

HiPC 2017 Preliminary Technical Program

18 Oct 2017

[Click here to download a PDF of this page](#)

Monday Dec 18, 2017 - Day 1

HiPC 2017 Workshops & Tutorials

Tuesday Dec 19, 2017 - Day 2

Keynote Address

End of Moore's Law: Or, a computer architect's mid-life crisis?

Parthasarathy Ranganathan (Google)

Technical Session 1: Graph Algorithms

Exact and Parallel Triangle Counting in Dynamic Graphs

Devavret Makkar (Georgia Institute of Technology), David A. Bader (Georgia Institute of Technology), Oded Green (Georgia Institute of Technology)

Shared-memory Graph Truss Decomposition

Humayun Kabir (The Pennsylvania State University), Kamesh Madduri (The Pennsylvania State University)

Approximate Computing Techniques for Iterative Graph Algorithms

Ajay Panyala (Pacific Northwest National Laboratory), Omer Subasi (Pacific Northwest National Laboratory), Mahantesh Halappanavar (Pacific Northwest National Laboratory), Ananth Kalyanaraman (Washington State University), Sriram Krishnamoorthy (Pacific Northwest National Lab)

Scalable Exact Parent Sets Identification in Bayesian Networks Learning with Apache Spark

Subhadeep Karan (University at Buffalo), Jaroslaw Zola (University at Buffalo)

Parallel Exact Dynamic Bayesian Network Structure Learning with Application to Gene Networks

Vasimuddin Md (Indian Institute of Technology Bombay), Srinivas Aluru (Georgia Institute of Technology)

FIX: A Distributed MIS Algorithm

Thejaka Amila Kanewala (Indiana University), Marcin Zalewski (Indiana University), Andrew Lumsdaine (Pacific Northwest Natl Lab and U Washington)

Technical Session 2: Architecture and Communication

Designing Registration Caching Free High-Performance MPI Library with Implicit On-Demand Paging (ODP) of InfiniBand

Mingzhe Li (The Ohio State University), Xiaoyi Lu (The Ohio State University), Hari Subramoni (The Ohio State University), Dhableswar Panda (The Ohio State University)

Last Level Collective Hardware Prefetching For Data-Parallel Applications

George Michelogiannakis (LBNL), John Shalf (Lawrence Berkeley National Laboratory)

Kernel-assisted Communication Engine for MPI on Emerging Manycore Processors

Jahanzeb Maqbool Hashmi (The Ohio State University), Khaled Hamidouche (The Ohio State University), Hari Subramoni (The Ohio State University), Dhableswar Panda (The Ohio State University)

Support for Power Efficient Proactive Cooling Mechanisms

Bilge Acun (University of Illinois at Urbana-Champaign), Eun Kyung Lee (IBM T.J. Watson Research Center), Yoonho Park (IBM T.J. Watson Research Center), Laxmikant Kale (University of Illinois at Urbana-Champaign)

Redundant Arithmetic based High Speed Carry Free Hybrid Adders with Built-In Scan Chain on FPGAs

Ayan Palchadhuri (IIT Kharagpur), Anindya Sundar Dhar (Indian Institute of Technology Kharagpur)

ConvLight: A Convolutional Neural Network Accelerator with Neuromorphic Photonic-based Computing

Dharanidhar Dang (Texas A&M University), Jyotikrishna Dass (Texas A&M University), Rabi Mahapatra (Texas A&M University)

Technical Session 3: Algorithms

Provably Efficient Scheduling of Dynamically Allocating Programs on Parallel Cache Hierarchies

Harsha Vardhan Simhadri (Microsoft Research), Guy Blelloch (Carnegie Mellon University), Phillip Gibbons (Carnegie Mellon University)

Further Explorations in State-Space Search for Optimal Task Scheduling

Michael Orr (University of Auckland), Oliver Sinnen (University of Auckland)

A Novel Approach for Job Scheduling Optimizations under Power Cap for ARM and Intel HPC Systems

Dineshkumar Rajagopal (Bull Atos Technologies), Daniele Tafani (Leibniz Supercomputing Centre), Yiannis Georgiou (Bull Atos Technologies), David Glesser (Bull Atos Technologies), Michael Ott (Leibniz Supercomputing Centre)

A Memory Congestion-aware MPI Process Placement for Modern NUMA Systems

Mulya Agung (Tohoku University), Muhammad Alfian Amrizal (Tohoku University), Kazuhiko Komatsu (Tohoku University), Ryusuke Egawa (Tohoku University), Hiroyuki Takizawa (Tohoku University)

