

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>

Wanling Gao · Jianfeng Zhan · Geoffrey Fox ·
Xiaoyi Lu · Dan Stanzione (Eds.)

Benchmarking, Measuring, and Optimizing

Second BenchCouncil International Symposium, Bench 2019
Denver, CO, USA, November 14–16, 2019
Revised Selected Papers

Editors

Wanling Gao 
Institute of Computing Technology,
Chinese Academy of Sciences
Beijing, China

Geoffrey Fox
School of Informatics, Computing,
and Engineering
Indiana University
Bloomington, IN, USA

Dan Stanzione
Texas Advanced Computing Center
The University of Texas at Austin
Austin, TX, USA

Jianfeng Zhan
Institute of Computing Technology,
Chinese Academy of Sciences
Beijing, China

Xiaoyi Lu
Department of Computer Science
and Engineering
The Ohio State University
Columbus, OH, USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-49555-8 ISBN 978-3-030-49556-5 (eBook)
<https://doi.org/10.1007/978-3-030-49556-5>

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

© Springer Nature Switzerland AG 2020

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. 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

Preface

This volume contains a selection of revised papers presented at Bench 2019: the Second BenchCouncil International Symposium, held in November, 2019, in Denver, CO, USA. To advance the state of the art of AI and other benchmarking, the International Open Benchmark Council (BenchCouncil) has three fundamental responsibilities. First, it releases influential benchmarks and index to prevent from chaotic competitions within AI, computer, finance, education, medical, and other technological industries. Second, it encourages data-driven, intelligence-inspired, and benchmark-based quantitative approaches to tackling multi-disciplinary challenges. Finally but not least, BenchCouncil incubates benchmark projects and hosts the BenchCouncil benchmark projects, and further encourages reliable and reproducible research using the BenchCouncil benchmark projects or incubator benchmark projects.

This year, BenchCouncil released four top level AI benchmarking projects, including AIBench – a scenario-distilling benchmarking methodology and an AI benchmark suite, HPC AI500 – a benchmark suite for HPC AI systems, Edge AIBench – an end-to-end edge AI benchmark Suite, and AIoTBench – an AI benchmark suite for benchmarking mobile and embedded device intelligence. Using AIBench as baseline, BenchCouncil hosted the 2019 AI System and Algorithm Challenge (<http://www.benchcouncil.org/competitions.html>) and organized the 2019 BenchCouncil International Symposium on Benchmarking, Measuring and Optimizing (Bench 2019) (<http://www.benchcouncil.org/bench19/index.html>). The 2019 AI Challenge consists of four challenge tracks: International AI System Challenge based on RISC-V, International AI System Challenge based on Cambricon Chip, International AI System Challenge based on X86 Platform, and International 3D Face Recognition Algorithm Challenge. The Bench 2019 symposium solicits papers that address hot topic issues in benchmarking, measuring, and optimizing systems. This book includes 31 regular papers from the Bench 2019 conference, which were selected from 79 submissions, yielding an acceptance rate of 39%. The tutorials about AI benchmarks for datacenter (AIBench), edge (Edge AIBench), and HPC (HPC AI500) were presented at the conference, but are not included in this book.

The call for papers for Bench 2019 attracted a number of high-quality submissions. During a rigorous review process, in which each paper was reviewed by at least three experts. In addition, we invited five keynote speakers, including Dr. Dan Stanzione from The University of Texas at Austin; Prof. Dhabaleswar K. (DK) Panda from The Ohio State University; Mr. Gilad Shainer from Mellanox; Prof. Geoffrey Fox from Indiana University; and Prof. Felix Wolf from TU Darmstadt in Germany. Bench 2019 also hosted five invited talks, including Dr. Zheng Cao from Alibaba; Dr. Dong Li from ICT, Chinese Academy of Sciences, and Seaway Technology Co., LTD.; Prof. Bo Wu from Colorado School of Mines; Dr. Weijia Xu from The University of Texas at Austin; and Dr. Gabriel Antoniu from Inria.

