

Founding Editors

Gerhard Goos

Karlsruhe Institute of Technology, Karlsruhe, Germany

Juris Hartmanis

Cornell University, Ithaca, NY, USA

Editorial Board Members

Elisa Bertino

Purdue University, West Lafayette, IN, USA

Wen Gao

Peking University, Beijing, China

Bernhard Steffen

TU Dortmund University, Dortmund, Germany

Gerhard Woeginger 

RWTH Aachen, Aachen, Germany

Moti Yung

Columbia University, New York, NY, USA

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

Chen Zheng · Jianfeng Zhan (Eds.)

Benchmarking, Measuring, and Optimizing

First BenchCouncil International Symposium, Bench 2018
Seattle, WA, USA, December 10–13, 2018
Revised Selected Papers

Editors

Chen Zheng
Chinese Academy of Sciences
Beijing, China

Jianfeng Zhan
Chinese Academy of Sciences
Beijing, China

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-32812-2 ISBN 978-3-030-32813-9 (eBook)
<https://doi.org/10.1007/978-3-030-32813-9>

LNCS Sublibrary: SL3 – Information Systems and Applications, incl. Internet/Web, and HCI

© Springer Nature Switzerland AG 2019, corrected publication 2021

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, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

BenchCouncil: Benchmarking and Promoting Innovative Techniques

The past decade witnessed many innovative techniques as they evolved our economy and society. The innovative techniques in Datacenter, AI, dedicated processors, and IoT are especially changing our world. For example, new wave AI techniques have changed every part of our daily life, such as facial recognition, self-driving cars, robots, financial services, personal assistants, etc. In recent years, this trend has accelerated, and there is an increasing interest in generating values from these emerging technologies. However, benchmarking and promoting innovative techniques are the very first steps in making the innovation successful, which is the mission of the International Open Benchmark Council (BenchCouncil).

BenchCouncil is a non-profit research institute, which aims to promote multi-disciplinary benchmarking research and practice and foster collaboration and interaction between industry and academia. The major goals of the BenchCouncil are as follows:

- Establish and maintain a repository of benchmark specifications for quantitative system and algorithm evaluation and analysis
- Review, shepherd, and release open-source benchmark implementations
- Organize conferences, workshops, and teleconferences fostering the transfer of knowledge between industry and academia in the areas of benchmarking
- Build a testbed to verify and promote the innovative techniques
- Organize challenges and competition using released benchmarks
- Publish the performance numbers to evaluate competitive techniques using the benchmarks

In order to fulfill its mission, the BenchCouncil organized the First International Symposium on Benchmarking, Measuring, and Optimizing (Bench 2018) (<http://prof.ict.ac.cn/Bench18/>), which was co-located with IEEE BigData 2018 (<http://cci.drexel.edu/bigdata/bigdata2018/>)—an IEEE International Conference on Big Data. The symposium solicits papers that address hot topical issues in benchmarking, measuring, and optimizing systems. This book includes 15 papers from the Bench 2018 conference, and 5 invited papers which report on the state-of-the-art AI benchmarking research and engineering efforts of the BenchCouncil. In addition, 6 Benchmark proposals were presented at the conference, but were not included in this book.

The call for papers for Bench 2018 attracted a number of high-quality submissions. During a rigorous review process, each paper was reviewed by at least three experts. In addition, we invited three keynote speakers:

- Prof. Geoffrey Fox (Indiana University): “Big Data Benchmarking: Applications and Systems”
- Prof. Vijay Janapa Reddi (Harvard University): “MLPerf: the Vision Behind an ML Benchmark Suite for Measuring the Performance of ML Software Frameworks, ML Hardware Accelerators, and ML Cloud and Edge Platforms”
- Dr. Arne Berre (SINTEF Digital): “Benchmarking for Digital Platforms with Big Data, IoT, AI, Cloud, HPC, and Cyber Security”

During the conference, we had 6 benchmark proposals covering the topics in Big Data, AI, Datacenter, Scalable Transaction, and Database. In order to discuss the implementation of benchmark and BenchCouncil workgroup specification, we also held the BenchCouncil Open Meeting to format the working groups.

We are very grateful for the efforts of all authors in relation to writing, revising, and presenting their papers at Bench 2018. We appreciate the indispensable support of the Bench 2018 Program Committee and thank them for their efforts and contributions in maintaining the high standards of the Bench 2018 Symposium.

