

# Guest Editors' Introduction: Special Issue on Data-Intensive Computing in the Clouds

Tevfik Kosar · Ioan Raicu

Published online: 24 March 2012  
© Springer Science+Business Media B.V. 2012

## 1 Introduction

Scientific as well as commercial applications generate increasingly large amounts of data, often referred as the “**data deluge**”, which necessitates collaboration and sharing among the national and international research and industrial organization. Simply purchasing high-capacity, high-performance storage and networking systems and adding them to the existing infrastructure of the collaborating institutions does not solve the underlying and highly challenging data handling problem. Users are often forced to spend a great deal of time and energy on solving basic data-handling issues, such as the physical location of data, how to access it, and/or how to move it to visualization and/or compute resources for further analysis. The management of data resources and

dataflow between the storage and compute resources is becoming the main bottleneck in front of efficient end-to-end computing. Analyzing, visualizing, and disseminating these large data sets has become a major challenge and data intensive computing is now considered as the “**fourth paradigm**” in scientific discovery after theoretical, experimental, and computational science.

This special issue on “**Data-Intensive Computing in the Clouds**” provides the scientific community a dedicated forum for introducing new research, development, and deployment efforts in running data-intensive computing workloads on Cloud Computing infrastructures. This special issue focuses on the use of cloud-based technologies to meet the new data intensive scientific challenges that are not well served by the current supercomputers, grids or compute-intensive clouds. We believe this issue will be an excellent place to help the community define the current state, determine future goals, and present architectures and services for future clouds supporting data intensive computing.

This special issue attracted 26 high quality submissions from around the world, out of which ten papers were accepted for publication. These papers cover a variety of topics, including MapReduce, security, storage, and I/O optimization in the Clouds. The remainder of this introduction serves as a guide to the content of this issue.

---

T. Kosar (✉)  
Department of Computer Science & Engineering,  
University at Buffalo (SUNY), Buffalo, NY, USA  
e-mail: tkosar@buffalo.edu

I. Raicu  
Department of Computer Science,  
Illinois Institute of Technology, Chicago, IL, USA  
e-mail: iraicu@cs.iit.edu

## 2 In This Issue

In “An Evaluation of the Cost and Performance of Scientific Workflows on Amazon EC2”, Gideon Juve et al. investigate ways in which data can be managed in the cloud, focusing on optimizing workflow systems. They ran experiments using three typical workflow applications (Montage, Broadband, and Epigenome) on Amazon’s EC2 cloud computing platform. They discuss the various storage and file systems used, describe the issues and problems they encountered deploying them on EC2, and analyze the resulting performance and cost of the workflows. Their results show that the cloud offers a convenient and flexible platform for deploying workflows with various storage systems. They found that the performance of storage systems such as GlusterFS is good on Amazon EC2. They also found that the cost of running workflows at small scale is not prohibitive, but the costs increases significantly when multiple virtual nodes are used in running larger workflows.

“Language and Runtime Support for Automatic Configuration and Deployment of Scientific Computing Software over Cloud Fabrics” by Chris Bunch et al. presents the design and implementation of Neptune, a simple, domain-specific language based on Ruby, that automates the configuration and deployment of popular scientific software frameworks over disparate cloud computing systems. Neptune integrates support for MPI, MapReduce, UPC, X10, StochKit, and others. They implemented Neptune as a software overlay for the AppScale cloud platform and extend AppScale with support for elasticity and hybrid execution for scientific computing applications. Neptune imposes no overhead on application execution yet significantly simplifies the application deployment process, enables portability across cloud systems, and promotes lock-in avoidance by specific cloud vendors.

