

Preface

In the new visions for cloud computing, accelerating business processes, increasing providers' profits, and reducing the price for end-users became imperative strategies of development. As computing and data storage needs grow exponentially, increasing the size of data centers brings important diseconomies of scale. Large-scale interconnected systems aim to aggregate and efficiently exploit the power of widely distributed resources. In this context, cloud computing has the potential to be the major solution for scalability, mobility, reliability, fault tolerance, and security for business and academic environments at the same time. Cloud computing became a suitable platform for big data processing and storage by embedding the major emerging trends at the research level and the technological level.

Resource management and task scheduling play an essential role, in cases where one is concerned with optimized use of resources. Ubiquitous networks are highly dynamic distributed systems and thus changes in overlay are frequent. On the other hand, the cloud systems are highly dynamic in their structure because the user requests must be respected as an agreement rule (SLA) and ensure QoS. When ubiquitous networks become clients for cloud systems, new algorithms for event and task scheduling and new methods for resource management should be designed in order to increase the performance of such systems. The adaptive methods used in this context are oriented toward: self-stabilizing, self-organizing, and autonomic systems; dynamic, adaptive, and machine learning-based distributed algorithms; fault tolerance, reliability, and availability, of distributed systems.

Following the success of the first ARMS-CC held in Paris in 2014, the second edition of the ARMS-CC workshop aimed to provide a venue for researchers, engineers, and practitioners involved in the development of new resource management methods, scheduling algorithms, and middleware tools for cloud computing. The goal is to offer an interactive and friendly yet professional forum for original research contributions describing novel ideas, groundbreaking results, or quantified system experiences in the context of the PODC Symposium. Papers presented at the ARMS-CC workshop can report and summarize on previous work, present early new results, or put forward new and outrageous ideas.

This volume contains the papers presented at ARMS-CC 2015: the Workshop on Adaptive Resource Management and Scheduling for Cloud Computing, held on July 19, 2015, in Donostia-San Sebastián, in conjunction with PODC 2015 (ACM Symposium on Principles of Distributed Computing). The papers of this volume have identified several important aspects of the problem addressed by ARMS-CC: self-* and autonomous cloud systems, cloud quality management and service level agreement (SLA), scalable computing, mobile cloud computing, cloud computing techniques for big data, high-performance cloud computing, resource management in big data platforms, scheduling algorithms for big data processing, cloud composition, federation,

bridging, and bursting, cloud resource virtualization and composition, load-balancing and co-allocation, fault tolerance, reliability, availability of cloud systems.

There were 24 submissions. Each submission was reviewed by at least one, and on average two, Program Committee members. The committee decided to accept 12 papers for publication in the LNCS proceedings and presentation at the workshop. The program also included one invited talk.

The invited talk paper was titled “Competitive Analysis of Task Scheduling Algorithms on a Fault-Prone Machine and the Impact of Resource Augmentation.” The authors analyzed the fault-tolerant properties of four popular scheduling algorithms: longest in system, shortest in system, longest processing time, and shortest processing time. They also investigated the effect of resource augmentation in their performance.

In the first paper, titled “Using Performance Forecasting to Accelerate Elasticity,” Paulo Moura et al. investigate the performance prediction of a service to dynamically adjust allocated resources. The authors present accurate predictive models that are based on workloads, using the universal scalability law.

Arani Bhattacharya et al. present a detailed model of the offloading problem incorporating lower energy usage on mobile devices, faster application execution, while operating even in unpredictable environments. Using simulation, the authors analyze the influence of these parameters on the offloading decision problem. The paper is titled “Parametric Analysis of Mobile Cloud Computing Frameworks Using Simulation Modeling.”

The third paper titled “Bandwidth-Aware Resource Optimization for SMT Processors” introduce the memory bandwidth and throughput behavior of various SPEC CPU2006 workloads in single threaded and simultaneous multi-threaded environment. Jasmine Madonna et al. show that once the bandwidth saturation occurs, the benefit of Simultaneous Multi-threaded (SMT) is more pronounced. They present a technique to leverage the memory-intensive nature of the workload towards a beneficial exploitation of SMT and throughput maximization.

