

HiPC 2017 Advance Program

(With Schedule)

25 Nov 2017 Edition

Monday Dec 18, 2017 – Day 1

HiPC 2017 **Workshops, Tutorials, BOF's**

Tuesday Dec 19, 2017 – Day 2

HiPC 2017 **SRS, Industry Exhibits, IRUS's, Sponsored Technical Sessions**

Keynote Address (8:30-9:30)

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

Parthasarathy Ranganathan (Google)

Technical Session 1: Graph Algorithms (10:00-12:00)

Exact and Parallel Triangle Counting in Dynamic Graphs

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

Shared-memory Graph Truss Decomposition

Humayun Kabir and Kamesh Madduri (The Pennsylvania State University)

Approximate Computing Techniques for Iterative Graph Algorithms

Ajay Panyala, Omer Subasi, and Mahantesh Halappanavar (Pacific Northwest National Laboratory), Ananth Kalyanaraman (Washington State University), Daniel Chavarría-Miranda (Trovares, Inc.), Sriram Krishnamoorthy (Pacific Northwest National Lab)

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

Subhadeep Karan and 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)

Parallel Asynchronous Distributed-Memory Maximal Independent Set Algorithm with Work Ordering

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

Technical Session 2: Architecture and Communication (1:00-3:00)

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

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

Last Level Collective Hardware Prefetching For Data-Parallel Applications

George Michelogiannakis and John Shalf (Lawrence Berkeley National Laboratory)

Kernel-assisted Communication Engine for MPI on Emerging Manycore Processors
Jahanzeb Maqbool Hashmi, Khaled Hamidouche, Hari Subramoni, and 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 Palchaudhuri and Anindya Sundar Dhar (Indian Institute of Technology Kharagpur)

ConvLight: A Convolutional Neural Network Accelerator with Neuromorphic Photonic-based Computing
Dharanidhar Dang, Jyotikrishna Dass, and Rabi Mahapatra (Texas A&M University)

Technical Session 3: Algorithms (3:15-5:15)

Provably Efficient Scheduling of Dynamically Allocating Programs on Parallel Cache Hierarchies
Harsha Vardhan Simhadri (Microsoft Research), Guy Blelloch and Phillip Gibbons (Carnegie Mellon University)

Further Explorations in State-Space Search for Optimal Task Scheduling
Michael Orr and 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 and David Glesser (Bull Atos Technologies), Michael Ott (Leibniz Supercomputing Centre)

A Memory Congestion-aware MPI Process Placement for Modern NUMA Systems
Mulya Agung, Muhammad Alfian Amrizal, Kazuhiko Komatsu, Ryusuke Egawa, and 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 and Philippe Helluy (Univ. of Strasbourg, INRIA), Cédric Bastoul (University of Strasbourg)

Plenary Panel: (5:30-7:30)

HPC in India

Featuring presentations by leading researchers from academia, industry and Indian government R&D labs.

Wednesday Dec 20, 2017 – Day 3

HiPC 2017 SRS, Industry Exhibits, IRUS's, Sponsored Technical Sessions

Keynote Address (8:30-9:30)

Machine Learning @ Amazon

Rajeev Rastogi (Amazon)

Technical Session 4: Big Data, Machine Learning and Optimization (10:00-12:00)

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

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

Parallel Dynamic Data Driven Approaches for Synthetic Aperture Radar

Adeesha Wijayasiri, Tania Banerjee, Sanjay Ranka, Sartaj Sahni, and Mark Schmalz (University of Florida – Gainesville Campus)

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

Jayanth Kalyanasundaram and Yogesh Simmhan (Indian Institute of Science)

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

Shashank Gugnani and Xiaoyi Lu (The Ohio State University), Franco Pestilli and Cesar Caiafa (Indiana University), Dhabaleswar 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 and 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 and Ming Gu (UC Berkeley), Julien Langou (University of Colorado Denver)

Technical Session 5: Graph Algorithms and GPU (1:00-3:00)

An X10 based Distributed Streaming Graph Database Engine

Miyuru Dayarathna (WSO2 Inc.), Sathya Bandara, Nandula Jayamaha, Mahen Herath, Achala Madhushan, and 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, Anshuman Goswami, and Davide Rossetti (NVIDIA Corporation), Manjunath Gorentla Venkata and Neena Inam (Oak Ridge National Laboratory), Chris J. Newburn (NVIDIA Corporation)

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

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

ReCALL: Reordered Cache aware LocaliTy based Graph Processing

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

Characterization of data movement requirements for sparse matrix computations on GPUs

Sureyya Emre Kurt, Vineeth Thumma, Changwan Hong, Aravind Sukumaran-Rajam, and Sadayappan P (The Ohio State University)

Applying Graph Analytics to Understand Compute Core Usage and Publication Trends in a Petascale Supercomputing Facility
Sangkeun Lee, Sudharshan S. Vazhkudai, and Raghul Gunasekaran (Oak Ridge National Laboratory)

Plenary Panel: (3:30-5:30)

Machine Learning and Mathematical Modeling – Towards Robust, Reliable and Certifiable Machine Learning

This event will offer HiPC attendees expert perspectives on mathematical foundations and the potential gaps in the underlying theory that must be overcome to provide reliable, robust, and certifiable outcomes from machine learning.

Thursday Dec 21, 2017 – Day 4

Keynote Address (8:30-9:30)

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

Ian Foster (Argonne National Lab and University of Chicago)

Technical Session 6: System Software (10:00-12:00)

Integrating External Resources with a Task-Based Programming Model

Zhihao Jia (Stanford University), Sean Treichler (NVIDIA Research), Galen Shipman (Los Alamos National Laboratory), Michael Bauer (NVIDIA Research), Noah Watkins and 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 (University of Texas at Austin), Tommy Minyard (The Texas Advanced Computing Center), James Browne (University of Texas at Austin)

Context-Aware Memory Profiling for Speculative Parallelism

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

Lifting Barriers Using Parallel Polyhedral Regions

Harenome Ranaivoarivony-Razanajato, Cédric Bastoul, and 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)

Efficient Fork-Join on GPUs through Warp Specialization

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

Technical Session 7: GPU Frameworks and Applications (1:00-3:00)

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

Ajai George, Sankar Manoj, and 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)

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

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

Parallelizing Hines Matrix Solver in Neuron Simulations on GPU

Dharma Teja Vooturi and 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 and 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 and Michela Becchi (North Carolina State University)