

HiPC 2018 Advance Program

(With Schedule)

10 Dec 2018 Edition

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

[Sunday Dec 16, 2018 – Pre-conference Satellite Events](#)

[Industry Tutorials 1 & 2](#) – Free to attendees but require pre-registration with presenter.

Monday Dec 17, 2018 – Day 1

[HiPC 2018 Workshops](#)

Tuesday Dec 18, 2018 - Day 2

[HiPC 2018 SRS](#), [Industry Exhibits](#), [Sponsored Technical Sessions](#)

Keynote Address: (8:30-9:30)

[Looking Under the Hood of Deep Neural Networks](#)

Balaraman Ravindran, Indian Institute of Technology, Madras

Technical Session 1: Learning (10:00-12:00)

Accelerating TensorFlow with Adaptive RDMA-based gRPC

Rajarshi Biswas, Xiaoyi Lu, and Dhabaleswar K. Panda (The Ohio State University, USA)

Balancing Stragglers Against Staleness in Distributed Deep Learning

Saurav Basu and Vaibhav Saxena (IBM Research, New Delhi, India), Rintu Panja (Indian Institute of Science, Bangalore, India) and Ashish Verma (IBM Research, USA)

Parallel Nonnegative CP Decomposition of Dense Tensors

Grey Ballard and Koby Hayashi (Wake Forest University, USA) and Ramakrishnan Kannan (Oak Ridge National Laboratory, USA)

Sampled Dense Matrix Multiplication for High-Performance Machine Learning

Israt Nisa, Aravind Sukumaran Rajam, Süreyya Emre Kurt, Changwan Hong and P Sadayappan (The Ohio State University, USA)

DeepHyper: Asynchronous Hyperparameter Search for Deep Neural Networks

Prasanna Balaprakash, Michael Salim, Thomas D. Uram, Venkat Vishwanath, and Stefan M. Wild (Argonne National Laboratory, USA)

Technical Session 2: Graph Algorithms (1:00-3:00)

Synchronization-Avoiding Graph Algorithms

Jesun Sahariar Firoz (Indiana University, USA), Marcin Zalewski (Pacific Northwest National Laboratory, USA), Thejaka Kanewala (Indiana University, USA) and Andrew Lumsdaine (Pacific Northwest National Laboratory and University of Washington, USA)

Shared-Memory Parallel Maximal Clique Enumeration

Apurba Das, Seyed-Vahid Sanei-Mehri and Srikanta Tirthapura (Iowa State University, USA)

Expediting Parallel Graph Connectivity Algorithms

Mihir Wadwekar and Kishore Kothapalli (International Institute of Information Technology, Hyderabad, India)

Adaptive Runtime Features For Distributed Graph Algorithms

Jesun Sahariar Firoz (Indiana University, USA), Marcin Zalewski and Joshua Suetterlein (Pacific Northwest National Laboratory, USA) and Andrew Lumsdaine (Pacific Northwest National Laboratory and University of Washington, USA)

Adaptive Pattern Matching with Reinforcement Learning for Dynamic Graphs

Hiroki Kanezashi (Tokyo Institute of Technology, Japan and IBM T.J. Watson Research Center, USA), Toyotaro Suzumura (IBM T.J. Watson Research Center, USA and Barcelona Supercomputing Center, Spain), Dario Garcia-Gasulla (Barcelona Supercomputing Center, Spain), Min-Hwan (Columbia University and IBM T.J. Watson Research Center, USA) and Satoshi Matsuoka (RIKEN Center for Computational Science and Tokyo Institute of Technology, Japan)

Probabilistic Sequential Consistency in Social Networks

Priyanka Singla (IISc Bangalore, India), Shubhankar Suman Singh (IIT Delhi, India), K. Gopinath (IISc Bangalore, India) and Smruti Sarangi (IIT Delhi, India)

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

Improving Provisioned Power Efficiency in HPC Systems with GPU-CAPP

Kramer Straube, Jason Lowe-Power, Christopher Nitta, Matthew Farrens and Venkatesh Akella (University of California, Davis, USA)

Compiling SIMT Programs on Multi- and Many-core Processors with Wide Vector Units: A Case Study with CUDA

Hancheng Wu, John Ravi and Michela Becchi (North Carolina State University, USA)

Lossless parallel implementation of a Turbo Decoder on GPU

Karthikeyan Natarajan and Nitin Chandrhoodan (Indian Institute of Technology, Madras, India)

OC-DNN: Exploiting Advanced Unified Memory Capabilities in CUDA 9 and Volta GPUs for Out-of-Core DNN Training

Ammar Ahmad Awan, Ching-Hsiang Chu, Hari Subramoni, Xiaoyi Lu and Dhaleswar K. Panda (The Ohio State University, USA)

Acceleration of an Adaptive Cartesian Mesh CFD Solver in the Current Generation Processor Architectures

Harichand M V (Vikram Sarabhai Space Centre, India), Bharatkumar Sharma (Nvidia Graphics Pvt Ltd, India), G Sudhakaran and V Ashok (Vikram Sarabhai Space Centre, India)

Data-parallel Training of Generative Adversarial Networks on HPC Systems for HEP Simulations

Sofia Vallecorsa (CERN, Switzerland), Diana Moise (Cray Inc., Switzerland), Federico Carminati and Gul Rukh Khattak (CERN, Switzerland)

