

2016 45th International Conference on Parallel Processing (ICPP 2016)

**Philadelphia, Pennsylvania, USA
16-19 August 2016**



**IEEE Catalog Number: CFP16127-POD
ISBN: 978-1-5090-2824-5**

**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:	CFP16127-POD
ISBN (Print-On-Demand):	978-1-5090-2824-5
ISBN (Online):	978-1-5090-2823-8
ISSN:	0190-3918

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

2016 45th International Conference on Parallel Processing

ICPP 2016

Table of Contents

Message from the General Chair	xii
Message from the Program Co-Chairs	xiii
Organizing and Program Committees.....	xiv

Session 1A: Data Center and Cloud 1

Efficient Virtual Network Embedding for Variable Size Virtual Machines in Fat-Tree Data Centers	1
<i>Jun Duan and Yuanyuan Yang</i>	
MIC: An Efficient Anonymous Communication System in Data Center Networks	11
<i>Tingwei Zhu, Dan Feng, Yu Hua, Fang Wang, Qingyu Shi, and Jiahao Liu</i>	
AppBag: Application-Aware Bandwidth Allocation for Virtual Machines in Cloud Environment	21
<i>Dian Shen, Junzhou Luo, Fang Dong, and Junxue Zhang</i>	
Performance Boosting Opportunities under Communication Imbalance in Power-Constrained HPC Clusters	31
<i>Leonardo Piga, Indrani Paul, and Wei Huang</i>	
RRect: A Novel Server-centric Data Center Network with High Availability	41
<i>Zhenhua Li and Yuanyuan Yang</i>	

Session 1B: Architecture 1

Making In-Memory Frequent Pattern Mining Durable and Energy Efficient	47
<i>Yi Lin, Po-Chun Huang, Duo Liu, Xiao Zhu, and Liang Liang</i>	
Run-Time Performance Estimation and Fairness-Oriented Scheduling Policy for Concurrent GPGPU Applications	57
<i>Qingda Hu, Jiwu Shu, Jie Fan, and Youyou Lu</i>	

Performance Analysis of GPU-Based Convolutional Neural Networks	67
<i>Xiaqing Li, Guangyan Zhang, H. Howie Huang, Zhufan Wang, and Weimin Zheng</i>	
Proxy-Guided Load Balancing of Graph Processing Workloads on Heterogeneous Clusters	77
<i>Shuang Song, Meng Li, Xinnian Zheng, Michael LeBeane, Jee Ho Ryoo, Reena Panda, Andreas Gerstlauer, and Lizy K. John</i>	
Piccolo: A Fast and Efficient Rollback System for Virtual Machine Clusters	87
<i>Lei Cui, Zhiyu Hao, Chonghua Wang, Haiqiang Fei, and Zhenquan Ding</i>	

Session 2A: Parallel Algorithms

Parallel k-Means++ for Multiple Shared-Memory Architectures	93
<i>Patrick Mackey and Robert R. Lewis</i>	
High Performance Parallel Algorithms for the Tucker Decomposition of Sparse Tensors	103
<i>Oguz Kaya and Bora Ucar</i>	
Parallel Tree Traversal for Nearest Neighbor Query on the GPU	113
<i>Moohyeon Nam, Jinwoong Kim, and Beomseok Nam</i>	
Resilient Application Co-scheduling with Processor Redistribution	123
<i>Anne Benoit, Loïc Pottier, and Yves Robert</i>	
Efficient Parallel Algorithms for k-Center Clustering	133
<i>Jessica McClintock and Anthony Wirth</i>	