Expander: Lock-free Cache for a Concurrent Data Structure

Pooja Aggarwal (IBM Research, India), Smruti Sarangi (IIT Delhi)

Adaptive Code Refinement: A Compiler Technique and Extensions to Generate Self-Tuning Applications

Maxime Schmitt (Univ. of Strasbourg, INRIA), Philippe Helluy (Univ. of Strasbourg, INRIA), Cédric Bastoul (University of Strasbourg)

Wednesday Dec 20, 2017 - Day 3

Keynote Address

Machine Learning @ Amazon

Rajeev Rastogi (Amazon)

Technical Session 4: Big Data, Machine Learning and Optimization

Parallel Deep Convolutional Neural Network Training by Exploiting the Overlapping of Computation and Communication

Sunwoo Lee (Northwestern University), Dipendra Jha (Northwestern University), Ankit Agrawal (Northwestern University), Alok Choudhary (Northwestern University), Wei-Keng Liao (Northwestern University)

Parallel Dynamic Data Driven Approaches for Synthetic Aperture Radar

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), Mark Schmalz (University of Florida - Gainesville Campus)

ARM Wrestling with Big Data: A Study of Commodity ARM Server for Big Data Workloads

Jayanth Kalyanasundaram (Indian Institute of Science), Yogesh Simmhan (Indian Institute of Science)

MPI-LiFE: Designing High-Performance Linear Fascicle Evaluation of Brain Connectome with MPI

Shashank Gugnani (The Ohio State University), Xiaoyi Lu (The Ohio State University), Franco Pestilli (Indiana University), Cesar Caiafa (Indiana University), Dhableswar Panda (The Ohio State University)

Reducing network congestion and global communication bottlenecks during aggregation on Torus and Dragonfly topologies for writing hierarchical data

Sidharth Kumar (SCI, University of Utah), Duong Hoang (SCI, University of Utah), Steve Petruzza (SCI, University of Utah and University of Rome "Tor Vergata"), John Edwards (Idaho State University), Valerio Pascucci (SCI, University of Utah)

Fast Parallel Randomized QR with Column Pivoting Algorithms for Reliable Low-rank Matrix Approximations

Jianwei Xiao (UC Berkeley), Ming Gu (UC Berkeley), Julien Langou (University of Colorado Denver)

Technical Session 5: Graph Algorithms and GPU

An X10 based Distributed Streaming Graph Database Engine

Miyuru Dayarathna (WSO2 Inc.), 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), Toyotaro Suzumura (IBM T.J. Watson Research Center)

GPU-centric Communication on NVIDIA GPU Clusters with InfiniBand: A Case Study with OpenSHMEM

Sreeram Potluri (NVIDIA Corporation), Anshuman Goswami (NVIDIA Corporation), Davide Rossetti (NVIDIA Corporation), Manjunath Gorentla Venkata (Oak Ridge National Laboratory), Neena Inam (Oak Ridge National Laboratory), Chris J. Newburn (NVIDIA Corporation)

Applying Graph Analytics to Understand Compute Core Usage and Publication Trends in a Petascale Supercomputing Facility

Sangkeun Lee (Oak Ridge National Laboratory), Sudharshan S. Vazhkudai (Oak Ridge National Laboratory), Raghul Gunasekaran (Oak Ridge National Laboratory)

Distributed Algorithm for High-Utility Subgraph Pattern Mining Over Big Data Platform

Alind Khare (IIIT-Delhi), Vikram Goyal (IIIT-Delhi), Srikant Baride (IIIT-Delhi), Michael McDermott (Georgia State University), Dhara Shah (Georgia State University), Sushil Prasad (Georgia State University)

ReCALL: Reordered Cache aware LocaliTy based Graph Processing

Kartik Lakhota (University of Southern California), Shreyas Singapura (University of Southern California), Rajgopal Kannan (University of Southern California), Viktor Prasanna (University of Southern California)

Characterization of data movement requirements for sparse matrix computations on GPUs

Sureyya 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), Sadayappan P (The Ohio State University)

Thursday Dec 21, 2017 - Day 4

Keynote Address

Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales

Ian Foster (Argonne National Lab and U. Chicago)