Wednesday Dec 19, Day 3

[HiPC 2018 SRS](#), [Industry Exhibits](#), [IRUS](#), [Sponsored Technical Sessions](#)

Keynote Address: (8:30-9:30)

The Future of Supercomputing

Marc Snir, University of Illinois at Urbana-Champaign

Technical Session 4: Linear Algebra and Fault Tolerance (10:00-12:00)

Making Strassen Matrix Multiplication Safe

Himeshi De Silva, John L. Gustafson and Weng-Fai Wong (National University of Singapore, Singapore)

Quantification, Trade-off Analysis, and Optimal Checkpoint Placement for Reliability and Availability

Omer Subasi, Ramakrishna Tipireddy and Sriram Krishnamoorthy (Pacific Northwest National Laboratory, USA)

A Novel Approach for Handling Soft Error in Conjugate Gradients

Marissa Renardy, Muhammed Emin Ozturk, Yukun Li, Gagan Agrawal and Ching-Shan Chou (The Ohio State University, USA)

Characterization of the Impact of Soft Errors on Iterative Methods

Burcu Ozcelik Mutlu (Pacific Northwest National Laboratory, USA and Polytechnic University of Catalonia, Spain), Gokcen Kestor and Joseph Manzano (Pacific Northwest National Laboratory, USA), Osman Unsal (Barcelona Supercomputing Center, Spain), Samrat Chatterjee and Sriram Krishnamoorthy (Pacific Northwest National Laboratory, USA)

Technical Session 5: Algorithms and Data Analysis (1:00-3:00)

Workflow Simulation Aware and Multi-Threading Effective Task Scheduling for Heterogeneous Computing

Vasilios Kelefouras (University of Plymouth, UK) and Karim Djemame (University of Leeds, West Yorkshire, UK)

Dynamic Count-Min Sketch for Analytical Queries over Continuous Data Streams

Xiaobo Zhu (Institute of Information Engineering, Chinese Academy of Sciences, Beijing and School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China), Guangjun Wu (Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China), Hong Zhang (National Computer Network Emergency Response technical Team/Coordination Center of China, Beijing, China), Shupeng Wang (Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China) and Bingnan Ma (National Computer Network Emergency Response technical Team/Coordination Center of China, Beijing, China)

Scalable Proximity-Based Methods for Large-Scale Analysis of Atom Probe Data

Hao Lu, Sudip K. Seal and Jonathan D. Poplawsky (Oak Ridge National Laboratory, USA)

A Shared-Memory Algorithm for Updating Single-Source Shortest Paths in Large Weighted Dynamic Networks

Sriram Srinivasan (University of Nebraska Omaha, USA), Sara Riaz (University of Oregon, USA), Sajal K. Das (Missouri University of Science and Technology, USA), Boyana Norris (University of Oregon, USA) and Sanjukta Bhowmick (University of North Texas, USA)

Vidya: Performing Code-Block I/O Characterization for Data Access Optimization

Hariharan Devarajan, Anthony Kougkas, Prajwal Challa and Xian-He Sun (Illinois Institute of Technology, USA)

Decentralized Privacy-preserving Timed Execution in Blockchain-based Smart Contract Platforms

Chao Li and Balaji Palanisamy (University of Pittsburgh, USA)

25th Year Celebration (3:15-5:15) – [See details here](#)

Conference Banquet (starting at 6:30)

Thursday Dec 20, 2018 – Day 4

HiPC 2018 Conference Tutorials 3 & 4 – Open to all attendees; see here for details.

Keynote Address: (8:30-9:30)

Secure High-Performance Computer Architectures: Challenges and Opportunities

Srini Devadas, Massachusetts Institute of Technology

Technical Session 6: Applications and System Tools (10:00-12:00)

Why do Users Kill HPC Jobs?

Venkatesh-Prasad Ranganath and Daniel Andresen (Kansas State University, USA)

Code and Data Transformations to Address Garbage Collector Performance in Big Data Processing

Damon Fenacci, Hans Vandierendonck and Dimitrios S. Nikolopoulos (Queen's University Belfast, Ireland)

Share-a-GPU: Providing Simple and Effective Time-Sharing on GPUs

Shaleen Garg, Kishore Kothapalli and Suresh Purini (International Institute of Information Technology, Hyderabad, India)

A Performance Prediction Framework for Irregular Applications

Gangyi Zhu and Gagan Agrawal (The Ohio State University, USA)

Achieving Performance and Programmability for MapReduce(-like) Frameworks

Jia Guo and Gagan Agrawal (The Ohio State University, USA)

Parallel Read Partitioning for Concurrent Assembly of Metagenomic Data

Vasudevan Rengasamy, Mahmut Kandemir, Paul Medvedev and Kamesh Madduri (The Pennsylvania State University, USA)

Conference Tutorial 3: (10:00-12:00)

Title: *[_HCLib: A Task-based Parallel Programming Model](#)*

Presenters:

Vivek Kumar, IIIT-Delhi, India

Vivek Sarkar, Georgia Tech., USA

Conference Tutorial 4: (1:00-3:00)

Title: *[A Language and Framework for Prototyping and Experimenting with Edge Oriented IoT](#)*

Presenter:

Muthucumar Maheswaran, McGill University, Canada