During the conference, BenchCouncil sponsored two awards to recognize important contributions in the area of benchmarking, measuring, and optimizing. The BenchCouncil Achievement Award recognizes a senior member who has made long-term contributions to benchmarking, measuring, and optimizing. Prof. Dr. Tony Hey, the Chief Data Scientist at Rutherford Appleton Laboratory STFC, was named the 2019 recipient of the International Open Benchmark Council (BenchCouncil) Achievement Award. The BenchCouncil Best Paper Award is recognizes a paper presented at the Bench conferences, which demonstrates potential impact on research and practice in benchmarking, measuring, and optimizing. Khaled Ibrahim, Samuel Williams, and Leonid Oliker from Lawrence Berkeley National Laboratory received the Bench 2019 Best Paper Award for their paper: “Performance Analysis of GPU Programming Models using the Roofline Scaling Trajectories.” In addition, 13 challenge teams from Georgia Institute of Technology, The Ohio State University, Google, etc., were honored with the 2019 AI Challenge Awards.

We are very grateful to the efforts of all authors related to writing, revising, and presenting their papers at Bench 2019 conference. We appreciate the indispensable support of the Bench 2019 Program Committee and thank them for their efforts and contributions in maintaining the high standards of the Bench 2019 symposium.

February 2020

Wanling Gao
Jianfeng Zhan

Organization

General Chairs

Dan Stanzione	Texas Advanced Computing Center, The University of Texas at Austin, USA
Xiaoyi Lu	The Ohio State University, USA

Program Chairs

Jianfeng Zhan	ICT, Chinese Academy of Sciences, University of Chinese Academy of Sciences, China
Geoffrey Fox	Indiana University, USA

Program Committee

Haiying Shen	University of Virginia, USA
Woongki Baek	UNIST, South Korea
Zheng Cao	Alibaba, China
Piotr Luszczek	University of Tennessee Knoxville, USA
Khaled Ibrahim	Lawrence Berkeley National Lab, USA
Li Zha	ICT, Chinese Academy of Sciences, China
Lei Wang	ICT, Chinese Academy of Sciences, China
Vladimir Getov	University of Westminster, UK
Gwangsun Kim	POSTECH, South Korea
Hyogi Sim	Oak Ridge National Laboratory, USA
Zhen Jia	Princeton University, USA
Weijia Xu	The University of Texas at Austin, USA
Bin Ren	Pacific Northwest National Laboratory, USA
Nikhil Jain	Lawrence Livermore National Laboratory, USA
Lucas Mello Schnorr	UFRGS, Brazil
Shantenu Jha	Rutgers University, USA
Bo Wu	Colorado School of Mines, USA
Jungang Xu	College of Computer and Control Engineering, University of Chinese Academy of Sciences, China
Saliya Ekanayake	Virginia Tech, USA
Zujie Ren	Zhejiang University, China
Ryan E. Grant	Sandia National Laboratories, USA
Zhihui Du	Tsinghua University, China
Guang R Gao	Computer Architecture and Parallel Systems Laboratory, University of Delaware, USA

Award Committee

Lizy Kurian John

Dhabaleswar K. (DK) Panda

Geoffrey Fox

Jianfeng Zhan

The University of Texas at Austin, USA

The Ohio State University, USA

Indiana University, USA

ICT, Chinese Academy of Sciences,

University of Chinese Academy of Sciences, China

Keynote Speech and Invited Talks

BenchCouncil Achievement Award

Tony Hey

Rutherford Appleton Laboratory STFC



Tony Hey has a doctorate in particle physics from the University of Oxford. After a career in physics that included research positions at Caltech, MIT and CERN, and a professorship at the University of Southampton, he became interested in parallel computing and moved into computer science. His group was one of the first to build and explore the development of parallel software for message-passing distributed memory computers. He was one of the authors of the first draft of the MPI message-passing standard. Tony led the UK eScience initiative in 2001 before joining Microsoft in 2005 as Vice-President for Technical Computing. He returned to work in the UK in 2015 as Chief Data Scientist at the Rutherford Appleton Laboratory and leads the ‘Scientific Machine Learning’ group. Tony is a fellow of the Association for Computing Machinery, the American Association for the Advancement of Science, and the Royal Academy of Engineering.

