2016 IEEE International Symposium on Workload Characterization (IISWC 2016)

Providence, Rhode Island, USA
25-27 September 2016

IEEE Catalog Number: CFP16236-POD
ISBN: 978-1-5090-3897-8

Copyright © 2016 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 publication is a representation of what appears in the IEEE Digital Libraries. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number: CFP16236-POD
ISBN (Print-On-Demand): 978-1-5090-3897-8
ISBN (Online): 978-1-5090-3896-1

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

Table of Contents

2016 IEEE International Symposium on Workload Characterization
IISWC 2016

Table of Contents ..... *iii*
Message from the General Chair ..... *vi*
Message from the Program Co-Chairs ..... *vii*
IISWC 2016 Organization ..... *viii*
IISWC 2016 Sponsors and Supporters ..... *ix*

Keynote Address I

Running on Empty: Getting Work Done on Battery-Free Energy Harvesting Platforms 1
Alanson Sample (Disney Research)

Keynote Address II

The Convergence of Physical/Digital Worlds: Implications on Workloads & Architecture 2
Ravishankar Iyer (Intel Corp.)

Session 1: Best Paper Nominees

TailBench: A Benchmark Suite and Evaluation Methodology for Latency-Critical Applications 3
Harshad Kasture, Daniel Sanchez (Massachusetts Institute of Technology)

Hetero-Mark, A Benchmark Suite for CPU-GPU Collaborative Computing 13
Yifan Sun, Xiang Gong, Amir Kavyan Ziabari, Leiming Yu, Xiangyu Li, Saoni Mukherjee, Carter McCardwell (Northeastern University), Alejandro Villegas (University of Málaga), David Kaeli (Northeastern University)

Measuring and Modeling On-Chip Interconnect Power on Real Hardware 23
Vignesh Adhinarayanan (Virginia Tech), Indrani Paul, Joseph L. Greathouse, Wei Huang (AMD Research), Ashutosh Pattnaik (Penn State University), Wu-chun Feng (Virginia Tech)

Session 2: Workload Characterization

Characterization of Quantum Workloads on SIMD Architectures 34
Robert Risque, Adwait Jog (College of William and Mary)

Characterizing the Workload of a Netflix Streaming Video Server 43
Jim Summers, Tim Brecht (University of Waterloo), Derek Eager (University of Saskatchewan), Alex Gutarin (Netflix)

Characterization and Mitigation of Power Contention across Multiprogrammed Workloads 55
Hiroshi Sasaki (Columbia University), Alper Buyuktosunoglu, Augusto Vega, Pradip Bose (IBM T. J. Watson Research Center)

Session 3: Operating Systems and Virtual Machines

Container Management as Emerging Workload for Operating Systems 65
Tatsushi Inagaki, Yohei Ueda, Moriyoshi Ohara (IBM Research - Tokyo)

Overhead of Deoptimization Checks in the V8 JavaScript Engine 75
Gabriel Southern, Jose Renau (University of California, Santa Cruz)

Workload Characterization for Microservices 85
Takanori Ueda, Takuya Nakaike, Moriyoshi Ohara (IBM Research - Tokyo)

Session 4: Benchmark Formation and Suites

PBench: A Benchmark Suite for Characterizing 3D Printing Prefabrication 95
Fan Yang, Feng Lin, Chen Song, Chi Zhou (University at Buffalo, SUNY), Zhanpeng Jin (Binghamton University, SUNY), Wenyao Xu (University at Buffalo, SUNY)

ANMLZoo: A Benchmark Suite for Exploring Bottlenecks in Automata Processing Engines and Architectures 105
Jack Wadden, Vinh Dang, Nathan Brunelle, Tommy Tracy II, Deyuan Guo, Elaheh Sadredini, Ke Wang, Chunkun Bo, Gabriel Robins, Mircea Stan, Kevin Skadron (University of Virginia)

SPEC-AX and PARSEC-AX: Extracting Accelerator Benchmarks from Microprocessor Benchmarks 117
Snehasish Kumar, William N. Sumner, Arrvindh Shriraman (Simon Fraser University)

Session 5: Hardware-Software Codesign

Rebalancing the Core Front-End through HPC Code Analysis 128
Ugljesa Milic (Barcelona Supercomputing Center, Universitat Politècnica de Catalunya), Paul Carpenter (Barcelona Supercomputing Center), Alejandro Rico (ARM Inc.), Alex Ramirez (Nvidia Corp.)