Session 2B: Architecture 2

Understanding the Architectural Characteristics of EDA Algorithms	139
<i>Xin Wang, Xiaofeng Ji, Yunping Lu, Yi Li, Weijia Zhou, Weihua Zhang, and Wenyun Zhao</i>	
Exploring Variation-Aware Fault-Tolerant Cache under Near-Threshold Computing	149
<i>Jing Wang, Yanjun Liu, Weigong Zhang, Kezhong Lu, Kenei Qiu, Xin Fu, and Tao Li</i>	
Tetris Write: Exploring More Write Parallelism Considering PCM Asymmetries	159
<i>Zheng Li, Fang Wang, Dan Feng, Yu Hua, Wei Tong, Jingning Liu, and Xiang Liu</i>	
ROP: Alleviating Refresh Overheads via Reviving the Memory System in Frozen Cycles	169
<i>Ping Huang, Wenjie Liu, Kun Tang, Xubin He, and Ke Zhou</i>	
Thread Similarity Matrix: Visualizing Branch Divergence in GPGPU Programs	179
<i>Zhibin Yu, Lieven Eeckhout, and Chengzhong Xu</i>	

Session 3A: Programming Techniques 1

One-Sided Interface for Matrix Operations Using MPI-3 RMA: A Case Study with Elemental	185
<i>Sayan Ghosh, Jeff R. Hammond, Antonio J. Peña, Pavan Balaji, Assefaw H. Gebremedhin, and Barbara Chapman</i>	
SWAP-Assembler 2: Optimization of De Novo Genome Assembler at Extreme Scale	195
<i>Jintao Meng, Sangmin Seo, Pavan Balaji, Yanjie Wei, Bingqiang Wang, and Shenzhong Feng</i>	
Programming Techniques for the Automata Processor	205
<i>Indranil Roy, Ankit Srivastava, and Srinivas Aluru</i>	
RCHC: A Holistic Runtime System for Concurrent Heterogeneous Computing	211
<i>Jinsu Park and Woongki Baek</i>	

Session 3B: Parallel Algorithms 2

An Unbounded Nonblocking Double-Ended Queue	217
<i>Matthew Graichen, Joseph Izraelevitz, and Michael L. Scott</i>	
Criticality-Aware Partitioning for Multicore Mixed-Criticality Systems	227
<i>Jian-Jun Han, Xin Tao, Dakai Zhu, and Hakan Aydin</i>	
A Parallel Hill-Climbing Refinement Algorithm for Graph Partitioning	236
<i>Dominique LaSalle and George Karypis</i>	
Massively-Parallel Lossless Data Decompression	242
<i>Evangelia Sitaridi, Rene Mueller, Tim Kaldewey, Guy Lohman, and Kenneth A. Ross</i>	

Session 4A: Data Cloud and Cloud 2

Improving Data Transfer Throughput with Direct Search Optimization	248
<i>Prasanna Balaprakash, Vitali Morozov, Rajkumar Kettimuthu, Kalyan Kumaran, and Ian Foster</i>	
MPI Overlap: Benchmark and Analysis	258
<i>Alexandre Denis and François Trahay</i>	
High Performance MPI Library for Container-Based HPC Cloud on InfiniBand Clusters	268
<i>Jie Zhang, Xiaoyi Lu, and Dhabaleswar K. (DK) Panda</i>	
AccuracyTrader: Accuracy-Aware Approximate Processing for Low Tail Latency and High Result Accuracy in Cloud Online Services	278
<i>Rui Han, Siguang Huang, Fei Tang, Fugui Chang, and Jianfeng Zhan</i>	

CoARC: Co-operative, Aggressive Recovery and Caching for Failures in Erasure Coded Hadoop	288
<i>Pradeep Subedi, Ping Huang, Tong Liu, Joseph Moore, Stan Skelton, and Xubin He</i>	

