

# **2018 IEEE 25th International Conference on High Performance Computing (HiPC 2018)**

**Bengaluru, India  
17-20 December 2018**



**IEEE Catalog Number: CFP18176-POD  
ISBN: 978-1-5386-8387-3**

**Copyright © 2018 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:	CFP18176-POD
ISBN (Print-On-Demand):	978-1-5386-8387-3
ISBN (Online):	978-1-5386-8386-6
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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2018 IEEE 25th International Conference on High Performance Computing (HiPC) **HiPC 2018**

## Table of Contents

Message from the General and Vice-General Co-Chairs	x
Message from the Program Chairs	xii
Message from the Steering Chair	xiv
HiPC 2018 Committees	xv

### Keynote 1

Looking Under the Hood of Deep Neural Networks	1
<i>Balaraman Ravindran (Indian Institute of Technology, Madras)</i>	

### Technical Session 1: Learning

Accelerating TensorFlow with Adaptive RDMA-Based gRPC	2
<i>Rajarshi Biswas (The Ohio State University), Xiaoyi Lu (The Ohio State University), and Dhabaleswar K. Panda (The Ohio State University)</i>	
Balancing Stragglers Against Staleness in Distributed Deep Learning	12
<i>Saurav Basu (IBM Research), Vaibhav Saxena (IBM Research), Rintu Panja (Indian Institute of Science), and Ashish Verma (IBM Research)</i>	
Parallel Nonnegative CP Decomposition of Dense Tensors	22
<i>Grey Ballard (Wake Forest University), Koby Hayashi (Wake Forest University), and Kannan Ramakrishnan (Oak Ridge National Labs)</i>	
Sampled Dense Matrix Multiplication for High-Performance Machine Learning	32
<i>Israt Nisa (The Ohio State University), Aravind Sukumaran-Rajam (The Ohio State University), Sureyya Emre Kurt (The Ohio State University), Changwan Hong (The Ohio State University), and P. Sadayappan (The Ohio State University)</i>	
DeepHyper: Asynchronous Hyperparameter Search for Deep Neural Networks	42
<i>Prasanna Balaprakash (Argonne National Laboratory), Michael Salim (Argonne National Laboratory), Thomas Uram (Argonne National Laboratory), Venkat Vishwanath (Argonne National Laboratory), and Stefan Wild (Argonne National Laboratory)</i>	

## Technical Session 2: Graph Algorithms

Synchronization-Avoiding Graph Algorithms .52.....	
<i>Jesun Sahariar Firoz (Indiana University Bloomington), Marcin Zalewski (Northwest Institute for Advanced Computing, Pacific Northwest National Lab), Thejaka Kanewala (Indiana University Bloomington), and Andrew Lumsdaine (Northwest Institute for Advanced Computing, Pacific Northwest National Lab)</i>	
Shared-Memory Parallel Maximal Clique Enumeration .62.....	
<i>Apurba Das (Iowa State University), Seyed-Vahid Sanei-Mehri (Iowa State University), and Srikanta Tirthapura (Iowa State University)</i>	
Expediting Parallel Graph Connectivity Algorithms .72.....	
<i>Kishore Kothapalli (IIIT Hyderabad) and Mihir Wadwekar (IIIT Hyderabad)</i>	
Adaptive Runtime Features for Distributed Graph Algorithms .82.....	
<i>Jesun Sahariar Firoz (Indiana University Bloomington), Marcin Zalewski (Pacific Northwest National Laboratory), Joshua Suetterlein (Pacific Northwest National Laboratory), and Andrew Lumsdaine (Pacific Northwest National Laboratory and University of Washington)</i>	
Adaptive Pattern Matching with Reinforcement Learning for Dynamic Graphs .92.....	
<i>Hiroki Kanezashi (Tokyo Institute of Technology), Toyotaro Suzumura (IBM), Dario Garcia-Gasulla (Barcelona Supercomputing Center), Min-Hwan Oh (Columbia University), and Satoshi Matsuoka (RIKEN Center for Computational Science)</i>	
Probabilistic Sequential Consistency in Social Networks .102.....	
<i>Priyanka Singla (IISc), Shubhankar Suman Singh (IIT Delhi), K. Gopinath (IISc), and Smruti Sarangi (IIT Delhi)</i>	

