

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, Lancaster, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Zürich, Switzerland

John C. Mitchell

Stanford University, Stanford, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Dortmund, Germany

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbruecken, Germany

More information about this series at <http://www.springer.com/series/7409>

Jianfeng Zhan · Rui Han
Chuliang Weng (Eds.)

Big Data Benchmarks, Performance Optimization, and Emerging Hardware

4th and 5th Workshops, BPOE 2014
Salt Lake City, USA, March 1, 2014
and Hangzhou, China, September 5, 2014
Revised Selected Papers

Editors

Jianfeng Zhan
ICT, Chinese Academy of Sciences
Beijing
China

Chuliang Weng
Shannon (IT) Lab.
Huawei
China

Rui Han
ICT, Chinese Academy of Sciences
Beijing
China

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-13020-0 ISBN 978-3-319-13021-7 (eBook)
DOI 10.1007/978-3-319-13021-7

Library of Congress Control Number: 2014953862

Springer Cham Heidelberg New York Dordrecht London
© Springer International Publishing Switzerland 2014

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

[Springer International Publishing AG Switzerland] is part of Springer Science+Business Media
(www.springer.com)

Preface

Today, huge amounts of data are being collected in many areas, which create new opportunities to understand phenomena in meteorology, health, finance, and many other sectors. Big Data is considered precious assets of companies, organizations, and even nations. Turning such big data into real treasures requires the support of big data systems and platforms. However, the sheer volume of big data requires significant storage capacity, transmission bandwidth, computation, and power consumption. It is expected that systems with unprecedented scales can resolve the problems caused by varieties of big data with daunting volumes.

The complexity, diversity, frequently changed workloads, and rapid evolution of big data systems raise great challenges in big data benchmarking. Without big data benchmarks, it is very difficult for big data owners to make a decision on which system is best for meeting with their specific requirements. They also face challenges on how to optimize the systems and their solutions for specific or even comprehensive workloads. Meanwhile, researchers are also working on innovative data management systems, hardware architectures, operating systems, and programming systems to improve performance in dealing with big data.

This book includes papers from two workshops, which are the fourth and fifth workshops on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE-4 and BPOE-5). BPOE-4 (http://prof.ict.ac.cn/bpoe_4_asplos/) is co-located with ASPLOS 2014 (<http://www.cs.utah.edu/asplos14/>), a premier conference on architecture support for operating systems and programming systems. BPOE-5 (http://prof.ict.ac.cn/bpoe_5_vldb/) is co-located with VLDB 2014 (<http://www.vldb.org/2014/>), a premier conference on data management, database and information systems. Both workshops focus on architecture and system support for big data systems, aiming at bringing researchers and practitioners from data management, architecture, and systems research communities together to discuss the research issues at the intersection of these areas.

The call for papers for these two workshops attracted a number of high-quality international submissions. Within a rigorous process, in which each paper was reviewed by at least four experts, we selected 6 papers out of 12 submissions for inclusion in the BPOE-04 and 10 papers out of 18 submissions in the BPOE-05, respectively. In addition, several prestigious keynote speakers were invited, including Prof. Lizy Kurian John at University of Texas at Austin (<http://users.ece.utexas.edu/~ljohn/>) whose topic was “Big Data Workloads: An Architect’s Perspective,” Prof. Dhableswar K. (DK) Panda at Ohio State University (<http://www.cse.ohio-state.edu/~panda/>) whose topic was “Accelerating Big Data Processing with RDMA-Enhanced Apache Hadoop,” Prof. Christos Kozyrakis at Stanford University (<http://csl.stanford.edu/~christos/>) whose topic was “Resource Efficient Cloud Computing,” and Dr. Jeff Stuecheli from IBM (<http://www.linkedin.com/pub/jeff-stuecheli/2/664/a0a>) whose topic was “Power Technology For a Smarter Future.”

We are very grateful to the efforts of all authors related to writing, revising, and presenting their papers at BPOE workshops. Finally, we appreciate the indispensable support of BPOE Program Committees and thank their efforts and contributions in maintaining the high standards of the BPOE workshop.

August 2014

Jianfeng Zhan
Rui Han
Chuliang Weng

Organization

Program Co-chairs

Jianfeng Zhan	ICT, Chinese Academy of Sciences and University of Chinese Academy of Sciences, China
Chuliang Weng	Shannon (IT) Lab, China
Rui Han	ICT, Chinese Academy of Sciences, China

Steering Committee

Christos Kozyrakis	Stanford University, USA
Xiaofang Zhou	University of Queensland, Australia
Dhabaleswar K. (DK) Panda	Ohio State University, USA
Aoying Zhou	East China Normal University, China
Raghunath Nambiar	Cisco, USA
Lizy Kurian John	University of Texas at Austin, USA
Xiaoyong Du	Renmin University of China, China
Ippokratis Pandis	IBM Almaden Research Center, USA
Xueqi Cheng	ICT, Chinese Academy of Sciences, China
Bill Jia	Facebook, USA
Lidong Zhou	Microsoft Research Asia, China
H. Peter Hofstee	IBM Austin Research Laboratory, USA
Alexandros Labrinidis	University of Pittsburgh, USA
Cheng-Zhong Xu	Wayne State University, USA
Guang R. Gao	University of Delaware, USA
Yunquan Zhang	ICT, Chinese Academy of Sciences, China