Session 4B: Cyberphysical Systems 1

Optimal Multi-taxi Dispatch for Mobile Taxi-Hailing Systems	294
<i>Guoju Gao, Mingjun Xiao, and Zhenhua Zhao</i>	
Fast RFID Polling Protocols	304
<i>Jia Liu, Bin Xiao, Xuan Liu, and Lijun Chen</i>	
Exploiting Real-Time Traffic Light Scheduling with Taxi Traces	314
<i>Zongjian He, Daqiang Zhang, Jiannong Cao, Xuefeng Liu, Xiaopeng Fan, and Chengzhong Xu</i>	
An Efficient Wireless Power Transfer System to Balance the State of Charge of Electric Vehicles	324
<i>Ankur Sarker, Chenxi Qiu, Haiying Shen, Andrea Gil, Joachim Taiber, Mashrur Chowdhury, Jim Martin, Mac Devine, and AJ Rindos</i>	
EchoLoc: Accurate Device-Free Hand Localization Using COTS Devices	334
<i>Huijie Chen, Fan Li, and Yu Wang</i>	

Session 5A: Parallel Algorithms 3

Randomly Optimized Grid Graph for Low-Latency Interconnection Networks	340
<i>Koji Nakano, Daisuke Takafuji, Satoshi Fujita, Hiroki Matsutani, Ikki Fujiwara, and Michihiro Koibuchi</i>	
Optimal Collision/Conflict-Free Distance-2 Coloring in Wireless Synchronous Broadcast/Receive Tree Networks	350
<i>Davide Frey, Hicham Lakhlef, and Michel Raynal</i>	
Help-Optimal and Language-Portable Lock-Free Concurrent Data Structures	360
<i>Bapi Chatterjee, Ivan Walulya, and Philippas Tsigas</i>	
DC-Top-k: A Novel Top-k Selecting Algorithm and Its Parallelization	370
<i>Zhengyuan Xue, Ruixuan Li, Heng Zhang, Xiwu Gu, and Zhiyong Xu</i>	
Efficient 2-Body Statistics Computation on GPUs: Parallelization & Beyond	380
<i>Napath Pitaksirianan, Zhila Nouri, and Yi-Cheng Tu</i>	

Session 5B: Storage Systems

Think Global, Act Local: A Buffer Cache Design for Global Ordering and Parallel Processing in the WAFL File System	386
<i>Peter R. Denz, Matthew Curtis-Maury, and Vinay Devadas</i>	
Improving RAID Performance Using an Endurable SSD Cache	396
<i>Chu Li, Dan Feng, Yu Hua, and Fang Wang</i>	
In Situ Storage Layout Optimization for AMR Spatio-temporal Read Accesses	406
<i>Houjun Tang, Suren Byna, Steve Harenberg, Wenzhao Zhang, Xiaocheng Zou, Daniel F. Martin, Bin Dong, Dharshi Devendran, Kesheng Wu, David Trebotich, Scott Klasky, and Nagiza F. Samatova</i>	
Managing I/O Interference in a Shared Burst Buffer System	416
<i>Sagar Thapaliya, Purushotham Bangalore, Jat Lofstead, Kathryn Mohror, and Adam Moody</i>	
Guaranteed Bang for the Buck: Modeling VDI Applications with Guaranteed Quality of Service	426
<i>Hao Wen, David H.C. Du, Milan Shetti, Doug Voigt, and Shanshan Li</i>	

Session 6A: Programming Techniques 2

Scalable Hierarchical Polyhedral Compilation	432
<i>Benoît Pradelle, Benoît Meister, Muthu Baskaran, Athanasios Konstantinidis, Thomas Henretty, and Richard Lethin</i>	
The Future(s) of Transactional Memory	442
<i>Jingna Zeng, João Barreto, Seif Haridi, Luís Rodrigues, and Paolo Romano</i>	
Declarative Tuning for Locality in Parallel Programs	452
<i>Sanjay Chatterjee, Nick Vrvilo, Zoran Budimlić, Kathleen Knobe, and Vivek Sarkar</i>	
Ensemble Toolkit: Scalable and Flexible Execution of Ensembles of Tasks	458
<i>Vivekanandan Balasubramanian, Antons Treikalis, Ole Weidner, and Shantenu Jha</i>	

