

2019 IEEE 26th International Conference on High Performance Computing, Data, and Analytics (HiPC 2019)

**Hyderabad, India
17-20 December 2019**



**IEEE Catalog Number: CFP19176-POD
ISBN: 978-1-7281-4536-5**

**Copyright © 2019 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:	CFP19176-POD
ISBN (Print-On-Demand):	978-1-7281-4536-5
ISBN (Online):	978-1-7281-4535-8
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

2019 IEEE 26th International Conference on High Performance Computing, Data, and Analytics (HiPC) **HiPC 2019**

Table of Contents

Message from General and Vice-General Co-Chairs	x
Message from the Program Chairs	xii
Message from the Steering Chair	xiv
HiPC 2019 Organization	xv
HiPC 2019 Steering Committee	xvi
HiPC 2019 Program Committee	xvii
HiPC 2019 Technical Program	xxi

Keynote 1

Data Flow Execution Models — A Third Opinion	1
<i>Vivek Sarkar (Georgia Institute of Technology)</i>	

Technical Session 1: Algorithms for Graphs and Emerging Platforms

HyDetect: A Hybrid CPU-GPU Algorithm for Community Detection	2
<i>Anwasha Bhowmik (Indian Institute of Science, Bangalore) and Sathish Vadhiyar (Indian Institute of Science, Bangalore)</i>	
Distributed Relational Algebra at Scale	12
<i>Thomas Gilray (University of Alabama, Birmingham) and Sidharth Kumar (University of Alabama, Birmingham)</i>	
Optimizing Breadth-First Search at Scale Using Hardware-Accelerated Space Consistency	23
<i>Khaled Ibrahim (Lawrence Berkeley National Laboratory, USA)</i>	
Shared-Memory Parallel Maximal Biclique Enumeration	34
<i>Apurba Das (National University of Singapore) and Srikanta Tirthapura (Iowa State University)</i>	
A Deterministic Multi-layered Partitioning Tool for Wire-Length Reduction of Monolithic 3D-IC	44
<i>Soumendu Ghorui (Heritage Institute of Technology, Kolkata), Sabyasachee Banerjee (Heritage Institute of Technology, Kolkata), and Subhashis Majumder (Heritage Institute of Technology, Kolkata)</i>	

Mapping Arbitrarily Sparse Two-Body Interactions on One-Dimensional Quantum Circuits .52.....
Arif Khan (Pacific Northwest National Laboratory), Mahantesh Halappanavar (Pacific Northwest National Laboratory), Tobias Hagge (Pacific Northwest National Laboratory), Karol Kowalski (Pacific Northwest National Laboratory), Alex Pothen (Purdue University), and Sriram Krishnamoorthy (Pacific Northwest National Laboratory)

Technical Session 2: Data Management and Visualization

k-NN Sampling for Visualization of Dynamic Data Using LION-tSNE .63.....
Bheekya Dharamsotu (University of Hyderabad), K. Swarupa Rani (University of Hyderabad), Salman Abdul Moiz (University of Hyderabad), and C. Raghavendra Rao (University of Hyderabad)

Analysis in the Data Path of an Object-Centric Data Management System .73.....
Richard Warren (The HDF Group), Jerome Soumagne (The HDF Group), Jingqing Mu (The HDF Group), Houjun Tang (Lawrence Berkeley National Laboratory), Suren Byna (Lawrence Berkeley National Laboratory), Bin Dong (Lawrence Berkeley National Laboratory), and Quincey Koziol (Lawrence Berkeley National Laboratory)

Exploring Metadata Search Essentials for Scientific Data Management .83.....
Wei Zhang (Texas Tech University), Suren Byna (Lawrence Berkeley National Laboratory), Chenxu Niu (Texas Tech University), and Yong Chen (Texas Tech University)

Designing a Profiling and Visualization Tool for Scalable and In-depth Analysis of High-Performance GPU Clusters .93.....
Pouya Kousha (The Ohio State University), Bharath Ramesh (The Ohio State University), Kaushik Kandadi Suresh (The Ohio State University), Ching-Hsiang Chu (The Ohio State University), Arpan Jain (The Ohio State University), Nick Sarkauskas (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhabaleswar K. Panda (The Ohio State University)