Yanfeng Zhang et al. present in “iMapReduce: A Distributed Computing Framework for Iterative Computation” a framework that supports iterative processing. Iterative computation is pervasive in many applications such as data mining or social network analysis. These iterative applications typically involve massive data sets

containing at least millions or billions of data records. This poses demand of distributed computing frameworks for processing massive data sets on a cluster of machines, such as MapReduce. iMapReduce is an extension to MapReduce, and allows users to specify the iterative computation with the separated map and reduce functions; it also provides the support of automatic iterative processing without the need of a user-defined driver program. More importantly, iMapReduce significantly improves the performance of iterative implementations by (1) reducing the overhead of creating new MapReduce jobs repeatedly, (2) eliminating the shuffling of static data, and (3) allowing asynchronous execution of map tasks. We implement an iMapReduce prototype based on Apache Hadoop, and show that iMapReduce can achieve a factor of 1.1 to 5 speedup over Hadoop for implementing iterative algorithms.

In “VMMB: Virtual Machine Memory Balancing for Unmodified Operating Systems”, Eom Young Ik et al. present VMMB (Virtual Machine Memory Balancer), a novel mechanism used to dynamically monitor the memory demand and periodically re-balance the memory among VMs. VMMB accurately measures the memory demand with low performance overhead and effectively allocates memory based on the memory demand and the QoS requirement of each VM. Because VMMB does not require modifying guest kernel, it is applicable even to guest OS whose source code is not available. Their experimental results on synthetic and realistic workloads show that VMMB can dramatically improve performance of workload that suffers from memory with low performance overhead (below 1%).

“Cloud Resource Usage—Heavy Tailed Distributions Invalidating Traditional Capacity Planning Models” by Charles Loboz et al. presents an analysis of daily resource usage on three storage clusters for 478 days. For each day, the authors consider the distribution of resource usage by customer accounts for five different resources: storage used, storage transactions executed, internal network transfer, egress transfer and inter-data-center transfer. All distributions were highly imbalanced and most distribution samples have tails heavier than log-normal, exponential, or normal distributions. These findings spell significant

problems for most models assuming normality. Mathematically: Central Limit Theorem does not apply to power-law distributions—so the ‘averaging’ effect cannot be counted on to help with modeling using traditional approach. Operationally: very high volatility found means that the ‘capacity buffers’ need to be large, leading to wasted capacity. Other, administrative, means need to be applied to reduce that. Overall the distributions of resource usage in cloud storage are so far from normal, even after usual transformations, that traditional approach to forecasting and capacity planning needs to be reconsidered.

In “Performance Evaluation of Range Queries in Key Value Stores”, Pouria Pirzadeh et al. compare three widely used Key-Value store systems: Cassandra, HBase and Voldemort in terms of their support for different types of query workloads. According to the authors, unlike HBase and Cassandra that have built-in support for range queries, Voldemort does not support this type of queries via its available API. For this matter, practical techniques are presented on top of Voldemort to support range queries. The authors present a performance evaluation based on mixed query workloads, in the sense that they contain a combination of short and long range queries, beside other types of typical queries on key-value stores such as lookup and update. The authors show that there are trade-offs in the performance of the selected system and scheme, and the types of the query workloads that can be processed efficiently.

“Digital Preservation in Grids and Clouds: A Middleware Approach” by Peter Wittek and Sandor Daranyi describes a middleware system that enables a flexible choice between a grid and a cloud for ad-hoc computations that arise during the execution of a preservation workflow and also for archiving digital objects. According to the authors, the choice between different infrastructures remains open during the lifecycle of the archive, ensuring a smooth switch between different solutions to accommodate the changing requirements of the organization that needs its digital assets preserved. The authors also offer insights on the costs, running times, and organizational issues of cloud computing, proving that the cloud alternative is particularly attractive for smaller organiza-

tions without access to a grid or with limited IT infrastructure.

Dinh Tien Tuan Anh et al.’s “City on the Sky: Flexible, Secure Data Sharing on the Cloud” presents a comprehensive set of practical scenarios where data sharing needs to be enforced by methods such as aggregation, windowed frame, and value constrains. The authors observe that existing basic access control mechanisms do not provide adequate flexibility to support effective data sharing in a secure and controlled manner. They propose a framework for cloud that extends popular XACML model significantly by integrating flexible access control decisions and data access in a seamless fashion. They have prototyped the framework and deployed it on commercial cloud environment for experimental runs to test the efficacy of our approach and evaluate the performance of the implemented prototype.

