

2018 IEEE International Symposium on Workload Characterization (IISWC 2018)

**Raleigh, North Carolina
30 September – 2 October 2018**



**IEEE Catalog Number: CFP18236-POD
ISBN: 978-1-5386-6781-1**

**Copyright © 2018 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:	CFP18236-POD
ISBN (Print-On-Demand):	978-1-5386-6781-1
ISBN (Online):	978-1-5386-6780-4

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

TABLE OF CONTENTS

μ SUITE: A BENCHMARK SUITE FOR MICROSERVICES	1
<i>Akshitha Sriraman ; Thomas F. Wenisch</i>	
AUTOMATAZOO: A MODERN AUTOMATA PROCESSING BENCHMARK SUITE	13
<i>Jack Wadden ; Tommy Tracy ; Elaheh Sadredini ; Lingxi Wu ; Chunkun Bo ; Jesse Du ; Yizhou Wei ; Jeffrey Udall ; Matthew Wallace ; Mircea Stan ; Kevin Skadron</i>	
VCOMPUTE BENCH: A VULKAN BENCHMARK SUITE FOR GPGPU ON MOBILE AND EMBEDDED GPUS	25
<i>Nadjib Mammeri ; Ben Juurlink</i>	
QUANTITATIVE OVERHEAD ANALYSIS FOR PYTHON	36
<i>Mohamed Ismail ; G. Edward Suh</i>	
DATA MOTIF-BASED PROXY BENCHMARKS FOR BIG DATA AND AI WORKLOADS	48
<i>Wanling Gao ; Jianfeng Zhan ; Lei Wang ; Chunjie Luo ; Zhen Jia ; Daoyi Zheng ; Chen Zheng ; Xiwen He ; Hainan Ye ; Haibin Wang ; Rui Ren</i>	
A NUMA-AWARE PROVABLY-EFFICIENT TASK-PARALLEL PLATFORM BASED ON THE WORK-FIRST PRINCIPLE	59
<i>Justin Deters ; Jiaye Wu ; Yifan Xu ; I-Ting Angelina Lee</i>	
HOT REGIONS IN SPEC CPU2017	71
<i>Qinzhe Wu ; Steven Flolid ; Shuang Song ; Junyong Deng ; Lizy K. John</i>	
PARCHMINT: A MICROFLUIDICS BENCHMARK SUITE	78
<i>Brian Crites ; Radhakrishna Sanka ; Joshua Lippai ; Jeffrey Mcdaniel ; Philip Brisk ; Douglas Densmore</i>	
CHOPSTIX: SYSTEMATIC EXTRACTION OF CODE-REPRESENTATIVE MICROBENCHMARKS	80
<i>Calvin Bulla ; Lluç Alvarez ; Miquel Moreto ; Ramon Bertran ; Alper Buyuktosunoglu ; Pradip Bose</i>	
CHARACTERIZING DNN MODELS FOR EDGE-CLOUD COMPUTING	82
<i>Chunwei Xia ; Jiacheng Zhao ; Huimin Cui ; Xiaobing Feng</i>	
GENE SEQUENCING: WHERE TIME GOES	84
<i>Meysam Roodi ; Andreas Moshovos</i>	
CHARACTERIZING SOURCES OF INEFFECTUAL COMPUTATIONS IN DEEP LEARNING NETWORKS	86
<i>Milod Nikolic ; Mostafa Mahmoud ; Andreas Moshovos</i>	
BENCHMARKING AND ANALYZING DEEP NEURAL NETWORK TRAINING	88
<i>Hongyu Zhu ; Mohamed Akrou ; Bojian Zheng ; Andrew Pelegris ; Anand Jayarajan ; Amar Phanishayee ; Bianca Schroeder ; Gennady Pekhimenko</i>	
CHARACTERISING ACROSS-STACK OPTIMISATIONS FOR DEEP CONVOLUTIONAL NEURAL NETWORKS	101
<i>Jack Turner ; José Cano ; Valentin Radu ; Elliot J. Crowley ; Michael O'boyle ; Amos Storkey</i>	
MEMORY REQUIREMENTS FOR CONVOLUTIONAL NEURAL NETWORK HARDWARE ACCELERATORS	111
<i>Kevin Siu ; Dylan Malone Stuart ; Mostafa Mahmoud ; Andreas Moshovos</i>	
PROFILING DNN WORKLOADS ON A VOLTA-BASED DGX-1 SYSTEM	122
<i>Saiful A. Mojumder ; Marcia S Louis ; Yifan Sun ; Amir Kavyan Ziabari ; José L. Abellán ; John Kim ; David Kaeli ; Ajay Joshi</i>	
HETEROGENEOUS MEMORY SUBSYSTEM FOR NATURAL GRAPH ANALYTICS	134
<i>Abraham Addisie ; Hiwot Kassa ; Opeoluwa Matthews ; Valeria Bertacco</i>	
MINIMIZING READ SEEKS FOR SMR DISK	146
<i>Mohammad Hossein Hajkazemi ; Mania Abdi ; Peter Desnoyers</i>	
REVIVING ZOMBIE PAGES ON SSDS	156
<i>Nima Elyasi ; Anand Sivasubramaniam ; Mahmut T. Kandemir ; Chita R. Das</i>	
SIMULATING PCI-EXPRESS INTERCONNECT FOR FUTURE SYSTEM EXPLORATION	168
<i>Mohammad Alian ; Krishna Parasuram Srinivasan ; Nam Sung Kim</i>	
INTERFERENCE FROM GPU SYSTEM SERVICE REQUESTS	179
<i>Arkaprava Basu ; Joseph L. Greathouse ; Guru Venkataramani ; Ján Veselý</i>	
TARTAN: EVALUATING MODERN GPU INTERCONNECT VIA A MULTI-GPU BENCHMARK SUITE	191
<i>Ang Li ; Shuaiwen Leon Song ; Jieyang Chen ; Xu Liu ; Nathan Tallent ; Kevin Barker</i>	

WHEN IS GRAPH REORDERING AN OPTIMIZATION? STUDYING THE EFFECT OF LIGHTWEIGHT GRAPH REORDERING ACROSS APPLICATIONS AND INPUT GRAPHS	203
---	------------

Vignesh Balaji ; Brandon Lucia

Author Index