Tuning Object-Centric Data Management Systems for Large Scale Scientific Applications .103.....
Houjun Tang (Lawrence Berkeley National Laboratory), Suren Byna (Lawrence Berkeley National Laboratory), Stephen Bailey (Lawrence Berkeley National Laboratory), Zarija Lukic (Lawrence Berkeley National Laboratory), Jialin Liu (Lawrence Berkeley National Laboratory), Quincey Koziol (Lawrence Berkeley National Laboratory), and Bin Dong (Lawrence Berkeley National Laboratory)

Replaceability Based Web Service Selection Approach .113.....
Lalit Purohit (Indian Institute of Technology Roorkee) and Sandeep Kumar (Indian Institute of Technology Roorkee)

Technical Session 3: Applications and Learning

Efficient Parallel Multi-bunch Beam-Beam Simulation in Particle Colliders .123.....
Ioannis Sakiotis (Old Dominion University), Kamesh Arumugam (NVIDIA), Desh Ranjan (Old Dominion University), Balsa Terzic (Old Dominion University), and Mohammad Zubair (Old Dominion University)

Bit-Wise and Multi-GPU Implementations of the DNA Recombination Algorithm .131.....	
	<i>Elnaz Tavakoli Yazdi (University of Arizona), Ankur Limaye (University of Arizona), Ali Akoglu (University of Arizona), Tosiron Adegbiya (University of Arizona), and Adam Buntzman (University of Arizona)</i>
Hierarchical Filter and Refinement System Over Large Polygonal Datasets on CPU-GPU .141.....	
	<i>Yiming Liu (Marquette University, USA), Jie Yang (Marquette University, USA), and Satish Puri (Marquette University, USA)</i>
Geostatistical Modeling and Prediction Using Mixed Precision Tile Cholesky Factorization .152.....	
	<i>Sameh Abdulah (King Abdullah University of Science and Technology), Hatem Ltaief (King Abdullah University of Science and Technology), Ying Sun (King Abdullah University of Science and Technology), Marc G. Genton (King Abdullah University of Science and Technology), and David E. Keyes (King Abdullah University of Science and Technology)</i>
Acceleration of Sparse Vector Autoregressive Modeling Using GPUs .163.....	
	<i>Shreenivas Bharadwaj Venkataramanan (University of California San Diego), Rahul Garg (Indian Institute of Technology Delhi), and Yogish Sabharwal (IBM Research, India)</i>
Fast and Accurate Learning of Knowledge Graph Embeddings at Scale .173.....	
	<i>Udit Gupta (Indian Institute of Science) and Sathish Vadhiyar (Indian Institute of Science)</i>

Keynote 2

Genome Sequencing for Disease Diagnosis: The Confluence of Biology and Computing .183.....	
	<i>Ramesh Hariharan (Strand Life Sciences)</i>

Technical Session 4: Accelerated Learning

On Linear Learning with Manycore Processors .184.....	
	<i>Eliza Wszola (ETH Zurich), Celestine Mender-Dünner (UC Berkeley), Martin Jaggi (EPFL, Switzerland), and Markus Püschel (ETH Zurich)</i>
SPEC2: SPECTral SParse CNN Accelerator on FPGAs .195.....	
	<i>Yue Niu (University of Southern California), Hanqing Zeng (University of Southern California), Ajitesh Srivastava (University of Southern California), Kartik Lakhota (University of Southern California), Rajgopal Kannan (US Army Research Lab-West), Yanzhi Wang (Northeastern University), and Viktor Prasanna (University of Southern California)</i>
Architecture-Centric Bottleneck Analysis for Deep Neural Network Applications .205.....	
	<i>Jihyun Ryoo (Pennsylvania State University), Mengran Fan (Pennsylvania State University), Xulong Tang (Pennsylvania State University), Huaipan Jiang (Pennsylvania State University), Meena Arunachalam (Intel), Sharada Naveen (Intel), and Mahmut T. Kandemir (Pennsylvania State University)</i>
Efficient Sparse Neural Networks Using Regularized Multi Block Sparsity Pattern on a GPU .215.....	
	<i>Dharma Teja Vooturi (International Institute of Information Technology Hyderabad) and Kishore Kothapalli (International Institute of Information Technology Hyderabad)</i>