Prof. Hey was honored for foundational contributions to distributed memory parallel machines, message-passing systems, and AI for sciences. Prof. Hey devised the first parallel benchmark suite – the ‘Genesis’ benchmarks – for performance evaluation of distributed memory parallel machines, and recently launches a large-scale data science benchmark.

Benchmarking Supercomputers in the Post-Moore Era

Dan Stanzione

The University of Texas at Austin

Abstract: In this talk, we will cover the increasing gaps between headline performance and application performance on Frontera and the last several generations of TACC supercomputers. We will also discuss the challenges of developing a new benchmark suite for the upcoming Leadership-Class Computing Facility, and solicit community input on capability benchmarks.



Bio: Dr. Dan Stanzione, Associate Vice President for Research at The University of Texas at Austin since 2018 and Executive Director of the Texas Advanced Computing Center (TACC) since 2014, is a nationally recognized leader in high performance computing. He is the principal investigator (PI) for a National Science Foundation (NSF) grant to deploy Frontera, which is the fastest supercomputer at any U.S. university. Stanzione is also the PI of TACC's Stampede2 and Wrangler systems, supercomputers for high performance computing and for data-focused applications, respectively. For six years he was co-PI of CyVerse, a large-scale NSF life sciences cyberinfrastructure. Stanzione was also a co-PI for TACC's Ranger and Lonestar supercomputers, large-scale NSF systems previously deployed at UT Austin. Stanzione received his bachelor's degree in electrical engineering and his master's degree and doctorate in computer engineering from Clemson University.

Benchmarks and Middleware for Designing Convergent HPC, Big Data and Deep Learning Software Stacks for Exascale Systems

Dhabaleswar K. (DK) Panda

The Ohio State University

Abstract: This talk will focus on challenges in designing benchmarks and middleware for convergent HPC, Deep Learning, and Big Data Analytics Software stacks for Exascale systems with millions of processors and accelerators. For the HPC domain, we will discuss about the OSU Micro-Benchmarks (OMB) Suite and associated middleware for designing runtime environments for MPI+X programming models by taking into account support for multi-core systems (x86, OpenPOWER, and ARM), high-performance networks, and GPGPUs (including GPUDirect RDMA). Features and sample performance numbers from the MVAPICH2 libraries (<http://mvapich.cse.ohio-state.edu>) will be presented. An overview of RDMA-based designs for Hadoop (HDFS, MapReduce, RPC, and HBase), Spark, and Memcached, together with the OSU HiBD benchmarks (<http://hibd.cse.ohio-state.edu>) will be presented for Big Data Analytics. For the Deep Learning domain, we will focus on a set of different benchmarks and profiling tools to deliver scalable DNN training with Horovod and TensorFlow using MVAPICH2-GDR MPI library (<http://hidl.cse.ohio-state.edu>).



Bio: Dhabaleswar K. (DK) Panda is a Professor and University Distinguished Scholar of Computer Science and Engineering at The Ohio State University. He has published over 450 papers in the area of high-end computing and networking. The MVAPICH2 (High Performance MPI and PGAS over InfiniBand, Omni-Path, iWARP, and RoCE) libraries, designed and developed by his research group (<http://mvapich.cse.ohio-state.edu>), are currently being used by more than 3,025 organizations worldwide (in 89 countries). More than 600,000 downloads of this software have taken place from the project's site. This software is empowering several InfiniBand clusters (including the 3rd, 5th, 8th, 15th, 16th, 19th, and 31st ranked ones) in the TOP500 list. The RDMA packages for Apache Spark, Apache Hadoop, and Memcached together with OSU HiBD benchmarks from his group (<http://hibd.cse.ohio-state.edu>) are also publicly available. These libraries are


currently being used by more than 315 organizations in 35 countries. More than 31,300 downloads of these libraries have taken place. High-performance and scalable versions of the Caffe and TensorFlow framework are available from <https://hidl.cse.ohio-state.edu>. Prof. Panda is an IEEE Fellow. More details about Prof. Panda are available at <http://www.cse.ohio-state.edu/panda>.