A.J. Rubio-Montero et al., in the fourth paper titled “User-Guided Provisioning in Federated Clouds for Distributed Calculations,” present a new framework to effectively schedule distributed calculations in cloud federations. The focus of this work is represented by the mechanisms that allow users to consolidate their own resource provisioning in cloud federations.

In the fifth paper, “Compute on the Go: A Case of Mobile-Cloud Collaborative Computing Under Mobility,” Himadri Sekhar Paul et al. present two heuristics for offloading decisions for workflow modules, where some of the tasks in the workflow are marked as native. The authors present simulation results showing that the performance of the proposed heuristics is as good as that of MAUI, and yet less expensive.

Catalin Negru et al. present a study of the impact of the heterogeneity of virtual machines on datacenter power consumption in data-intensive applications. The experimental results, based on Amazon EC2 instances, show that the power consumption is proportional with the degree of heterogeneity.

In the seventh paper, “Implementing the Cloud Software to Data Approach for OpenStack Environments,” Lenos Vakanas et al. propose an innovative software for data service that allows virtual machines in the form of running instances or images to be migrated between OpenStack environments.

The next paper titled “Is Cloud Self-Organization Feasible?” discusses why cloud self-organization is not only desirable, but also critical for the future of cloud computing. Dan Marinescu et al. analyze major challenges and discuss practical principles for cloud self-organization, underlining the advantages of a self-organization model based on coalition formation and combinatorial auctions.

Beniamino Di Martino et al., in the ninth paper “Cloud Services Composition Through Cloud Patterns,” present the concept of cloud pattern that has emerged as a way to describe the composition and orchestration of cloud services in order to satisfy particular application requirements. They propose a methodology for the discovery and composition of cloud services, guided by cloud patterns.

The performance evaluation of Hadoop under failure when applying several schedulers is proposed by Shadi Ibrahim et al. in the next paper, titled “An Eye on the Elephant in the Wild: A Performance Evaluation of Hadoop’s Schedulers Under Failures.” The results reveal several drawbacks of Hadoop’s current mechanism in prioritizing failed tasks.

In the 11th paper, “Partitioning Graph Databases by Using Access Patterns,” Volkan Tüfkeçi and Can Özturan design and implement a framework that both partitions a graph database and provides a fully functional distributed graph database system. The access pattern is based on partitioning.

The last paper, “Cloud Search-Based Applications for Big Data — Challenges and Methodologies for Acceleration,” by George Suciu et al. analyzes cloud techniques that can be used for a faster search of large volumes of data. The authors propose a practical implementation by using the EXALEAD CloudView platform.

We express our gratitude and thanks to all members of the Technical Program Committee, for their work in finalizing the reviews on time, as well as the authors for submitting their papers to ARMS-CC 2015. We address our personal warm regards to the PODC 2015 organizers, especially the workshop chair, Elad Schiller, for all the support and advice offered during the workshop organization. The editors would like to thank Alfred Hofmann, Anna Kramer, and Elke Werner for the editorial assistance and excellent collaboration. We appreciate the support offered by the EasyChair system in handling the paper submission, review process, and communication with authors and reviewers. We thank them for their support in organizing the ARMS-CC Workshop.

July 2015

Florin Pop
Maria Potop-Butucaru

Adaptive Resource Management and Scheduling for
Cloud Computing

Second International Workshop, ARMS-CC 2015, Held in
Conjunction with ACM Symposium on Principles of
Distributed Computing, PODC 2015, Donostia-San
Sebastián, Spain, July 20, 2015, Revised Selected
Papers

Pop, F.; Potop-Butucaru, M. (Eds.)

2015, XII, 187 p. 77 illus. in color., Softcover

ISBN: 978-3-319-28447-7