

2017 IEEE 24th International Conference on High Performance Computing (HiPC 2017)

**Jaipur, India
18 – 21 December 2017**



**IEEE Catalog Number: CFP17176-POD
ISBN: 978-1-5386-2294-0**

**Copyright © 2017 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:	CFP17176-POD
ISBN (Print-On-Demand):	978-1-5386-2294-0
ISBN (Online):	978-1-5386-2293-3
ISSN:	1094-7256

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

2017 IEEE 24th International Conference on High Performance Computing HiPC 2017

Table of Contents

Message from the General and Vice-General Chairs	xi
Message from the Program Chair	xiii
Message from the Steering Chair	xv
HiPC 2017 Committees	xvi
HiPC 2017 Technical Program	xxi

Keynote 1

End of Moore’s Law: Or, a Computer Architect’s Mid-life Crisis?	1
<i>Parthasarathy Ranganathan (Google)</i>	

Technical Session 1: Graph Algorithms

Exact and Parallel Triangle Counting in Dynamic Graphs	2
<i>Devavret Makkar (Georgia Institute of Technology), David A. Bader (Georgia Institute of Technology), and Oded Green (Georgia Institute of Technology)</i>	
Shared-Memory Graph Truss Decomposition	13
<i>Humayun Kabir (The Pennsylvania State University) and Kamesh Madduri (The Pennsylvania State University)</i>	
Approximate Computing Techniques for Iterative Graph Algorithms	23
<i>Ajay Panyala (Pacific Northwest National Laboratory), Omer Subasi (Pacific Northwest National Laboratory), Mahantesh Halappanavar (Pacific Northwest National Laboratory), Ananth Kalyanaraman (Washington State University), Daniel Chavarría-Miranda (Trovares Inc.), and Sriram Krishnamoorthy (Pacific Northwest National Laboratory)</i>	
Scalable Exact Parent Sets Identification in Bayesian Networks Learning with Apache Spark	33
<i>Subhadeep Karan (University at Buffalo) and Jaroslaw Zola (University at Buffalo)</i>	
Parallel Exact Dynamic Bayesian Network Structure Learning with Application to Gene Networks	42
<i>Md. Vasimuddin (Indian Institute of Technology Bombay) and Srinivas Aluru (Georgia Institute of Technology)</i>	

Parallel Asynchronous Distributed-Memory Maximal Independent Set Algorithm with Work Ordering	52
<i>Thejaka Kanewala (Indiana University), Marcin Zalewski (Pacific Northwest National Laboratory & University of Washington), and Andrew Lumsdaine (Pacific Northwest National Laboratory & University of Washington)</i>	

Technical Session 2: Architecture and Communication

Designing Registration Caching Free High-Performance MPI Library with Implicit On-Demand Paging (ODP) of InfiniBand	62
<i>Mingzhe Li (The Ohio State University), Xiaoyi Lu (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhabaleswar K. (DK) Panda (The Ohio State University)</i>	
Last Level Collective Hardware Prefetching For Data-Parallel Applications	72
<i>George Michelogiannakis (Lawrence Berkeley National Laboratory) and John Shalf (Lawrence Berkeley National Laboratory)</i>	
Kernel-Assisted Communication Engine for MPI on Emerging Manycore Processors	84
<i>Jahanzeb Maqbool Hashmi (The Ohio State University), Khaled Hamidouche (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhabaleswar K. (DK) Panda (The Ohio State University)</i>	
Support for Power Efficient Proactive Cooling Mechanisms	94
<i>Bilge Acun (IBM T. J. Watson Research Center), Eun Kyung Lee (IBM T. J. Watson Research Center), Yoonho Park (IBM T. J. Watson Research Center), and Laxmikant V. Kale (University of Illinois at Urbana-Champaign)</i>	
Redundant Arithmetic Based High Speed Carry Free Hybrid Adders with Built-In Scan Chain on FPGAs	104
<i>Ayan Palchoudhuri (Indian Institute of Technology Kharagpur) and Anindya Sundar Dhar (Indian Institute of Technology Kharagpur)</i>	
ConvLight: A Convolutional Accelerator with Memristor Integrated Photonic Computing	114
<i>Dharanidhar Dang (Texas A&M University), Jyotikrishna Dass (Texas A&M University), and Rabi Mahapatra (Texas A&M University)</i>	