In “Bayesian Cognitive Model in Scheduling Algorithm for Data Intensive Computing”, Wei Wang and Guosun Zeng present a trust mechanism-based task-scheduling model for data-intensive computing. Referring to the trust relationship models of social persons, the authors build a trust relationship among computing nodes, and evaluate the trustworthiness of nodes by utilizing the Bayesian cognitive method. They propose a Trust-Dynamic Level Scheduling (Trust-DLS) algorithm by integrating the trustworthiness of nodes into a Dynamic Level Scheduling (DLS) algorithm. They structure a benchmark to span a range of data intensive computing characteristics for evaluation of the proposed method. Their theoretical analysis and simulations prove that the Trust-DLS algorithm can efficiently meet the requirement of data intensive workloads in trust, sacrificing fewer time costs, and assuring the execution of tasks in a security way in large-scale data intensive computing environment.

“iCanCloud: A Flexible and Scalable Cloud Infrastructure Simulator” by Alberto Nunez et al. introduces and validates a novel simulator of cloud infrastructures, called iCanCloud, with remarkable features such as flexibility, scalability, performance and usability. This simulator is targeted to conduct large experiments, and provides a flexible and fully customizable global hypervisor for integrating any cloud brokering policy. It

reproduces the instance types provided by a given cloud infrastructure, and contains a user-friendly GUI for configuring and launching simulations, that goes from a single VM to large cloud computing systems composed of thousands of machines.

### 3 About the Editors

*Tevfik Kosar* is an associate professor of Computer Science & Engineering at University at Buffalo (SUNY). He holds a Ph.D. degree in Computer Science from University of Wisconsin-Madison under the guidance of Prof. Miron Livny. Dr. Kosar's main research interests lie in the cross-section of petascale distributed systems, eScience, Grids, Clouds, and collaborative computing with a focus on large-scale data-intensive distributed applications. He is the primary designer and developer of the Stork distributed data scheduling system which has been adopted by many national and international institutions, and the lead investigator of the state-wide PetaShare distributed storage network in Louisiana. He has published more than fifty academic papers in leading journals and conferences. Some of the awards received by Dr. Kosar include NSF CAREER Award (for his work on data-aware distributed computing), LSU Rainmaker Award, LSU Flagship Faculty Award, Baton Rouge Business Report's Top 40 Under 40 Award, 1012 Corridor's Young Scientist Award, College of Basic Science's Research Award, and CCT Faculty of the Year Award. Dr. Kosar's work on data intensive computing has been funded by NSF, DOE, ONR, DoEd, SURA, and Louisiana Board of Regents.

*Ioan Raicu* is an assistant professor in the Department of Computer Science at Illinois Institute of Technology, as well as a guest research faculty in the Math and Computer Science Division at Argonne National Laboratory. He is also the founder and director of the Data-Intensive Distributed Systems Laboratory at IIT. He has received the prestigious NSF CAREER award for his innovative work on distributed file systems for exascale computing. Prior to joining IIT, he was a NSF/CRA Computation Innovation Fellow at Northwestern University, and obtained his Ph.D. in CS from University of Chicago under Dr. Ian Foster's guidance in 2009. He is a 3-year award winner of the GSRP Fellowship from NASA Ames Research Center. His research work and interests are in the general area of distributed systems. He is particularly interested in resource management in large scale distributed systems with a focus on many-task computing, data intensive computing, cloud computing, grid computing, supercomputing, and many-core computing. Over the past decade, he has co-authored over 50 peer reviewed articles, book chapters, books, theses, and dissertations, which received over 2100 citations. His H-index is 19, G-Index is 45, and E-Index is 37. His work has been funded by NASA, DOE, and NSF. He has also founded and chaired several workshops, such as IEEE/ACM MTAGS, IEEE DataCloud/DataCloud-SC, and ACM ScienceCloud. He is on the editorial board of the Springer JoCCASA, as well as a guest editor for the IEEE TPDS, SPJ, and JoGC. He has had leadership roles in several high profile conferences, such as HPDC, CCGrid, Grid, eScience, and ICAC.