Program Committee

Onur Mutlu	Carnegie Mellon University, USA
Xu Liu	Rice University, USA
Meikel Poess	Oracle Corporation, USA
Dejun Jiang	ICT, Chinese Academy of Sciences, China
Yueguo Chen	Renmin University, China
Rene Mueller	IBM, Almaden Research Center, USA
Xiaoyi Lu	Ohio State University, USA
Yongqiang He	Dropbox, USA
Edwin Sha	University of Texas at Dallas, USA
Kun Wang	IBM Research China, China
Rong Chen	Shanghai Jiao Tong University, China

Jens Teubner	TU Dortmund University, Germany
Yinliang Yue	ICT, Chinese Academy of Sciences, China
Mauricio Breternitz	AMD Research, China
Seetharami Seelam	IBM, USA
Zhenyu Guo	MSRA
Farhan Tauheed	EPFL, Switzerland
Gansha Wu	Intel, China
Bingsheng He	Nanyang Technological University, Singapore
Zhibin Yu	SIAT, Chinese Academy of Sciences, China
Lei Wang	ICT, Chinese Academy of Sciences, China
Yuanchun Zhou	CNIC, Chinese Academy of Sciences, China
Tilman Rabl	University of Toronto, Canada
Weijia Xu	TACC, University of Texas at Austin, USA
Mingyu Chen	ICT, Chinese Academy of Sciences, China
Jian Ouyang	Baidu, China
Wentao Qu	Google, USA
Guangyan Zhang	Tsinghua University, China
Cheqing Jin	East China Normal University, China
Jiuyang Tang	National University of Defense Technology, China
Farhan Tauheed	EPFL, Switzerland
Xiaoyu Zhang	CSHUST, USA
Lijie Wen	School of Software, Tsinghua University, China
Rong Chen	Shanghai Jiao Tong University, China

Contents

Topical Section Headings: Benchmarking

On Big Data Benchmarking	3
<i>Rui Han, Xiaoyi Lu, and Jiangtao Xu</i>	
A Micro-benchmark Suite for Evaluating Hadoop MapReduce on High-Performance Networks	19
<i>Dipti Shankar, Xiaoyi Lu, Md. Wasi-ur-Rahman, Nusrat Islam, and Dhabaleswar K. (DK) Panda</i>	
MemTest: A Novel Benchmark for In-memory Database	34
<i>Qiangqiang Kang, Cheqing Jin, Zhao Zhang, and Aoying Zhou</i>	
DSIMBench: A Benchmark for Microarray Data Using R	47
<i>Shicai Wang, Ioannis Pandis, Ibrahim Emam, David Johnson, Florian Guitton, Axel Oehmichen, and Yike Guo</i>	
A Benchmark to Evaluate Mobile Video Upload to Cloud Infrastructures. . . .	57
<i>Afsin Akdogan, Hien To, Seon Ho Kim, and Cyrus Shahabi</i>	
Benchmarking Replication and Consistency Strategies in Cloud Serving Databases: HBase and Cassandra	71
<i>Huajin Wang, Jianhui Li, Haiming Zhang, and Yuanchun Zhou</i>	

Topical Section Headings: Workload Characterization

I/O Characterization of Big Data Workloads in Data Centers	85
<i>Fengfeng Pan, Yinliang Yue, Jin Xiong, and Daxiang Hao</i>	
Characterizing Workload of Web Applications on Virtualized Servers	98
<i>Xiajun Wang, Song Huang, Song Fu, and Krishna Kavi</i>	

Topical Section Headings: Performance Optimization and Evaluation

Performance Benefits of DataMPI: A Case Study with BigDataBench	111
<i>Fan Liang, Chen Feng, Xiaoyi Lu, and Zhiwei Xu</i>	
InvarNet-X: A Comprehensive Invariant Based Approach for Performance Diagnosis in Big Data Platform	124
<i>Pengfei Chen, Yong Qi, Di Hou, and Huachong Sun</i>	

Tuning Hadoop Map Slot Value Using CPU Metric	141
<i>Kamal Kc and Vincent W. Freeh</i>	
A Study of SQL-on-Hadoop Systems	154
<i>Yueguo Chen, Xiongpai Qin, Haoqiong Bian, Jun Chen, Zhaoan Dong, Xiaoyong Du, Yanjie Gao, Dehai Liu, Jiaheng Lu, and Huijie Zhang</i>	
Predoop: Preempting Reduce Task for Job Execution Accelerations.	167
<i>Yi Liang, Yufeng Wang, Minglu Fan, Chen Zhang, and Yuqing Zhu</i>	
Record Placement Based on Data Skew Using Solid State Drives	181
<i>Jun Suzuki, Shivaram Venkataraman, Sameer Agarwal, Michael Franklin, and Ion Stoica</i>	
Efficient HTTP Based I/O on Very Large Datasets for High Performance Computing with the Libdaxiv Library	194
<i>Adrien Devresse and Fabrizio Furano</i>	
Topical Section Headings: Emerging Hardware	
Exploring Opportunities for Non-volatile Memories in Big Data Applications . . .	209
<i>Wei Wei, Dejun Jiang, Jin Xiong, and Mingyu Chen</i>	
Author Index	221