Technical Session 3: Algorithms

Provably Efficient Scheduling of Dynamically Allocating Programs on Parallel Cache Hierarchies	124
<i>Guy E. Blelloch (Carnegie Mellon University), Phillip B. Gibbons (Carnegie Mellon University), and Harsha Vardhan Simhadri (Microsoft Research India)</i>	
Further Explorations in State-Space Search for Optimal Task Scheduling	134
<i>Michael Orr (University of Auckland) and Oliver Sinnen (University of Auckland)</i>	
A Novel Approach for Job Scheduling Optimizations Under Power Cap for ARM and Intel HPC Systems	142
<i>Dineshkumar Rajagopal (Bull Atos Technologies), Daniele Tafani (Leibniz Supercomputing Centre), Yiannis Georgiou (Bull Atos Technologies), David Glesser (Bull Atos Technologies), and Michael Ott (Leibniz Supercomputing Centre)</i>	

A Memory Congestion-Aware MPI Process Placement for Modern NUMA Systems	152
<i>Mulya Agung (Tohoku University), Muhammad Alfian Amrizal (Tohoku University), Kazuhiko Komatsu (Tohoku University), Ryusuke Egawa (Tohoku University), and Hiroyuki Takizawa (Tohoku University)</i>	
Expander: Lock-Free Cache for a Concurrent Data Structure	162
<i>Pooja Aggarwal (IBM Research Labs Bangalore) and Smruti R. Sarangi (IIT Delhi India)</i>	
Adaptive Code Refinement: A Compiler Technique and Extensions to Generate Self-Tuning Applications	172
<i>Maxime Schmitt (University of Strasbourg, INRIA), Philippe Helluy (University of Strasbourg, INRIA), and Cédric Bastoul (University of Strasbourg)</i>	

Keynote 2

Machine Learning @ Amazon	182
<i>Rajeev Rastogi (Amazon)</i>	

Technical Session 4: Big Data, Machine Learning and Optimization

Parallel Deep Convolutional Neural Network Training by Exploiting the Overlapping of Computation and Communication	183
<i>Sunwoo Lee (Northwestern University), Dipendra Jha (Northwestern University), Ankit Agrawal (Northwestern University), Alok Choudhary (Northwestern University), and Wei-keng Liao (Northwestern University)</i>	
Parallel Dynamic Data Driven Approaches for Synthetic Aperture Radar	193
<i>Adeesha Wijayasiri (University Of Florida - Gainesville Campus), Tania Banerjee (University Of Florida - Gainesville Campus), Sanjay Ranka (University Of Florida - Gainesville Campus), Sartaj Sahni (University Of Florida - Gainesville Campus), and Mark Schmalz (University Of Florida - Gainesville Campus)</i>	
ARM Wrestling with Big Data: A Study of Commodity ARM64 Server for Big Data Workloads	203
<i>Jayanth Kalyanasundaram (Indian Institute of Science) and Yogesh Simmhan (Indian Institute of Science)</i>	
MPI-LiFE: Designing High-Performance Linear Fascicle Evaluation of Brain Connectome with MPI	213
<i>Shashank Gugnani (The Ohio State University), Xiaoyi Lu (The Ohio State University), Franco Pestilli (Indiana University), Cesar Caiafa (Indiana University), and Dhabaleswar K. (DK) Panda (The Ohio State University)</i>	
Reducing Network Congestion and Synchronization Overhead During Aggregation of Hierarchical Data	223
<i>Sidharth Kumar (SCI, University of Utah), Duong Hoang (SCI, University of Utah), Steve Petruzza (SCI, University of Utah), John Edwards (Idaho State University), and Valerio Pascucci (SCI, University of Utah)</i>	
Fast Parallel Randomized QR with Column Pivoting Algorithms for Reliable Low-Rank Matrix Approximations	233
<i>Jianwei Xiao (UC Berkeley), Ming Gu (UC Berkeley), and Julien Langou (University of Colorado Denver)</i>	