InfiniBand In-Network Computing Technology for Scalable HPC/AI

Gilad Shainer

Mellanox

Abstract: The ever-increasing demands for higher computation performance drive the creation of new datacenter accelerators and processing units. Previously CPUs and GPUs were the main sources for compute power. The exponential increase in data volume and in problems complexity, drove the creation of a new processing unit the I/O processing unit or IPU. IPUs are interconnect elements that include In-Network Computing engines, engines that can participate in the application run time, and analyze application data as it being transferred within the data center, or at the edge. The combination of CPUs, GPUs, and IPUs, creates the next generation of data center and edge computing architectures. The first generations of IPUs are already in use in leading HPC and Deep learning data centers, have been integrated into multiple MPI frameworks, NVIDIA NCCL, Charm++, and others, and have demonstrated accelerate performance by nearly 10X.



Bio: Gilad Shainer serves as Mellanox's Senior Vice President of Marketing, focusing on high-performance computing. Mr. Shainer joined Mellanox in 2001 as a design engineer and later served in senior marketing management roles since 2005. Mr. Shainer serves as the chairman of the HPC-AI Advisory Council organization, he serves as the president of UCF and CCIX consortiums, a board member in the OpenCAPI and OpenFabrics organizations, a member of IBTA and contributor to the PCISIG PCI-X and PCIe specifications. Mr. Shainer holds multiple patents in the field of high-speed networking. He is a recipient of 2015 R&D100 award for his contribution to the CORE-Direct In-Network Computing technology and the 2019 R&D100 award for his contribution to the UCX technology. Gilad Shainer holds MSc degree and BSc degree in Electrical Engineering from the Technion Institute of Technology in Israel.

Benchmarking Perspectives on Emerging HPC Workloads

Geoffrey Fox

Indiana University



Bio: Geoffrey Charles Fox (<https://www.engineering.indiana.edu/>, <http://www.dsc.soic.indiana.edu/>, gcf@indiana.edu). Fox received a PhD in Theoretical Physics from Cambridge University where he was Senior Wrangler. He is now a distinguished Professor of Engineering, Computing, and Physics at Indiana University where he is director of the Digital Science Center. He previously held positions at Caltech, Syracuse University, and Florida State University after being a postdoc at the Institute for Advanced Study at Princeton, Lawrence Berkeley Laboratory, and Peterhouse College Cambridge. He has supervised the PhD of 73 students and published around 1,300 papers (over 500 with at least 10 citations) in physics and computing with an hindex of 78 and over 35,000 citations. He is a Fellow of APS (Physics) and ACM (Computing) and works on the interdisciplinary interface between computing and applications. Current work is in Biology, Pathology, Sensor Clouds and Ice-sheet Science, Image processing, Deep Learning, and Particle Physics. His architecture work is built around High-performance Computing enhanced Software Defined Big Data Systems on Clouds and Clusters. The analytics focuses on scalable parallel machine learning. He is an expert on streaming data and robot-cloud interactions. He is involved in several projects to enhance the capabilities of Minority Serving Institutions. He has experience in online education and its use in MOOCs for areas like Data and Computational Science.

Lightweight Requirements Engineering for Exascale Co-design

Felix Wolf

Department of Computer Science of TU Darmstadt, Germany

Abstract: Given the tremendous cost of an exascale system, its architecture must match the requirements of the applications it is supposed to run as precisely as possible. Conversely, applications must be designed such that building an appropriate system becomes feasible, motivating the idea of co-design. In this process, a fundamental aspect of the application requirements are the rates at which the demands for different resources grow as a code is scaled to a larger machine. However, if the anticipated scale exceeds the size of available platforms this demand can no longer be measured. This is clearly the case when designing an exascale system. Moreover, creating analytical models to predict these requirements is often too laborious especially when the number and complexity of target applications is high. In this paper, we show how automated performance modeling can be used to quickly predict application requirements for varying scales and problem sizes.