Technical Session 6: System Software

Integrating External Resources with a Task-Based Programming Model

Zihao Jia (Stanford University), Sean Treichler (NVIDIA Research), Galen Shipman (Los Alamos National Laboratory), Michael Bauer (NVIDIA Research), Noah Watkins (UC Santa Cruz), Carlos Maltzahn (UC Santa Cruz), Pat McCormick (Los Alamos National Laboratory), Alex Aiken (Stanford University)

Enabling Dependability-Driven Resource Use and Message Log-Analysis for Cluster System Diagnosis

Edward Chuah (The Alan Turing Institute & The University of Warwick), Arshad Jhumka (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), Bill Barth (The University of Texas at Austin), Tommy Minyard (The Texas Advanced Computing Center), James Browne (University of Texas)

DAOS for Extreme-scale Systems in Scientific Applications

Michael Breitenfeld (The HDF Group), Neil Fortner (The HDF Group), Jordan Henderson (The HDF Group), Jerome Soumagne (The HDF Group), Mohamad Chaarawi (Intel), Johann Lombardi (Intel), Quincey Koziol (Lawrence Berkeley National Laboratory)

Context-Aware Memory Profiling for Speculative Parallelism

Changsu Kim (POSTECH), Juhyun Kim (POSTECH), Juwon Kang (POSTECH), Jae W. Lee (Seoul National University), Hanjun Kim (POSTECH)

Lifting Barriers Using Parallel Polyhedral Regions

Harenome Ranaivoarivony-Razanajato (CAMUS team, INRIA Nancy Grand-Est and University of Strasbourg), Cédric Bastoul (CAMUS team, INRIA Nancy Grand-Est and University of Strasbourg), Vincent Loechner (CAMUS team, INRIA Nancy Grand-Est and University of Strasbourg)

Exploiting Common Neighborhoods to Optimize MPI Neighborhood Collectives

Seyed Hessamedin Mirsadeghi (Queen's University), Jesper Larsson Traff (Vienna University of Technology/TU Wien), Pavan Balaji (Argonne National Laboratory), Ahmad Afsahi (Queen's University)

Technical Session 7: GPU Frameworks and Applications

Thrust++: Extending Thrust Framework for Better Abstraction and Performance

Ajai George (BITS Pilani K K Birla Goa Campus), Sankar Manoj (BITS Pilani K K Birla Goa Campus), Sanket Gupte (BITS Pilani K K Birla Goa Campus), Sayantan Mitra (Siemens Technology & Services Pvt Ltd, Bangalore, India), Santonu Sarkar (BITS Pilani K K Birla Goa Campus)

Efficient Fork-Join on GPUs through Warp Specialization

Arpith Jacob (IBM T.J. Watson Research Center), Alexandre Eichenberger (IBM T.J. Watson Research Center), Hyojin Sung (IBM T.J. Watson Research Center), Samuel Antao (IBM Research UK), Gheorghe-Teodor Bercea (IBM T.J. Watson Research Center), Carlo Bertolli (IBM TJ Watson Research Center), Alexey Bataev (IBM Research), Tian Jin (IBM Research), Tong Chen (IBM research), Zehra Sura (IBM T.J. Watson Research Center, Yorktown Heights, NY 10598), Rokos Georgios (IBM T.J. Watson Research Center), Kevin O'Brien (IBM T.J. Watson Research Center)

A Novel Implementation of 2D3V Particle-In-Cell (PIC) Algorithm for Kepler GPU Architectures

Harshil Shah (DA-IICT, Gandhinagar), Siddharth Kamaria (DA-IICT, Gandhinagar), Riddhesh Markandeya (DA-IICT, Gandhinagar), Miral Shah (DA-IICT, Gandhinagar), Bhaskar Chaudhury (DA-IICT, Gandhinagar)

Parallelizing Hines Matrix Solver in Neuron Simulations on GPU

Dharma Teja Vooturi (International Institute of Information Technology, Hyderabad), Kishore Kothapalli (International Institute of Information Technology, Hyderabad), Upinder Bhalla (National Centre for Biological Sciences)

Building Halo Merger Trees from the Q Continuum Simulation

Esteban Rangel (Northwestern University), Nicholas Frontiere (University of Chicago), Salman Habib (Argonne National Laboratory), Katrin Heitmann (Argonne National Laboratory), Wei-Keng Liao (Northwestern University), Ankit Agrawal (Iowa State University), Alok Choudhary (Northwestern University)

A Memory-Efficient GPU Method for Hamming and Levenshtein Distance Similarity

Andrew Todd (University of Missouri), Marziyeh Nourian (North Carolina State University), Michela Becchi (North Carolina State University)