Technical Session 5: Graph Algorithms and GPU

An X10-Based Distributed Streaming Graph Database Engine	243
<i>Miyuru Dayarathna (WSO2 Inc. / JST CREST), Sathya Bandara (University of Moratuwa), Nandula Jayamaha (University of Moratuwa), Mahen Herath (University of Moratuwa), Achala Madhushan (University of Moratuwa), Sanath Jayasena (University of Moratuwa), and Toyotaro Suzumura (IBM T.J. Watson Research Center)</i>	
GPU-Centric Communication on NVIDIA GPU Clusters with InfiniBand: A Case Study with OpenSHMEM ..	253
<i>Sreeram Potluri (NVIDIA Corporation), Anshuman Goswami (NVIDIA Corporation), Davide Rossetti (NVIDIA Corporation), C.J. Newburn (NVIDIA Corporation), Manjunath Gorentla Venkata (Oak Ridge National Laboratory), and Neena Imam (Oak Ridge National Laboratory)</i>	
Distributed Algorithm for High-Utility Subgraph Pattern Mining Over Big Data Platforms	263
<i>Alind Khare (IIIT-D), Vikram Goyal (IIIT-Delhi), Srikanth Baride (IIIT-Delhi), Sushil K. Prasad (Georgia State University), Michael McDermott (Georgia State University), and Dhara Shah (Georgia State University)</i>	
ReCALL: Reordered Cache Aware Locality Based Graph Processing	273
<i>Kartik Lakhota (University of Southern California), Shreyas Singapura, Rajgopal Kannan (US Army Research Lab), and Viktor Prasanna (University of Southern California)</i>	
Characterization of Data Movement Requirements for Sparse Matrix Computations on GPUs	283
<i>Süreyya Emre Kurt (The Ohio State University), Vineeth Thumma (The Ohio State University), Changwan Hong (The Ohio State University), Aravind Sukumaran-Rajam (The Ohio State University), and P. Sadayappan (The Ohio State University)</i>	
Applying Graph Analytics to Understand Compute Core Usage and Publication Trends in a Petascale Supercomputing Facility	294
<i>Sangkeun Lee (Oak Ridge National Laboratory), Sudharshan S. Vazhkudai (Oak Ridge National Laboratory), and Raghul Gunasekaran (Oak Ridge National Laboratory)</i>	

Keynote 3

Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales	306
<i>Ian Foster (Argonne National Lab and University of Chicago)</i>	

Technical Session 6: System Software

Integrating External Resources with a Task-Based Programming Model	307
<i>Zhihao Jia (Stanford University), Sean Treichler (NVIDIA), Galen Shipman (Los Alamos National Laboratory), Michael Bauer (NVIDIA), Noah Watkins (UC Santa Cruz), Carlos Maltzahn (UC Santa Cruz), Patrick McCormick (Los Alamos National Laboratory), and Alex Aiken (Stanford University)</i>	

