

Guest Editorial: Automated Techniques for Migrating to the Cloud (II)

Ching-Hsien Hsu¹ and John Grundy²

¹Department of Computer Science and Information Engineering,
Chung Hua University, Taiwan
chh@chu.edu.tw

²Swinburne University of Technology, Australia
jgrundy@swin.edu.au

Introduction

The ever-growing popularity of cloud computing platforms provides hybrid architectures and dynamic intelligence for tomorrow's complex software applications. The "hybrid cloud" has emerged as an exciting new paradigm that includes private cloud, public cloud and community cloud infrastructure to provide IT services anytime and anywhere. It has recently also raised relevant interest in both the academic and the industrial research communities as a very promising application field for intelligent cloud computing and services. The interest in hybrid cloud environments has also been shown by many industrial and standardization efforts accomplished in the last years. Both academy and industry widely recognize the crucial role of software automation to facilitate the development of efficient and intelligent cloud services, by reducing implementation and deployment costs and time to market.

This two-part special issue is in response to the increasing convergence between automated software technologies and cloud computing; and the challenges and opportunities in this context. The second part of this special issue of Automated Software Engineering (ASE) covers different aspects of the problem, both from the theoretical to practical side. After a large open call, which received over two dozen submissions, an international editorial committee selected eight high quality research papers for publication. Each paper was reviewed by at least three expert reviewers. Many papers underwent several rounds of review before final acceptance. The research papers selected for this special issue represent key recent progress in the field. They include work on cloud architectures, service extraction, composition and deployment, big data intelligence, cloud databases, performance benchmarking, resource allocation and monitoring. All of these papers not only provide novel ideas and state-of-the-art techniques in the field, but also stimulate future research in the sustainable environment.

Cloud Architectures

With the recent paradigm shift of cloud computing, deployment of operating systems (OSs) onto a large-scale computer network is becoming necessary. Given that there are usually numerous nodes with various functions in a cloud computing system, it is usually necessary to deploy different operating systems onto different nodes. The paper by Kuen-Min Lee, Wei-Guang Teng, Jin-Neng Wu, Kuo-Ming Huang, Yao-Hsing Ko and Ting-Wei Hou entitled "Multicast and Customized Deployment of Large-Scale Operating Systems" proposes a new multicast deployment approach to significantly improve deployment efficiency. Their new approach can leverage existing configurations of its unicast counterpart. Key features of the proposed approach include support for a reliable multicast protocol, a heterogeneous infrastructure, and cloud hypervisor environments. To evaluate the feasibility of the proposed approach in practical applications, CentOS and Ubuntu are used when implementing the proposed deployment approach on several tens of nodes. Empirical studies show that both the required time for the entire distribution process, and the network bandwidth consumption are significantly reduced as compared with conventional unicast approaches.

Service Extraction and Composition

With a variety of providers large and small delivering a number of cloud-based services, cloud computing is evolving into an important service delivery infrastructure. One of the challenges in this evolution is how to provide the necessary fault handling for migrating long-running or computation-

intensive application services into shared open cloud infrastructures. The paper by Zhichun Jia, Rong Chen, Xing Xing, Junjie Xu and Yiwu Xie entitled “SFDCloud: Top-k Service Faults Diagnosis in Cloud Computing” presents a diagnostic architecture and a SFDCloud method to automatically diagnose the faults during the execution of web service composition in the cloud. The proposed diagnostic architecture decouples diagnosis service components, so that diagnosis resources can be shared and reused on the cloud platform. Scalability is a feature of the method and allows a diagnosis service to integrate functional components as a further means to deal with different complex web services. Experimental results show that the proposed method is effective in diagnosing faults in web service composition of various scales.

Big Data Services

With the size of datasets growing rapidly, large-scale data processing has become more and more important. The paper by Wu-Chun Chung, Hung-Pin Lin, Shih-Chang Chen, Mon-Fong Jiang and Yeh-Ching Chung entitled “JackHare: a framework for SQL to NoSQL translation using MapReduce” proposes the JackHare framework. This provides a comprehensive solution including SQL query compiler, JDBC driver and a systematical method using MapReduce for processing unstructured data in a NoSQL database. JackHare is developed based on Hadoop and HBase to store the data that originally resides in the relational database and designed the corresponding MapReduce methods based on the logic of SQL queries. To ease the learning curve of employing a NoSQL database like HBase, SQL users can import the JDBC driver into SQL client software to manipulate the large-scale data without being aware of the NoSQL database. The authors’ experimental results show that JackHare outperforms both Hive and MySQL in most cases and is also a scalable solution for various sizes of data.

Performance Benchmarking

In recent years, graphs have become an important workload in cloud systems due to the rapid increase of applications that produce data in the form of graphs. The paper by Miyuru Dayarathna and Toyotaro Suzumura entitled “Graph Database Benchmarking on Cloud Environments with XGDBench” introduces XGDBench, a graph database-benchmarking platform for cloud computing systems. This describes the architecture of XGDBench as well as how its graph traversal-based workloads are implemented. The XGDBench has been designed with the aim of creating an extensible platform for graph database benchmarking which makes it suitable for benchmarking future HPC systems. The authors extend the Yahoo! Cloud Serving Benchmark (YCSB) to the area of graph database benchmarking by creation of XGDBench. The authors evaluated the applicability of the MAG model for graph database benchmarking and also conducted a performance comparison of graph databases to demonstrate the capabilities of XGDBench for graph database benchmarking.

Conclusions

All of the above papers address either software technologies, cloud services models or propose novel application models in the various cloud and big data fields. They also will help to promote further related research and technology improvements in automated software technologies for cloud computing. This special issue serves as a landmark source for education, information, and reference for professors, researchers and graduate students interested in updating their knowledge about or active in cloud computing, software technologies and management, and novel application models for cloud computing platforms and big data.

The guest editors would like to express their sincere gratitude to Dr. Robert J. Hall (EiC, ASE), for giving us the opportunity to prepare this special issue. In addition, we are deeply indebted to the many reviewers for their professional effort, insight and hard work put into commenting on the selected articles, which reflect the essence of this special issue. Last but not least, we are grateful to all of the authors for their contributions and for undertaking the many revisions of their manuscripts, without which this special issue could not have been produced.