June 2019

Chen Zheng
Jianfeng Zhan

Organization

Program Chair

Jianfeng Zhan

ICT, Chinese Academy of Sciences, and University
of Chinese Academy of Sciences, China

Program Committee

Lizy K. John

The University of Texas at Austin, USA

Vijay Janapa Reddi

University of Texas at Austin and Google, USA

Wenguang Chen

Tsinghua University, China

Woongki Baek

UNIST, South Korea

Yunji Chen

ICT, Chinese Academy of Sciences, China

Piotr Luszczek

University of Tennessee, USA

Rui Hou

Institute of Information Engineering, Chinese Academy
of Sciences, China

Yueguo Chen

Renmin University of China, China

H. Peter Hofstee

IBM, USA

Cheqing Jin

East China Normal University, China

Trevor E. Carlson

National University of Singapore (NUS), Singapore

Jun-zhao Du

Xidian University, China

Hyogi Sim

Oak Ridge National Laboratory, USA

Weining Qian

East China Normal University, China

Zheng Cao

Alibaba, China

Matthias Nicola

IBM, Germany

Zhihui Du

Tsinghua University, China

Chunjie Luo

ICT, Chinese Academy of Sciences, China

Bin Ren

College of William and Mary, USA

Hua Chen

ICT, Chinese Academy of Sciences, China

Jiaquan Gao

Nanjing Normal University, China

Jiannan Ouyang

Facebook, USA

Xueming Si

Fudan University, China

YanJun Wu

Institute of Software Chinese Academy of Sciences,
China

Lucas Mello Schnorr

Federal University of Rio Grande do Sul (UFRGS),
Brazil

Yunquan Zhang

ICT, Chinese Academy of Sciences, China

Shengzhong Feng

Shenzhen Institutes of Advanced Technology,
Chinese Academy of Sciences, China

Zhen Jia

Princeton University, USA

Zhibin Yu

Shenzhen Institutes of Advanced Technology,
Chinese Academy of Sciences, China

Dapeng Wang	Institute of Software Chinese Academy of Sciences, China
Bo Wu	Colorado School of Mines, USA
Shaoliang Peng	National University of Defense Technology, China
Yong Qi	Xi'an Jiaotong University, China
Gwangsun Kim	Arm Inc., Austin, USA
Rui Ren	ICT, Chinese Academy of Sciences, China
Zujie Ren	Hangzhou Dianzi University, China
Xiaoyi Lu	The Ohio State University, USA
Lei Wang	ICT, Chinese Academy of Sciences, China
Jianhui Li	Computer Network Information Center, Chinese Academy of Sciences, China
Jianwu Wang	University of Maryland, USA
Kenli Li	Hunan University, China
Suzanne Rivoire	Sonoma State University, USA
Rui Han	ICT, Chinese Academy of Sciences, China
Jungang Xu	Chinese Academy of Sciences, China
Xu Liu	College of William and Mary, USA
Tong Wu	National Institute of Metrology, China
Li Zha	University of the Chinese Academy of Sciences, China
Lei Liu	ICT, Chinese Academy of Sciences, China
Chen Zheng	ICT, Chinese Academy of Sciences, China

Contents

AI Benchmarking

AI Bench: Towards Scalable and Comprehensive Datacenter AI Benchmarking	3
<i>Wanling Gao, Chunjie Luo, Lei Wang, Xingwang Xiong, Jianan Chen, Tianshu Hao, Zihan Jiang, Fanda Fan, Mengjia Du, Yunyou Huang, Fan Zhang, Xu Wen, Chen Zheng, Xiwen He, Jiahui Dai, Hainan Ye, Zheng Cao, Zhen Jia, Kent Zhan, Haoning Tang, Daoyi Zheng, Biwei Xie, Wei Li, Xiaoyu Wang, and Jianfeng Zhan</i>	
HPC AI500: A Benchmark Suite for HPC AI Systems	10
<i>Zihan Jiang, Wanling Gao, Lei Wang, Xingwang Xiong, Yuchen Zhang, Xu Wen, Chunjie Luo, Hainan Ye, Xiaoyi Lu, Yunquan Zhang, Shengzhong Feng, Kenli Li, Weijia Xu, and Jianfeng Zhan</i>	
Edge AI Bench: Towards Comprehensive End-to-End Edge Computing Benchmarking	23
<i>Tianshu Hao, Yunyou Huang, Xu Wen, Wanling Gao, Fan Zhang, Chen Zheng, Lei Wang, Hainan Ye, Kai Hwang, Zujie Ren, and Jianfeng Zhan</i>	
AIoT Bench: Towards Comprehensive Benchmarking Mobile and Embedded Device Intelligence	31
<i>Chunjie Luo, Fan Zhang, Cheng Huang, Xingwang Xiong, Jianan Chen, Lei Wang, Wanling Gao, Hainan Ye, Tong Wu, Runsong Zhou, and Jianfeng Zhan</i>	
A Survey on Deep Learning Benchmarks: Do We Still Need New Ones? . . .	36
<i>Qin Zhang, Li Zha, Jian Lin, Dandan Tu, Mingzhe Li, Fan Liang, Ren Wu, and Xiaoyi Lu</i>	