Enabling Dependability-Driven Resource Use and Message Log-Analysis for Cluster System Diagnosis	317
<i>Edward Chuah (The Alan Turing Institute & The University of Warwick), Arshad Jhumka (The University of Warwick), Samantha Alt (Intel Corporation), Theo Damoulas (The Alan Turing Institute & The University of Warwick), Nentawe Gurumdimma (The University of Jos), Marie-Christine Sawley (Intel Corporation), William L. Barth (Texas Advanced Computing Center), Tommy Minyard (Texas Advanced Computing Center), and James C. Browne (The University of Texas at Austin)</i>	
Context-Aware Memory Profiling for Speculative Parallelism	328
<i>Changsu Kim (POSTECH), Juhyun Kim (POSTECH), Juwon Kang (POSTECH), Jae W. Lee (Seoul National University), and Hanjun Kim (POSTECH)</i>	
Lifting Barriers Using Parallel Polyhedral Regions	338
<i>Harenome Razanajato (University of Strasbourg and INRIA), Cédric Bastoul (University of Strasbourg and INRIA), and Vincent Loechner (University of Strasbourg and INRIA)</i>	
Exploiting Common Neighborhoods to Optimize MPI Neighborhood Collectives	348
<i>Seyed Hessamedin Mirsadeghi (Queen's University), Jesper Larsson Träff (Vienna University of Technology (TU Wien)), Pavan Balaji (Argonne National Laboratory), and Ahmad Afsahi (Queen's University)</i>	
Efficient Fork-Join on GPUs Through Warp Specialization	358
<i>Arpith Chacko Jacob (IBM T.J. Watson Research Center), Alexandre E Eichenberger (IBM T.J. Watson Research Center), Hyojin Sung (IBM T.J. Watson Research Center), Samuel F Antao (IBM T.J. Watson Research Center), Gheorghe-Teodor Bercea (IBM T.J. Watson Research Center), Carlo Bertolli (IBM T.J. Watson Research Center), Alexey Bataev (IBM T.J. Watson Research Center), Tian Jin (IBM T.J. Watson Research Center), Tong Chen (IBM T.J. Watson Research Center), Zehra Sura (IBM T.J. Watson Research Center), Georgios Rokos (IBM T.J. Watson Research Center), and Kevin O'Brien (IBM T.J. Watson Research Center)</i>	

Technical Session 7: GPU Frameworks and Applications

Thrust++: Extending Thrust Framework for Better Abstraction and Performance	368
<i>Ajai V. George (BITS Pilani KK Birla Goa Campus), Sankar Manoj (BITS Pilani KK Birla Goa Campus), Sanket R. Gupte (BITS Pilani KK Birla Goa Campus), Sayantan Mitra (Siemens Technology & Services Pvt Ltd., Bangalore, India), and Santonu Sarkar (BITS Pilani KK Birla Goa Campus)</i>	
A Novel Implementation of 2D3V Particle-in-Cell (PIC) Algorithm for Kepler GPU Architecture	378
<i>Harshil Shah (DA-IICT, Gandhinagar), Siddharth Kamaria (DA-IICT, Gandhinagar), Riddhesh Markandeya (DA-IICT, Gandhinagar), Miral Shah (DA-IICT, Gandhinagar), and Bhaskar Chaudhury (DA-IICT, Gandhinagar)</i>	
Parallelizing Hines Matrix Solver in Neuron Simulations on GPU	388
<i>Dharma Teja Vooturi (International Institute of Information Technology - Hyderabad), Kishore Kothapalli (International Institute of Information Technology - Hyderabad), and Upinder Singh Bhalla (National Center for Biological Sciences - Tata Institute of Fundamental Research.)</i>	

Building Halo Merger Trees from the Q Continuum Simulation	398
<i>Esteban Rangel (Northwestern University), Nicholas Frontiere</i>	
<i>(University of Chicago), Salman Habib (Argonne National Laboratory),</i>	
<i>Katrin Heitmann (Argonne National Laboratory), Wei-keng Liao</i>	
<i>(Northwestern University), Ankit Agrawal (Northwestern University),</i>	
<i>and Alok Choudhary (Northwestern University)</i>	
A Memory-Efficient GPU Method for Hamming and Levenshtein Distance Similarity	408
<i>Andrew Todd (University of Missouri), Marziyeh Nourian (North Carolina</i>	
<i>State University), and Michela Becchi (North Carolina State</i>	
<i>University)</i>	
Author Index	419