Quantitative Characterization of the Software Layer of a HW/SW Co-Designed Processor 138
José Cano, Rakesh Kumar (University of Edinburgh), Aleksandar Brankovic (Intel), Demos Pavlou, Kyriakos Stavrou (11pets), Enric Gibert (Pharmacelera), Alejandro Martínez (ARM), Antonio González (Universitat Politècnica de Catalunya)

Fathom: Reference Workloads for Modern Deep Learning Methods 148
Robert Adolf, Saketh Rama, Brandon Reagen, Gu-Yeon Wei, David Brooks (Harvard University)

Session 6: GPGPUs and Heterogeneous Computing

ID-Cache: Instruction and Memory Divergence Based Cache Management for GPUs 158
Akhil Arunkumar, Shin-Ying Lee, Carole-Jean Wu (Arizona State University)

Evaluating the Effect of Last-Level Cache Sharing on Integrated GPU-CPU Systems with Heterogeneous Applications 168
Victor García (Universitat Politècnica de Catalunya / Barcelona Supercomputing Center), Juan Gómez-Luna (Universidad de Córdoba), Thomas Grass (Universitat Politècnica de Catalunya / Barcelona Supercomputing Center), Alejandro Rico (ARM Inc.), Eduard Ayguade (Universitat Politècnica de Catalunya / Barcelona Supercomputing Center), Antonio J. Peña (Barcelona Supercomputing Center)

GPU Concurrency Choices in Graph Analytics 178
Masab Ahmad, Omer Khan (University of Connecticut)

Session 7: Memory and Storage

Memory Controller Design Under Cloud Workloads 188
Mostafa Mahmoud, Andreas Moshovos (University of Toronto)

A Simulation Analysis of Reliability in Primary Storage Deduplication 199
Min Fu (Huazhong University of Science and Technology), Patrick P. C. Lee (The Chinese University of Hong Kong), Dan Feng (Huazhong University of Science and Technology), Zuoning Chen (National Engineering Research Center for Parallel Computer), Yu Xiao (Huazhong University of Science and Technology)

Quantifying the Performance Impact of Large Pages on In-Memory Big-Data Workloads 209
Jinsu Park, Myeonggyun Han, Woongki Baek (UNIST)

Poster Session

Analyzing Power Consumption and Characterizing User Activities on Smartwatches: Summary 219
Emirhan Poyraz, Gokhan Memik (Northwestern University)

Resilience Characterization of a Vision Analytics Application Under Varying Degrees of Approximation ... 221
Radha Venkatagiri (University of Illinois at Urbana-Champaign), Karthik Swaminathan, Chung-Ching Lin (IBM Research), Liang Wang (University of Virginia), Alper Buyuktosunoglu, Pradip Bose (IBM Research), Sarita Adve (University of Illinois at Urbana-Champaign)

Identifying Representative Regions of Parallel HPC Applications: a Cross-architectural Evaluation 223
Alexandra Ferrerón (Universidad de Zaragoza), Radhika Jagtap, Roxana Ruşitoru (ARM Ltd., U.K.)

Power-Aware Characterization and Mapping of Workloads on CPU-GPU Processors 225
Kapil Dev, Xin Zhan, Sherief Reda (Brown University)

Treelogy: A Benchmark Suite for Tree Traversal Applications 227
Nikhil Hegde, Jianqiao Liu, Milind Kulkarni (Purdue University)

Characterizing Memory Bottlenecks in GPGPU Workloads 229
Saumay Dublish, Vijay Nagarajan, Nigel Topham (University of Edinburgh)

Author Index 231