## Technical Session 3: GPUs

Improving Provisioned Power Efficiency in HPC Systems with GPU-CAPP .112.....	
<i>Kramer Straube (University of California, Davis), Jason Lowe-Power (University of California, Davis), Christopher Nitta (University of California, Davis), Matthew Farrens (University of California, Davis), and Venkatesh Akella (University of California, Davis)</i>	
Compiling SIMT Programs on Multi- and Many-Core Processors with Wide Vector Units: A Case Study with CUDA .123.....	
<i>Hancheng Wu (North Carolina State University), John Ravi (North Carolina State University), and Michela Becchi (North Carolina State University)</i>	
Lossless Parallel Implementation of a Turbo Decoder on GPU .133.....	
<i>Karthikeyan Natarajan (Indian Institute of Technology, Madras) and Nitin Chandrachoodan (Indian Institute of Technology, Madras)</i>	
OC-DNN: Exploiting Advanced Unified Memory Capabilities in CUDA 9 and Volta GPUs for Out-of-Core DNN Training .143.....	
<i>Ammar Ahmad Awan (The Ohio State University), Ching-Hsiang Chu (The Ohio State University), Hari Subramoni (The Ohio State University), Xiaoyi Lu (The Ohio State University), and Dhableswar K. (DK) Panda (The Ohio State University)</i>	

Acceleration of an Adaptive Cartesian Mesh CFD Solver in the Current Generation Processor Architectures .153.....	
<i>Harichand M V (Vikram Sarabhai Space Centre), Bharatkumar Sharma (Nvidia Graphics Pvt Ltd), G Sudhakaran (Vikram Sarabhai Space Centre), and V Ashok (Vikram Sarabhai Space Centre)</i>	
Data-Parallel Training of Generative Adversarial Networks on HPC Systems for HEP Simulations .162.....	
<i>Sofia Vallecorsa (CERN Switzerland), Diana Moise (Cray Inc. Switzerland), Federico Carminati (CERN Switzerland), and Gul Rukh Khattak (CERN Switzerland)</i>	

## Keynote 2

The Future of Supercomputing .172.....	
<i>Marc Snir (University of Illinois at Urbana-Champaign)</i>	

## Technical Session 4: Linear Algebra and Fault Tolerance

Making Strassen Matrix Multiplication Safe .173.....	
<i>Himeshi De Silva (National University of Singapore), John L. Gustafson (National University of Singapore), and Weng-Fai Wong (National University of Singapore)</i>	
Quantification, Trade-off Analysis, and Optimal Checkpoint Placement for Reliability and Availability .183.....	
<i>Omer Subasi (Pacific Northwest National Laboratory), Ramakrishna Tipireddy (Pacific Northwest National Laboratory), and Sriram Krishnamoorthy (Pacific Northwest National Laboratory)</i>	
A Novel Approach for Handling Soft Error in Conjugate Gradients .193.....	
<i>Muhammed Emin Ozturk (The Ohio State University), Marissa Renardy (The Ohio State University), Yukun Li (The Ohio State University), Gagan Agrawal (The Ohio State University), and Ching-Shan Chou (The Ohio State University)</i>	
Characterization of the Impact of Soft Errors on Iterative Methods .203.....	
<i>Burcu Ozelik Mutlu (Pacific Northwest National Laboratory &amp; Polytechnic University of Catalonia), Gokcen Kestor (Pacific Northwest National Laboratory), Joseph Manzano (Pacific Northwest National Laboratory), Osman Unsal (Barcelona Supercomputing Center), Samrat Chatterjee (Pacific Northwest National Laboratory), and Sriram Krishnamoorthy (Pacific Northwest National Laboratory)</i>	