Bio: Felix Wolf is Full Professor at the Department of Computer Science of TU Darmstadt in Germany, where he leads the Laboratory for Parallel Programming. He works on methods, tools, and algorithms that support the development and deployment of parallel software systems in various stages of their life cycle. Prof. Wolf received his PhD degree from RWTH Aachen University in 2003. After working more than two years as a postdoc at the Innovative Computing Laboratory of the University of Tennessee, he was appointed research group leader at Jülich Supercomputing Centre. Between 2009 and 2015, he was head of the Laboratory for Parallel Programming at the German Research School for Simulation Sciences in Aachen and Full Professor at RWTH Aachen University. Prof. Wolf has published more than a hundred refereed articles on parallel computing, several of which have received awards.

FloraBench: An End-to-End Application Benchmark Suite for Datacenter

Zheng Cao

Alibaba

Abstract: The topic is FloraBench: an end-to-end application benchmark suite for datacenter. This talk abstracts the realistic application scenario of Alibaba and provides an application benchmark for datacenter computing – FloraBench. This Benchmark aims to identify the characteristics and bottlenecks of business E-commerce applications, and further optimize the performance of large-scale clusters.




Bio: Dr. Zheng Cao is a Senior Staff Engineer of the Alibaba Group and leads the architecture team of Alibaba Infrastructure Service BU. He received his PhD from ICT, Chinese Academy of Sciences. Before joining Alibaba Group, he served as a Professor at the Institute of Computing Technology, CAS and was one of the core architects of Dawning 5000, Dawning 6000 (ranked the 2nd in the TOP500 list), and Dawning 7000 supercomputer systems. His team is working on the Alibaba's workload analysis, software-hardware codesign, and datacenter architecture. He is a member of Advanced Computing Expert Committee, and Blockchain Expert Committee of China Computer Federation.

Towards Benchmarking AIOT Device Based on MCU

Dong Li

ICT, Chinese Academy of Sciences Seaway Technology Co., Ltd.

Abstract: The topic is “Towards Benchmarking AIOT Device based on MCU”. This talk introduces MCU-based AIOT device and discusses the benchmarking requirements and goals. Seaway RTOS for AIOT devices provide KB-level Seaway RTOS kernel, KB-level runtime, and KB-level EdgeStack, and allow only one application for the whole end-Edge-cloud system.




Bio: Dr. Dong Li is an Associate Professor at Wireless Sensor Network Laboratory, ICT, Chinese Academy of Sciences and Seaway Technology Co., Ltd.

Harmonizing High-Level Abstraction and High Performance for Graph Mining

Bo Wu

Colorado School of Mines

Abstract: Graph mining algorithms that aim at identifying structural patterns in graphs are typically more complex than graph computation algorithms such as breadth first search. Researchers have implemented several systems with high-level and flexible interfaces customized for tackling graph mining problems. However, we found that for triangle counting, one of the simplest graph mining problems, such systems can be several times slower than a single-threaded implementation of a straightforward algorithm. In this talk, I will reveal the root causes of the severe inefficiency of state-of-the-art graph mining systems and the challenges to address the performance problems. I will describe AutoMine, a system we developed to automatically generate both specialized algorithms and high-performance low-level code for arbitrary patterns.



Bio: Bo Wu is an Associate Professor in the Department of Computer Science at Colorado School of Mines. His research focuses on leveraging compiler and runtime techniques to build efficient software systems for large-scale graph analytics and machine learning applications on heterogeneous platforms. He received the Best Paper Award at SC'15, an NSF CRII Award, an NSF Early Career Award, and an NSF SPX Award.

Deep Learning on HPC: Performance Factors and Lessons Learned

Weijia Xu

The University of Texas at Austin

Abstract: In this talk, we report several ongoing efforts for deploying and running deep learning applications using high performance computing clusters at Texas Advanced Computing Center. From both lessons learned through practices and designed experiments, we discuss several factors affecting the deep learning performances, both accuracy and execution time, at various stages of analysis pipeline from low level data storage to high level deep learning framework. The talk will end with discussions and future outlooks on development, deployment, and benchmark deep learning applications at scale.



