

2019 IEEE International Symposium on Workload Characterization (IISWC 2019)

**Orlando, Florida, USA
3 – 5 November 2019**



**IEEE Catalog Number: CFP19236-POD
ISBN: 978-1-7281-4046-9**

**Copyright © 2019 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:	CFP19236-POD
ISBN (Print-On-Demand):	978-1-7281-4046-9
ISBN (Online):	978-1-7281-4045-2

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

A CLOSER LOOK AT LIGHTWEIGHT GRAPH REORDERING	1
<i>Priyank Faldu ; Jeff Diamond ; Boris Grot</i>	
DETECTING LAST-LEVEL CACHE CONTENTION IN WORKLOAD COLOCATION WITH META LEARNING.....	14
<i>Huanxing Shen ; Cong Li</i>	
PERFORMANCE AWARE CONVOLUTIONAL NEURAL NETWORK CHANNEL PRUNING FOR EMBEDDED GPUS	24
<i>Valentin Radu ; Kuba Kaszyk ; Yuan Wen ; Jack Turner ; José Cano ; Elliot J. Crowley ; Björn Franke ; Amos Storkey ; Michael O'Boyle</i>	
CHARACTERIZING THE DEPLOYMENT OF DEEP NEURAL NETWORKS ON COMMERCIAL EDGE DEVICES.....	35
<i>Ramyad Hadidi ; Jiashen Cao ; Yilun Xie ; Bahar Asgari ; Tushar Krishna ; Hyesoon Kim</i>	
TRIMMING THE TAIL FOR DETERMINISTIC READ PERFORMANCE IN SSDS	49
<i>Nima Elyasi ; Changho Choi ; Anand Sivasubramaniam ; Jingpei Yang ; Vijay Balakrishnan</i>	
OPTIMIZING HYPERPLANE SWEEP OPERATIONS USING ASYNCHRONOUS MULTI-GRAIN GPU TASKS.....	59
<i>Anirudh Mohan Kaushik ; Ashwin M. Aji ; Muhammad Amber Hassaan ; Noel Chalmers ; Noah Wolfe ; Scott Moe ; Sooraj Puthoor ; Bradford M. Beckmann</i>	
EFFICACY OF STATISTICAL SAMPLING ON CONTEMPORARY WORKLOADS: THE CASE OF SPEC CPU2017	70
<i>Sarabjeet Singh ; Manu Awasthi</i>	
AUTONOMOUS DATA-RACE-FREE GPU TESTING	81
<i>Tuan Ta ; Xianwei Zhang ; Anthony Gutierrez ; Bradford M. Beckmann</i>	
SNU-NPB 2019: PARALLELIZING AND OPTIMIZING NPB IN OPENCL AND CUDA FOR MODERN GPUS	93
<i>Youngdong Do ; Hyungmo Kim ; Pyeongseok Oh ; Daeyoung Park ; Jaejin Lee</i>	
WORKLOAD-AWARE DRAM ERROR PREDICTION USING MACHINE LEARNING	106
<i>Lev Mukhanov ; Konstantinos Tovletoglou ; Hans Vandierendonck ; Dimitrios S. Nikolopoulos ; Georgios Karakonstantis</i>	
MULTI-BIT UPSETS VULNERABILITY ANALYSIS OF MODERN MICROPROCESSORS.....	119
<i>Athanasios Chatzidimitriou ; George Papadimitriou ; Christos Gavanis ; George Katsoridas ; Dimitris Gizopoulos</i>	
DEEP LEARNING LANGUAGE MODELING WORKLOADS: WHERE TIME GOES ON GRAPHICS PROCESSORS.....	131
<i>Ali Hadi Zadeh ; Zissis Poulos ; Andreas Moshovos</i>	
EVALUATION OF NON-VOLATILE MEMORY BASED LAST LEVEL CACHE GIVEN MODERN USE CASE BEHAVIOR	143
<i>Alexander Hankin ; Tomer Shapira ; Karthik Sangaiah ; Michael Lui ; Mark Hempstead</i>	
ONE SIZE DOESN'T FIT ALL: QUANTIFYING PERFORMANCE PORTABILITY OF GRAPH APPLICATIONS ON GPUS.....	155
<i>Tyler Sorensen ; Sreepathi Pai ; Alastair F. Donaldson</i>	
BHIVE: A BENCHMARK SUITE AND MEASUREMENT FRAMEWORK FOR VALIDATING X86-64 BASIC BLOCK PERFORMANCE MODELS	167
<i>Yishen Chen ; Ajay Brahmakshatriya ; Charith Mendis ; Alex Renda ; Eric Atkinson ; Ondřej Sýkora ; Saman Amarasinghe ; Michael Carbin</i>	
SIMDHT-BENCH: CHARACTERIZING SIMD-AWARE HASH TABLE DESIGNS ON EMERGING CPU ARCHITECTURES	178
<i>Dipti Shankar ; Xiaoyi Lu ; Dhabaleswar K. DK Panda</i>	
CHARACTERIZING DEEP LEARNING TRAINING WORKLOADS ON ALIBABA-PAI	189
<i>Mengdi Wang ; Chen Meng ; Guoping Long ; Chuan Wu ; Jun Yang ; Wei Lin ; Yangqing Jia</i>	
AN OVERFLOW-FREE QUANTIZED MEMORY HIERARCHY IN GENERAL-PURPOSE PROCESSORS.....	203
<i>Marzieh Lenjani ; Patricia Gonzalez ; Elaheh Sadredini ; M Arif Rahman ; Mircea R. Stan</i>	
FASTER THAN FLASH: AN IN-DEPTH STUDY OF SYSTEM CHALLENGES FOR EMERGING ULTRA-LOW LATENCY SSDS	216
<i>Sungjoon Koh ; Junhyeok Jang ; Changrim Lee ; Miryeong Kwon ; Jie Zhang ; Myoungsoo Jung</i>	