Cloud

Benchmarking VM Startup Time in the Cloud	53
<i>Samiha Islam Abrita, Moumita Sarker, Faheem Abrar, and Muhammad Abdullah Adnan</i>	
An Open Source Cloud-Based NoSQL and NewSQL Database Benchmarking Platform for IoT Data.	65
<i>Arjun Pandya, Chaitanya Kulkarni, Kunal Mali, and Jianwu Wang</i>	

Scalability Evaluation of Big Data Processing Services in Clouds	78
<i>Xin Zhou, Congfeng Jiang, Yeliang Qiu, Tiantian Fan, Yumei Wang, Liangbin Zhang, Jian Wan, and Weisong Shi</i>	
PAIE: A Personal Activity Intelligence Estimator in the Cloud	91
<i>Yingjie Shi, Fang Du, Yanyan Zhang, Zhi Li, and Tao Zhang</i>	
DCMIX: Generating Mixed Workloads for the Cloud Data Center	105
<i>Xingwang Xiong, Lei Wang, Wanling Gao, Rui Ren, Ke Liu, Chen Zheng, Yu Wen, and Yi Liang</i>	
Machine-Learning Based Spark and Hadoop Workload Classification Using Container Performance Patterns	118
<i>Mikhail Genkin, Frank Dehne, Pablo Navarro, and Siyu Zhou</i>	
Testing Raft-Replicated Database Systems	131
<i>Guohao Ding, Weining Qian, Peng Cai, Tianze Pang, and Qiong Zhao</i>	
Big Data	
Benchmarking for Transaction Processing Database Systems in Big Data Era	147
<i>Chunxi Zhang, Yuming Li, Rong Zhang, Weining Qian, and Aoying Zhou</i>	
UMDISW: A Universal Multi-Domain Intelligent Scientific Workflow Framework for the Whole Life Cycle of Scientific Data.	159
<i>Qi Sun, Yue Liu, Wenjie Tian, Yike Guo, and Bocheng Li</i>	
MidBench: Multimodel Industrial Big Data Benchmark.	172
<i>Yijian Cheng, Mengqian Cheng, Hao Ge, Yuhe Guo, Yuanzhe Hao, Xiaoguang Sun, Xiongpai Qin, Wei Lu, Yueguo Chen, and Xiaoyong Du</i>	
Modelling and Prediction	
Power Characterization of Memory Intensive Applications: Analysis and Implications	189
<i>Yeliang Qiu, Congfeng Jiang, Tiantian Fan, Yumei Wang, Liangbin Zhang, Jian Wan, and Weisong Shi</i>	
Multi-USVs Coordinated Detection in Marine Environment with Deep Reinforcement Learning	202
<i>Ruiying Li, Rui Wang, Xiaohui Hu, Kai Li, and Haichang Li</i>	
EC-Bench: Benchmarking Onload and Offload Erasure Coders on Modern Hardware Architectures	215
<i>Hayang Shi, Xiaoyi Lu, and Dhabaleswar K. Panda</i>	

Algorithm and Implementations

Benchmarking SpMV Methods on Many-Core Platforms	233
<i>Biwei Xie, Zhen Jia, and Yungang Bao</i>	
Benchmarking Parallel K-Means Cloud Type Clustering from Satellite Data	248
<i>Carlos Barajas, Pei Guo, Lipi Mukherjee, Susan Hoban, Jianwu Wang, Daeho Jin, Aryya Gangopadhyay, and Matthias K. Gobbert</i>	
Correction to: MiDBench: Multimodel Industrial Big Data Benchmark	C1
<i>Yijian Cheng, Mengqian Cheng, Hao Ge, Yuhe Guo, Yuanzhe Hao, Xiaoguang Sun, Xiongpai Qin, Wei Lu, Yueguo Chen, and Xiaoyong Du</i>	
Author Index	261