

2023 IEEE International Symposium on Workload Characterization (IISWC 2023)

**Ghent, Belgium
1 – 3 October 2023**



**IEEE Catalog Number: CFP23236-POD
ISBN: 979-8-3503-0318-6**

**Copyright © 2023 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP23236-POD
ISBN (Print-On-Demand):	979-8-3503-0318-6
ISBN (Online):	979-8-3503-0317-9
ISSN:	2835-222X

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

2023 IEEE International Symposium on Workload Characterization (IISWC) IISWC 2023

Table of Contents

Message from General Chair	ix
Message from Program Chairs	x
Organizing Committee	xi
Program Committee	xiii
Steering Committee	xv
Sponsors	xvi
Keynote Addresses	xvii

Best Paper Nominees

LumiBench: A Benchmark Suite for Hardware Ray Tracing	1
<i>Lufei Liu (University of British Columbia, Canada), Mohammadreza Saed (University of British Columbia, Canada), Yuan Hsi Chou (University of British Columbia, Canada), Davit Grigoryan (University of British Columbia, Canada), Tyler Nowicki (Huawei Technologies, Canada), and Tor M. Aamodt (University of British Columbia, Canada)</i>	
Vector-Processing for Mobile Devices: Benchmark and Analysis	15
<i>Alireza Khadem (University of Michigan, USA), Daichi Fujiki (Keio University, Japan), Nishil Talati (University of Michigan, USA), Scott Mahlke (University of Michigan; Nvidia Research, USA), and Reetuparna Das (University of Michigan, USA)</i>	
Making Sense of Using a SmartNIC to Reduce Datacenter Tax from SLO and TCO Perspectives	28
<i>Jinghan Huang (University of Illinois at Urbana-Champaign, USA), Jiaqi Lou (University of Illinois at Urbana-Champaign, USA), Yan Sun (University of Illinois at Urbana-Champaign, USA), Tianchen Wang (University of Illinois at Urbana-Champaign, USA), Eun Kyung Lee (IBM Research, USA), and Nam Sung Kim (University of Illinois at Urbana-Champaign, USA)</i>	
Balancing Accuracy and Evaluation Overhead in Simulation Point Selection	43
<i>Björn Gottschall (Norwegian University of Science and Technology (NTNU), Norway), Silvio Campelo de Santana (Norwegian University of Science and Technology (NTNU), Norway), and Magnus Jahre (Norwegian University of Science and Technology (NTNU), Norway)</i>	

Media and IoT

Do Video Encoding Workloads Stress the Microarchitecture?	54
<i>Steffen Jensen (University of Texas at Austin), Jaekyu Lee (Arm), Dam Sunwoo (Arm), Matt Horsnell (Arm), Matthew Siggs (British Columbia Institute of Technology), Jeeho Ryoo (British Columbia Institute of Technology), and Lizy John (University of Texas at Austin)</i>	
ESS: Repeatable Evaluation of Energy Harvesting Subsystems for Industry-Grade IoT Platforms	65
<i>Fatemeh Ghasemi (Norwegian University of Science and Technology (NTNU)), Lukas Liedtke (Norwegian University of Science and Technology (NTNU)), and Magnus Jahre (Norwegian University of Science and Technology (NTNU))</i>	
PBVR: Physically Based Rendering in Virtual Reality	77
<i>Yavuz Selim (North Carolina State University, USA) and Huiyang Zhou (North Carolina State University, USA)</i>	

Potpourri

Facile: Fast, Accurate, and Interpretable Basic-Block Throughput Prediction	87
<i>Andreas Abel (Saarland University, Germany), Shrey Sharma (Saarland University, Germany), and Jan Reineke (Saarland University, Germany)</i>	
Rebasing Microarchitectural Research with Industry Traces	100
<i>Josué Feliu (Universitat Politècnica de València, Spain), Arthur Perais (Univ. Grenoble Alpes, CNRS, France), Daniel A. Jiménez (Texas A&M University, USA), and Alberto Ros (University of Murcia, Spain)</i>	
Performance Implications of Async Memcpy and UVM: A Tale of Two Data Transfer Modes	115
<i>Ruihao Li (The University of Texas at Austin), Sanjana Yadav (The University of Texas at Austin), Qinzhe Wu (The University of Texas at Austin), Krishna Kavi (University of North Texas), Gayatri Mehta (University of North Texas), Neeraja J. Yadwadkar (The University of Texas at Austin; VMware Research), and Lizy K. John (The University of Texas at Austin)</i>	
Decepticon: Attacking Secrets of Transformers	128
<i>Mujahid Al Rafi (University of California, USA), Yuan Feng (University of California, USA), Fan Yao (University of Central Florida, USA), Meng Tang (University of California, USA), and Hyeran Jeon (University of California, USA)</i>	

Artificial Intelligence

Tale of Two Cs: Computation vs. Communication Scaling for Future Transformers on Future Hardware	140
<i>Suchita Pati (University of Wisconsin-Madison; Advanced Micro Devices, Inc., USA), Shaizeen Aga (Advanced Micro Devices, Inc., USA), Mahzabeen Islam (Advanced Micro Devices, Inc., USA), Nuwan Jayasena (Advanced Micro Devices, Inc., USA), and Matthew D. Sinclair (University of Wisconsin-Madison; Advanced Micro Devices, Inc., USA)</i>	