Memory and Interconnect Optimizations for Peta-Scale Deep Learning Systems .225.....	
<i>Swagath Venkataramani (IBM T.J. Watson Research Center, USA),</i>	
<i>Vijayalakshmi Srinivasan (IBM T. J. Watson Research Center, USA),</i>	
<i>Jungwook Choi (IBM T. J. Watson Research Center, USA), Philip</i>	
<i>Heidelberger (IBM T. J. Watson Research Center, USA), Leland Chang</i>	
<i>(IBM T. J. Watson Research Center, USA), and Kailash Gopalakrishnan</i>	
<i>(IBM T. J. Watson Research Center, USA)</i>	
Accelerating Data Loading in Deep Neural Network Training .235.....	
<i>Chih-Chieh Yang (IBM T. J. Watson Research Center, USA) and Guojing</i>	
<i>Cong (IBM T. J. Watson Research Center, USA)</i>	

Keynote 3

Delivering the Future of High-Performance Computing .246.....	
<i>Mark Papermaster (AMD)</i>	

Technical Session 5: Storage and Communication

IsoKV: An Isolation Scheme for Key-Value Stores by Exploiting Internal Parallelism in SSD .247.....	
<i>Heerak Lim (Seoul National University, Korea), Hwajung Kim (Seoul</i>	
<i>National University, Korea), Kihyeon Myung (Seoul National University,</i>	
<i>Korea), Heon Young (Seoul National University, Korea), and Yongseok</i>	
<i>Son (Chung-Ang University, Korea)</i>	
SCOR-KV: SIMD-Aware Client-Centric and Optimistic RDMA-Based Key-Value Store for Emerging CPU Architectures .257.....	
<i>Dipti Shankar (The Ohio State University), Xiaoyi Lu (The Ohio State</i>	
<i>University), and Dhabaleswar K. Panda (The Ohio State University)</i>	
High-Performance Adaptive MPI Derived Datatype Communication for Modern Multi-GPU Systems .267.....	
<i>Ching-Hsiang Chu (The Ohio State University), Jahanzeb Maqbool Hashmi</i>	
<i>(The Ohio State University), Kawthar Shafie Khorassani (The Ohio State</i>	
<i>University), Hari Subramoni (The Ohio State University), and</i>	
<i>Dhabaleswar K. Panda (The Ohio State University)</i>	
Online Management of Hybrid DRAM-NVMM Memory for HPC .277.....	
<i>Reza Salkhordeh (Johannes Gutenberg University Mainz, Germany) and</i>	
<i>André Brinkmann (Johannes Gutenberg University Mainz, Germany)</i>	
User-Level Scheduled Communications for MPI .290.....	
<i>Derek Schafer (Tennessee Tech University), Sheikh Ghafoor (Tennessee</i>	
<i>Tech University), Daniel Holmes (University of Edinburgh), Martin</i>	
<i>Ruefenacht (University of Tennessee at Chattanooga), and Anthony</i>	
<i>Skjellum (University of Tennessee at Chattanooga)</i>	
Evaluating the Impact of Energy Efficient Networks on HPC Workloads .301.....	
<i>Giorgis Georgakoudis (Lawrence Livermore National Laboratory), Nikhil</i>	
<i>Jain (Lawrence Livermore National Laboratory), Takatsugu Ono (Kyushu</i>	
<i>University), Koji Inoue (Kyushu University), Shinobu Miwa (The</i>	
<i>University of Electro-Communications), and Abhinav Bhatele (Lawrence</i>	
<i>Livermore National Laboratory; University of Maryland)</i>	

Keynote 4

The New World of Heterogeneous AI/ML High Performance Computing with Intel FPGAs Mark	311
<i>José Roberto Alvarez (Intel)</i>	