Session 6B: Cyberphysical Systems 2

TECH: A Thermal-Aware and Cost Efficient Mechanism for Colocation Demand Response	464
<i>Ziqi Zhao, Fan Wu, Shaolei Ren, Xiaofeng Gao, Guihai Chen, and Yong Cui</i>	
A Scalability Comparison Study of Data Management Approaches for Smart Metering Systems	474
<i>Houssem Chihoub and Christine Collet</i>	
A Comparison of Accelerator Architectures for Radio-Astronomical Signal-Processing Algorithms	484
<i>John W. Romein</i>	

MobiSensing: Exploiting Human Mobility for Multi-application Mobile Data Sensing with Low User Intervention	490
<i>Kang Chen and Haiying Shen</i>	

Session 7A: Performance Modeling

Sparse Matrix Format Selection with Multiclass SVM for SpMV on GPU	496
<i>Akrem Benatia, Weixing Ji, Yizhuo Wang, and Feng Shi</i>	
On the Impact of Widening Vector Registers on Sequence Alignment	506
<i>Jeffrey Daily, Ananth Kalyanaraman, Sriram Krishnamoorthy, and Bin Ren</i>	
The Case for Cross-Component Power Coordination on Power Bounded Systems	516
<i>Rong Ge, Xizhou Feng, Yangyang He, and Pengfei Zou</i>	
Performance Maximization via Frequency Oscillation on Temperature Constrained Multi-core Processors	526
<i>Shi Sha, Wujie Wen, Ming Fan, Shaolei Ren, and Gang Quan</i>	
RMD: A Resemblance and Mergence Based Approach for High Performance Deduplication	536
<i>Panfeng Zhang, Ping Huang, Xubin He, Hua Wang, Lingyu Yan, and Ke Zhou</i>	

Session 7B: GPU Applications

GFLink: An In-Memory Computing Architecture on Heterogeneous CPU-GPU Clusters for Big Data	542
<i>Cen Chen, Kenli Li, Aijia Ouyang, Zhuo Tang, and Keqin Li</i>	
Partial Flattening: A Compilation Technique for Irregular Nested Parallelism on GPGPUs	552
<i>Ming-Hsiang Huang and Wuu Yang</i>	
RegTT: Accelerating Tree Traversals on GPUs by Exploiting Regularities	562
<i>Feng Zhang, Peng Di, Hao Zhou, Xiangke Liao, and Jingling Xue</i>	
Optimizing GPU Register Usage: Extensions to OpenACC and Compiler Optimizations	572
<i>Xiaonan Tian, Dounia Khaldi, Deepak Eachempati, Rengan Xu, and Barbara Chapman</i>	
HppCnn: A High-Performance, Portable Deep-Learning Library for GPGPUs	582
<i>Yi Yang, Min Feng, and Srimat Chakradhar</i>	

Session 8A: Applications

Locality-Aware Laplacian Mesh Smoothing	588
<i>Guillaume Aupy, JeongHyung Park, and Padma Raghavan</i>	
Fault Tolerant Support Vector Machines	598
<i>Sameh Shohdy, Abhinav Vishnu, and Gagan Agrawal</i>	
Parallel Two-Dimensional Unstructured Anisotropic Delaunay Mesh	
Generation of Complex Domains for Aerospace Applications	608
<i>Juliette Pardue and Andrey Chernikov</i>	

Session 8B: Scalable Software

PARVMEC: An Efficient, Scalable Implementation of the Variational Moments	
Equilibrium Code	618
<i>Sudip K. Seal, Steven P. Hirshman, Andreas Wingen, Robert S. Wilcox, Mark R. Cianciosa, and Ezekial A. Unterberg</i>	
RepEx: A Flexible Framework for Scalable Replica Exchange Molecular	
Dynamics Simulations	628
<i>Antons Treikalis, Andre Merzky, Haoyuan Chen, Tai-Sung Lee, Darrin M. York, and Shantenu Jha</i>	
PCAF: Scalable, High Precision k-NN Search Using Principal Component	
Analysis Based Filtering	638
<i>Huan Feng, David Eyers, Steven Mills, Yongwei Wu, and Zhiyi Huang</i>	
Author Index	648