MMBench: Benchmarking End-to-End Multi-Modal DNNs and Understanding Their Hardware-Software Implications	154
<i>Cheng Xu (Shanghai Jiao Tong University, China), Xiaofeng Hou (Shanghai Jiao Tong University, China), Jiacheng Liu (Shanghai Jiao Tong University, China), Chao Li (Shanghai Jiao Tong University, China), Tianhao Huang (Shanghai Jiao Tong University, China), Xiaozhi Zhu (Shanghai Jiao Tong University, China), Mo Niu (Shanghai Jiao Tong University, China), Lingyu Sun (Shanghai Jiao Tong University, China), Peng Tang (Shanghai Jiao Tong University, China), Tongqiao Xu (Shanghai Jiao Tong University, China), Kwang-Ting Cheng (Hongkong University of Science and Technology, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
mNPUsim: Evaluating the Effect of Sharing Resources in Multi-core NPUs	167
<i>Soojin Hwang (KAIST, Republic of Korea), Sunho Lee (KAIST, Republic of Korea), Jungwoo Kim (KAIST, Republic of Korea), Hongbeen Kim (KAIST, Republic of Korea), and Jaehyuk Huh (KAIST, Republic of Korea)</i>	
Quantization for Bayesian Deep Learning: Low-Precision Characterization and Robustness	180
<i>Jun-Liang Lin (The Pennsylvania State University; Intel), Ranganath Krishnan (Intel), Keyur Rugeanathbhai Ranipa (Intel), Mahesh Subedar (Intel), Vrushabh Sanghavi (Intel), Meena Arunachalam (AMD), Omesh Tickoo (Intel), Ravishankar Iyer (Intel), and Mahmut Taylan Kandemir (The Pennsylvania State University)</i>	

Posters

Dataset Announcement: MS-BioGraphs, Trillion-Scale Public Real-World Sequence Similarity Graphs	193
<i>Mohsen Koochi Esfahani (Queen's University Belfast, United Kingdom; University of Sistan, Iran), Paolo Boldi (Università degli Studi di Milano, Italy), Hans Vandierendonck (Queen's University Belfast, United Kingdom), Peter Kilpatrick (Queen's University Belfast, United Kingdom), and Sebastiano Vigna (Università degli Studi di Milano, Italy)</i>	
ERAS: A Flexible and Scalable Framework for Seamless Integration of RTL Models with Structural Simulation Toolkit	196
<i>Shubham Nema (North Carolina State University, United States), Shiva Kaushik Chundururu (North Carolina State University, United States), Charan Kodigal (North Carolina State University, United States), Gwendolyn Voskuilen (Sandia National Laboratories, United States), Arun F. Rodrigues (Sandia National Laboratories, United States), Scott Hemmert (Sandia National Laboratories, United States), Ben Feinberg (Sandia National Laboratories, United States), Hyekeun Lee (North Carolina State University, United States), Amro Awad (North Carolina State University, United States), and Clayton Hughes (Sandia National Laboratories, United States)</i>	
CInC: Workload Characterization In Context of Resource Contention	201
<i>Cesar Gomes (Tufts University) and Mark Hempstead (Tufts University)</i>	

VelociTI: An Architecture-Level Performance Modeling Framework for Trapped Ion Quantum Computers	206
<i>Alexander Hankin (Harvard University; Intel Labs, USA), Abdulrahman Mahmoud (Harvard University, USA), Mark Hempstead (Tufts University, USA), David Brooks (Harvard University, USA), and Gu-Yeon Wei (Harvard University, USA)</i>	
Evaluating Homomorphic Operations on a Real-World Processing-In-Memory System	211
<i>Harshita Gupta (ETH Zürich), Mayank Kabra (ETH Zürich), Juan Gómez-Luna (ETH Zürich), Kanellopoulos Konstantinos (ETH Zürich), and Onur Mutlu (ETH Zürich)</i>	
Energy Efficiency of Out-of-Order CPUs: Comparative Study and Microarchitectural Hotspot Characterization of RISC-V Designs	216
<i>Odysseas Chatzopoulos (University of Athens, Greece), George Papadimitriou (University of Athens, Greece), Wing Shek Wong (Intel, USA), and Dimitris Gizopoulos (University of Athens, Greece)</i>	
Empirical Composite Workload Analysis for RSU-Assisted Computation Offloading in Connected Vehicle Services	221
<i>BaekGyu Kim (DGIST, South Korea) and Deepak Gangadharan (IIIT Hyderabad, India)</i>	
How Addresses are Made	223
<i>Xiaoyue Chen (Uppsala University), Pavlos Aimoniotis (Uppsala University), and Stefanos Kaxiras (Uppsala University)</i>	
Optimising GPGPU Execution Through Runtime Micro-Architecture Parameter Analysis	226
<i>Giuseppe M. Sarda (KU Leuven; imec, Belgium), Nimish Shah (KU Leuven, Belgium), Debjyoti Bhattacharjee (imec, Belgium), Peter Debacker (imec, Belgium), and Marian Verhelst (KU Leuven; imec, Belgium)</i>	
NVMSurvey: Recent Advances and Comparative Analysis of Emerging Non-Volatile Memories (eNVMs)	229
<i>Sandesh Ghimire (Amherst College, USA), Shinsaku Kataoka (Amherst College, USA), and Lillian Pentecost (Amherst College, USA)</i>	
CompressionGPT: Evaluating Fault Tolerance of a Compressed Large Language Model	232
<i>Neil Kapur (Amherst College, United States), America Rangel (Amherst College, United states), and Lillian Pentecost (Amherst College, United States)</i>	
Author Index	235