BRANCH PREDICTION IS NOT A SOLVED PROBLEM: MEASUREMENTS, OPPORTUNITIES, AND FUTURE DIRECTIONS	228
<i>Chit-Kwan Lin ; Stephen J. Tarsa</i>	
HOLDCSIM: A HOLISTIC SIMULATOR FOR DATA CENTERS	239
<i>Fan Yao ; Kathy Ngyugen ; Sai Santosh Dayapule ; Jingxin Wu ; Bingqian Lu ; Suresh Subramaniam ; Guru Venkataramani</i>	
OPTIMIZING GPU CACHE POLICIES FOR MI WORKLOADS	243
<i>Johnathan Alsop ; Matthew D. Sinclair ; Srikanth Bharadwaj ; Alexandru Dutu ; Anthony Gutierrez ; Onur Kayiran ; Michael LeBeane ; Brandon Potter ; Sooraj Puthoor ; Xianwei Zhang ; Tsung Tai Yeh ; Bradford M. Beckmann</i>	
PERSISTENT MEMORY WORKLOAD CHARACTERIZATION: A HARDWARE PERSPECTIVE	249
<i>Xiao Liu ; Bhaskar Jupudi ; Pankaj Mehra ; Jishen Zhao</i>	
FINGERPRINTING ANOMALOUS COMPUTATION WITH RNN FOR GPU-ACCELERATED HPC MACHINES	253
<i>Pengfei Zou ; Ang Li ; Kevin Barker ; Rong Ge</i>	
PERFORMANCE-DRIVEN PROGRAMMING OF MULTI-TFLOP DEEP LEARNING ACCELERATORS	257
<i>Swagath Venkataramani ; Jungwook Choi ; Vijayalakshmi Srinivasan ; Kailash Gopalakrishnan ; Leland Chang</i>	
BARRIER SYNCHRONIZATION VS. VOLTAGE NOISE: A QUANTITATIVE ANALYSIS	263
<i>Zamshed I. Chowdhury ; S. Karen Khatamifard ; Zhaoyong Zheng ; Tali Moreshet ; R. Iris Bahar ; Ulya R. Karpuzcu</i>	
CHARACTERIZING THE PERFORMANCE/ACCURACY TRADEOFF OF HIGH-PRECISION APPLICATIONS VIA AUTO-TUNING	268
<i>Ruidong Gu ; Paul Beata ; Michela Becchi</i>	
Author Index	