Bio: Dr. Weijia Xu is a research scientist and lead the Scalable Computational Intelligence group at Texas Advanced Computing Center at The University of Texas at Austin. He received his PhD from Computer Science Department at UT Austin and has been an experienced data scientist. Dr. Xu's main research interest is to enable data-driven discoveries through developing new computational methods and applications that facilitate the data-to-knowledge transfer process. Dr. Xu leads the group that supports large scale data driven analysis and machine learning applications using computing resources at TACC. His projects have been funded through various federal and state agencies including NIH, NSF, City of Austin, and USDA. He has served in Program Committees for several workshops and conferences in Big Data, Cloud Computing, and HPC areas.

Towards a Methodology for Benchmarking Edge Processing Frameworks

Gabriel Antoniu

Inria

Abstract: With the spectacular growth of the Internet of Things, edge processing emerged as a relevant means to offload data processing and analytics from centralized Clouds to the devices that serve as data sources (often provided with some processing capabilities). While a large plethora of frameworks for edge processing were recently proposed, the distributed systems community has no clear means today to discriminate between them. Some preliminary surveys exist, focusing on a feature-based comparison. We claim that a step further is needed, to enable a performance-based comparison. To this purpose, the definition of a benchmark is a necessity. In this talk, we make this step by discussing the definition of a methodology for benchmarking Edge processing frameworks.



Bio: Dr. Gabriel Antoniu is a Senior Research Scientist at Inria, Rennes. He leads the KerData research team, focusing on storage and I/O management for Big Data processing on scalable infrastructures (clouds, HPC systems). His main current interests regard HPC-Big Data convergence for data storage and processing aspects. He currently serves as Vice Executive Director of JLESC – Joint Inria-Illinois-ANL-BSC-JSC-RIKEN/AICS Laboratory for Extreme-Scale Computing on behalf of Inria. He received his PhD degree in Computer Science in 2001 from ENS Lyon. He leads several international projects in partnership with Microsoft Research, IBM, Argonne National Lab, the University of Illinois at Urbana Champaign, and Huawei. He served as program chair for the IEEE Cluster conference in 2014 and 2017 and regularly serves as a Program Committee member of major conferences in the area of HPC, Cloud Computing and Big Data (SC, HPDC, CCGRID, Cluster, Big Data, etc.). He has acted as advisor for 19 PhD theses and has co-authored over 140 international publications in the aforementioned areas.

Contents

Best Paper Session

Performance Analysis of GPU Programming Models Using the Roofline Scaling Trajectories	3
<i>Khaled Z. Ibrahim, Samuel Williams, and Leonid Oliker</i>	
GraphBench: A Benchmark Suite for Graph Computing Systems	20
<i>Lei Wang and Minghe Yu</i>	
Early Experience in Benchmarking Edge AI Processors with Object Detection Workloads	32
<i>Yujie Hui, Jeffrey Lien, and Xiaoyi Lu</i>	

AI Challenges on Cambricon Using AIBench

XDN: Towards Efficient Inference of Residual Neural Networks on Cambricon Chips	51
<i>Guangli Li, Xueying Wang, Xiu Ma, Lei Liu, and Xiaobing Feng</i>	
Performance Analysis of Cambricon MLU100	57
<i>Jiansong Li and Zihan Jiang</i>	
Exploring the Performance Bound of Cambricon Accelerator in End-to-End Inference Scenario	67
<i>Yifan Wang, Chundian Li, and Chen Zeng</i>	
Improve Image Classification by Convolutional Network on Cambricon	75
<i>Peng He, Ge Chen, Kai Deng, Ping Yao, and Li Fu</i>	

AI Challenges on RISC-V Using AIBench

RVTensor: A Light-Weight Neural Network Inference Framework Based on the RISC-V Architecture	85
<i>Pengpeng Hou, Jiageng Yu, Yuxia Miao, Yang Tai, Yanjun Wu, and Chen Zhao</i>	
AIRV: Enabling Deep Learning Inference on RISC-V	91
<i>Yangyang Kong</i>	