Technical Session 6: Storage, Fault tolerance, and Resilience

MLBS: Transparent Data Caching in Hierarchical Storage for Out-of-Core HPC Applications	312
<i>Tariq Alturkestani (King Abdullah University of Science and Technology, Saudi Arabia), Thierry Tonellot (Saudi Aramco, Saudi Arabia), Hatem Ltaief (King Abdullah University of Science and Technology, Saudi Arabia), Rached Abdelkhalak (King Abdullah University of Science and Technology, Saudi Arabia), Vincent Etienne (Saudi Aramco, Saudi Arabia), and David Keyes (King Abdullah University of Science and Technology, Saudi Arabia)</i>	
Reducing False Node Failure Predictions in HPC	323
<i>Alvaro Frank (Johannes Gutenberg Universität Mainz), Dai Yang (Technical University of Munich), Andre Brinkmann (Johannes Gutenberg Universität Mainz), Martin Schulz (Technical University of Munich), and Tim Süß (University of Applied Science Fulda)</i>	
Ground-Truth Prediction to Accelerate Soft-Error Impact Analysis for Iterative Methods	333
<i>Burcu O. Mutlu (Pacific Northwest National Laboratory, USA), Gokcen Kestor (Pacific Northwest National Laboratory, USA), Adrian Cristal (Polytechnic University of Catalonia, Spain), Osman Unsal (Barcelona Supercomputing Center, Spain), and Sriram Krishnamoorthy (Pacific Northwest National Laboratory, USA)</i>	
Efficient Memory Pool Allocation Algorithm for CNN Inference	345
<i>Arun Abraham (Samsung R&D Institute India – Bangalore), Manas Sahni (Samsung R&D Institute India – Bangalore), and Akshay Parashar (Samsung R&D Institute India – Bangalore)</i>	
A Linux Kernel Scheduler Extension for Multi-core Systems	353
<i>Aleix Roca (Barcelona Supercomputing Center), Samuel Rodríguez (Barcelona Supercomputing Center), Albert Segura (Barcelona Supercomputing Center), Kevin Marquet (Univ Lyon, INSA Lyon, Inria), and Vicenç Beltran (Barcelona Supercomputing Center)</i>	
uMMAP-IO: User-Level Memory-Mapped I/O for HPC	363
<i>Sergio Rivas-Gomez (KTH Royal Institute of Technology), Alessandro Fanfarillo (National Center for Atmospheric Research), Sebastien Valat (Atos), Christophe Laferriere (Atos), Philippe Couvee (Atos), Sai Narasimhamurthy (Seagate Systems UK), and Stefano Markidis (KTH Royal Institute of Technology)</i>	

Technical Session 7: Parallel and Data Frameworks

DeepSparse: A Task-Parallel Framework for SparseSolvers on Deep Memory Architectures	373
<i>Md Afibuzzaman (Michigan State University), Fazlay Rabbi (Michigan State University), M. Yusuf Özkaya (Georgia Institute of Technology), Hasan Metin Aktulga (Michigan State University), and Umit V. Çatalyürek (Georgia Institute of Technology)</i>	
Worksharing Tasks: An Efficient Way to Exploit Irregular and Fine-Grained Loop Parallelism	383
<i>Marcos Maroñas (Barcelona Supercomputing Center (BSC)), Kevin Sala (Barcelona Supercomputing Center (BSC)), Sergi Mateo (Barcelona Supercomputing Center (BSC)), Eduard Ayguadé (Barcelona Supercomputing Center (BSC)), Universitat Politècnica de Catalunya (UPC), and Vicenç Beltran (Barcelona Supercomputing Center (BSC))</i>	
Empirical Analysis of Hardware-Assisted GPU Virtualization	395
<i>Anshuj Garg (Indian Institute of Technology Bombay), Purushottam Kulkarni (Indian Institute of Technology Bombay), Uday Kurkure (VMware, USA), Hari Sivaraman (VMware, USA), and Lan Vu (VMware, USA)</i>	

Author Index	407
--------------------	-----