## Technical Session 5: Algorithms and Data Analysis

Workflow Simulation Aware and Multi-threading Effective Task Scheduling for Heterogeneous Computing .215	
<i>Vasilios Kelefouras (University of Plymouth) and Karim Djemame</i>	

Dynamic Count-Min Sketch for Analytical Queries Over Continuous Data Streams .225.....	
<i>Xiaobo Zhu (Institute of Information Engineering, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Guangjun Wu (Institute of Information Engineering, Chinese Academy of Sciences), Hong Zhang (National Computer Network Emergency Response technical Team/Coordination Center of China), Shupeng Wang (Institute of Information Engineering, Chinese Academy of Sciences), and Bingnan Ma (National Computer Network Emergency Response technical Team/Coordination Center of China)</i>	
Scalable Proximity-Based Methods for Large-Scale Analysis of Atom Probe Data .235.....	
<i>Hao Lu (Oak Ridge National Laboratory), Sudip Seal (Oak Ridge National Laboratory), and Jonathan D. Poplawsky (Center for Nanophase Materials Sciences)</i>	
A Shared-Memory Parallel Algorithm for Updating Single-Source Shortest Paths in Large Dynamic Networks .245.....	
<i>Sriram Srinivasan (University of Nebraska), Sara Riazi (University of Oregon, Eugene), Boyana Norris (University of Nebraska), Sajal K. Das (Missouri University of Science and Technology, Rolla, ), and Sanjukta Bhowmick (University of North Texas)</i>	
Vidya: Performing Code-Block I/O Characterization for Data Access Optimization .255.....	
<i>Hariharan Devarajan (Illinois Institute of Technology Chicago), Anthony Kougkas (Illinois Institute of Technology Chicago), Prajwal Challa (Illinois Institute of Technology Chicago), and Xian-He Sun (Illinois Institute of Technology Chicago)</i>	
Decentralized Privacy-Preserving Timed Execution in Blockchain-Based Smart Contract Platforms .265.....	
<i>Chao Li (University of Pittsburgh) and Balaji Palanisamy (University of Pittsburgh)</i>	

### **Keynote 3**

Secure High-Performance Computer Architectures: Challenges and Opportunities .275.....	
<i>Srini Devadas (Massachusetts Institute of Technology)</i>	

### **Technical Session 6: Applications and System Tools**

Why do Users Kill HPC Jobs? .276.....	
<i>Venkatesh-Prasad Ranganath (Kansas State University) and Daniel Andresen (Kansas State University)</i>	
Code and Data Transformations to Address Garbage Collector Performance in Big Data Processing .284.....	
<i>Damon Fenacci (Queen's University Belfast), Hans Vandierendonck (Queen's University Belfast), and Dimitrios Nikolopoulos (Queen's University Belfast)</i>	
Share-a-GPU: Providing Simple and Effective Time-Sharing on GPUs .294.....	
<i>Shaleen Garg (IIIT Hyderabad), Kishore Kothapalli (IIIT Hyderabad), and Suresh Purini (IIIT Hyderabad)</i>	

A Performance Prediction Framework for Irregular Applications .304.....  
*Gangyi Zhu (Ohio State University) and Gagan Agrawal (Ohio State University)*

Achieving Performance and Programmability for MapReduce(-Like) Frameworks .314.....  
*Jia Guo (Ohio State University) and Gagan Agrawal (Ohio State University)*

Parallel Read Partitioning for Concurrent Assembly of Metagenomic Data .324.....  
*Vasudevan Rengasamy (The Pennsylvania State University), Mahmut T. Kandemir (The Pennsylvania State University), Paul Medvedev (The Pennsylvania State University), and Kamesh Madduri (The Pennsylvania State University)*

**Author Index 335.** .....