infrastructure is difficult, but we must consider that the complexity of the problem is further exacerbated by the need for the infrastructure to react in a timely and unsupervised manner with an ever-changing workload, characterized by the unpredictable oscillations in application demands through the joining and leaving of new applications that are deployed or dropped. The goal is to devise models, techniques, and algorithms that can support a self-management system that is able to adapt, manage, and cope with changes, without the need for human supervision.

2. Special Issue

All submissions to this Special Issue were reviewed by at least three experts. After a strict review, four papers were accepted, as introduced below.

All the papers focus on distributed systems management, but the considered scenarios range from multi-cloud to fog-like and Peer-to-Peer (P2P) architectures, with two papers proposing solutions that are power- and energy-aware. The proposed techniques are wide ranging, from modeling based on stochastic Petri-networks to heuristics and machine learning algorithms. We believe that this selection highlights the vast heterogeneity that characterizes modern distributed infrastructures, thereby proposing novel solutions for their management and highlighting the broad scope of the research tackling these problems.

The paper "Multi-level elasticity for wide-area Data Streaming Systems: A reinforcement learning approach" by Russo Gabriele et al. [1] addresses the issue of data stream processing (DSP) in the context of a multi cloud/fog. The reference architecture is hierarchically distributed according to the guidelines of the considered Multi-Level Elastic and Distributed DSP Framework (E2DF). In particular, the paper focuses on the control module of the framework and it proposes a Reinforcement Learning algorithm, combined with a heuristic approach to define multi-level policies for the placement and the management of the operators in a distributed stream processing system.

The paper "SLoPCloud: An Efficient Solution for Locality Problem in Peer-to-Peer Cloud Systems" by Gharib Mohammed et al. [2] is characterized by a P2P approach to distributed cloud systems. In particular, the paper tackles the problem of mapping the overlay network over the physical networking infrastructure. The mapping problem considers a hypercube-shaped overlay and the metric of interest to be optimized is the communication Round Trip Time (RTT). The final goal is to build a hypercubic overlay network with a locality-aware approach, thereby evaluating how this optimized infrastructure can provide a performance benefit.

The paper "Modeling and Evaluation of Power-Aware Software Rejuvenation in Cloud Systems" by Fakhrolmobasheri Sharifeh et al. [3] considers the problem of software rejuvenation in cloud systems. The paper focuses on the trade-off between the power consumption related to software ageing (typically due to failures and degradation in the Quality of Service, often reffered to as QoS degradation), and the power consumption related to a frequent rejuvenation process (that requires the shut down and restarting of parts of the system). The paper advocated for coordination between the rejuvenation process and the lifecycle of VMs (including migrations), which may significantly affect the power consumption. The authors propose a model, based on Stochastic Petri Networks, to determine when the rejuvenation should be performed. The experimental results confirm the benefit of this approach in terms of reduced failures and lower power consumption.

Finally, the paper "Reducing the Operational Cost of Cloud Data Centers through Renewable Energy" by Laganà Demetrio et al. [4] proposes a new framework, namely the EcoMultiCloud, that aims at increasing the sustainability of a multi-cloud infrastructure considering not only management techniques aimed at a power reduction, but also taking into account the utilization of renewable energy sources, as well as cooling costs. The considered architecture is scalable for a geographically distributed infrastructure, owing to the proposed hierarchical approach, where the main control mechanism for the management of the infrastructure is based on the decisions about VMs migrations (i.e., when, from where, and to where).

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Russo Russo, G.; Nardelli, M.; Cardellini, V.; Lo Presti, F. Multi-Level Elasticity for Wide-Area Data Streaming Systems: A Reinforcement Learning Approach. *Algorithms* **2018**, *11*, 134. [CrossRef]
2. Gharib, M.; Malekimajd, M.; Movaghar, A. SLoPCloud: An Efficient Solution for Locality Problem in Peer-to-Peer Cloud Systems. *Algorithms* **2018**, *11*, 150. [CrossRef]
3. Fakhrolmobasheri, S.; Ataie, E.; Movaghar, A. Modeling and Evaluation of Power-Aware Software Rejuvenation in Cloud Systems. *Algorithms* **2018**, *11*, 160. [CrossRef]
4. Laganà, D.; Mastroianni, C.; Meo, M.; Renga, D. Reducing the Operational Cost of Cloud Data Centers through Renewable Energy. *Algorithms* **2018**, *11*, 145. [CrossRef]