AI Challenges on X86 Using AIBench

PSL: Exploiting Parallelism, Sparsity and Locality to Accelerate Matrix Factorization on x86 Platforms 101
Weixin Deng, Pengyu Wang, Jing Wang, Chao Li, and Minyi Guo

The Implementation and Optimization of Matrix Decomposition Based Collaborative Filtering Task on X86 Platform 110
Tianshu Hao and Ziping Zheng

An Efficient Implementation of the ALS-WR Algorithm on x86 CPUs 116
Maosen Chen, Tun Chen, and Qianyun Chen

Accelerating Parallel ALS for Collaborative Filtering on Hadoop 123
Yi Liang, Shaokang Zeng, Yande Liang, and Kaizhong Chen

AI Challenges on 3D Face Recognition Using AIBench

Improving RGB-D Face Recognition via Transfer Learning from a Pretrained 2D Network 141
Xingwang Xiong, Xu Wen, and Cheng Huang

An Implementation of ResNet on the Classification of RGB-D Images 149
Tongyan Gong and Huiqian Niu

Utilization of Resnet in RGB-D Facial Recognition Problems. 156
Xi Xiong

Benchmark

Building the DataBench Workflow and Architecture 165
Todor Ivanov, Timo Eichhorn, Arne Jørgen Berre, Tomás Pariente Lobo, Ivan Martinez Rodriguez, Ricardo Ruiz Saiz, Barbara Pernici, and Chiara Francalanci

Benchmarking Solvers for the One Dimensional Cubic Nonlinear Klein Gordon Equation on a Single Core 172
B. K. Muite and Samar Aseeri

AI

Deep Reinforcement Learning for Auto-optimization of I/O Accelerator Parameters. 187
Trong-Ton Pham and Dennis Mintah Djan

Causal Learning in Question Quality Improvement 204
Yichuan Li, Ruocheng Guo, Weiying Wang, and Huan Liu

SparkAIBench: A Benchmark to Generate AI Workloads on Spark	215
<i>Zifeng Liu, Xiaojiang Zuo, Zeqing Li, and Rui Han</i>	

Big Data

Benchmarking Database Ingestion Ability with Real-Time Big Astronomical Data	225
<i>Qing Tang, Chen Yang, Xiaofeng Meng, and Zhihui Du</i>	

A Practical Data Repository for Causal Learning with Big Data	234
<i>Lu Cheng, Ruocheng Guo, Raha Moraffah, K. Selçuk Candan, Adrienne Raglin, and Huan Liu</i>	

Datacenter

LCIO: Large Scale Filesystem Aging.	251
<i>Matthew Bachstein, Feiyi Wang, and Sarp Oral</i>	

BOPS, A New Computation-Centric Metric for Datacenter Computing	262
<i>Lei Wang, Wanling Gao, Kaiyong Yang, and Zihan Jiang</i>	

Anomaly Analysis and Diagnosis for Co-located Datacenter Workloads in the Alibaba Cluster	278
<i>Rui Ren, Jinheng Li, Lei Wang, Yan Yin, and Zheng Cao</i>	

Performance Analysis and Workload Generator

SSH-Backed API Performance Case Study	295
<i>Anagha Jamthe, Mike Packard, Joe Stubbs, Gilbert Curbelo III, Roseline Shapi, and Elias Chalhoub</i>	

NTP: A Neural Net Topology Profiler	306
<i>Pravin Chandran, Raghavendra Bhat, Juby Jose, Viswanath Dibbur, and Prakash Sirra Ajith</i>	

MCC: A Predictable and Scalable Massive Client Load Generator.	319
<i>Wenqing Wu, Xiao Feng, Wenli Zhang, and Mingyu Chen</i>	

Scientific Computing and Metrology

Apache Spark Streaming, Kafka and HarmonicIO: A Performance Benchmark and Architecture Comparison for Enterprise and Scientific Computing	335
<i>Ben Blamey, Andreas Hellander, and Salman Toor</i>	

Benchmark Researches from the Perspective of Metrology 348
 Kun Yang, Tong Wu, Qingfei Shen, Weiqun Cui, and Guichun Zhang

Author Index 361