

2020 IEEE 27th International Conference on High Performance Computing, Data, and Analytics (HiPC 2020)

**Virtual Conference
16 – 18 December 2020**



**IEEE Catalog Number: CFP20176-POD
ISBN: 978-1-6654-4650-1**

**Copyright © 2020 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:	CFP20176-POD
ISBN (Print-On-Demand):	978-1-6654-4650-1
ISBN (Online):	978-1-6654-2292-5
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

2020 IEEE 27th International Conference on High Performance Computing, Data, and Analytics (HiPC) **HiPC 2020**

Table of Contents

Message from the General Co-chairs	x
Message from the Program Chairs	xii
HiPC 2020 Organization	xiv
HiPC 2020 Technical Program Committee	xvi
Keynote 1: Katherine Yelick	xix
Keynote 2: Animashree Anandkumar	xx
Keynote 3: Fabrizio Petrini	xxi
Industry Sponsors	xxii
HiPC 2020 Technical Program	xxiii

Best Paper Session

SimGQ: Simultaneously Evaluating Iterative Graph Queries	1
<i>Chengshuo Xu (University of California, Riverside, USA), Abbas Mazloumi (University of California, Riverside, USA), Xiaolin Jiang (University of California, Riverside, USA), and Rajiv Gupta (University of California, Riverside, USA)</i>	
WarpCore: A Library for Fast Hash Tables on GPUs	11
<i>Daniel Jünger (Johannes Gutenberg University), Robin Kobus (Johannes Gutenberg University), André Müller (Johannes Gutenberg University), Christian Hundt (NVIDIA), Kai Xu (Shandong University), Weiguo Liu (Shandong University), and Bertil Schmidt (Johannes Gutenberg University)</i>	

Session 1: Applications

Towards High Performance, Portability, and Productivity: Lightweight Augmented Neural Networks for Performance Prediction	21
<i>Ajitesh Srivastava (University of Southern California, USA), Naifeng Zhang (University of Southern California, USA), Rajgopal Kannan (US Army Research Lab-West, USA), and Viktor K. Prasanna (University of Southern California, USA)</i>	

Performance Optimization and Scalability Analysis of the MGB Hydrological Model .31.....
Henrique Rennó de Azeredo Freitas (National Institute for Space Research (INPE)), Celso Luiz Mendes (National Institute for Space Research (INPE)), and Aleksandar Ilic (INESC-ID, Instituto Superior Técnico)

Exploring Task Parallelism for the Multilevel Fast Multipole Algorithm .41.....
Michael Lingg (Michigan State University, Array of Engineers), Stephen Hughey (Michigan State University, ARA), Doga Dikbayir (Michigan State University), Balasubramaniam Shanker (Michigan State University), and Hasan Metin Aktulga (Michigan State University)

SparsePipe: Parallel Deep Learning for 3D Point Clouds .51.....
Keke Zhai (University of Florida), Pan He (University of Florida), Tania Banerjee (University of Florida), Anand Rangarajan (University of Florida), and Sanjay Ranka (University of Florida)

HyPR: Hybrid Page Ranking on Evolving Graphs .62.....
Hemant Kumar Giri (Indian Institute of Information Technology Guwahati), Mridul Haque (Indian Institute of Information Technology Guwahati), and Dip Sankar Banerjee (Indian Institute of Information Technology Guwahati)

Distributing Sparse Matrix/Graph Applications in Heterogeneous Clusters – an Experimental Study .72.....
Charilaos Tzovas (Humboldt University of Berlin), Maria Predari (Humboldt University of Berlin), and Henning Meyerhenke (Humboldt University of Berlin)

Session 2: Scalable Data Science

Processor Pipelining Method for Efficient Deep Neural Network Inference on Embedded Devices .82.....
Akshay Parashar (Samsung R&D Institute Bangalore, India), Arun Abraham (Samsung R&D Institute Bangalore, India), Deepak Chaudhary (Samsung R&D Institute Bangalore, India), and Vikram Nelvooy Rajendiran (Samsung R&D Institute Bangalore, India)

Avoiding Communication in Logistic Regression .91.....
Aditya Devarakonda (Johns Hopkins University, USA) and James Demmel (University of California, Berkeley, USA)

A Parallel and Scalable Framework for Insider Threat Detection .101.....
Abdoulaye Diop (University of Versailles Paris Saclay/Li-Parad/MDLS/Atos-Evidian), Nahid Emad (University of Versailles Paris Saclay/Li-Parad/MDLS/), and Thierry Winter (Atos-Evidian)

Blink: Towards Efficient RDMA-Based Communication Coroutines for Parallel Python Applications .111.....
Aamir Shafi (The Ohio State University, USA), Jahanzeb Maqbool Hashmi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhableswar K. (DK) Panda (The Ohio State University, USA)

Content-Defined Merkle Trees for Efficient Container Delivery .121.....	<i>Yuta Nakamura (DePaul University), Raza Ahmad (DePaul University), and Tanu Malik (DePaul University)</i>
Model Checking as a Service using Dynamic Resource Scaling .131.....	<i>Surya Teja Palavalasa (International Institute of Information Technology, Hyderabad), Yuvaraj Singh (International Institute of Information Technology, Hyderabad), Adhish Singla (International Institute of Information Technology, Hyderabad), Suresh Purini (International Institute of Information Technology, Hyderabad), and Venkatesh Choppella (International Institute of Information Technology, Hyderabad)</i>

Session 3: Algorithms

Parallel Hierarchical Clustering using Rank-Two Nonnegative Matrix Factorization .141.....	<i>Lawton Manning (Wake Forest University), Grey Ballard (Wake Forest University), Ramakrishnan Kannan (Oak Ridge National Laboratory), and Haesun Park (Georgia Institute of Technology)</i>
Pipelined Preconditioned Conjugate Gradient Methods for Distributed Memory Systems .151.....	<i>Manasi Tiwari (Department of Computational and Data Sciences, Indian Institute of Science, Bengaluru) and Sathish Vadhiyar (Department of Computational and Data Sciences, Indian Institute of Science, Bengaluru)</i>
Fair Allocation of Asymmetric Operations in Storage Systems .161.....	<i>Thomas Keller (Rice University) and Peter Varman (Rice University)</i>
A GPU Algorithm for Earliest Arrival Time Problem in Public Transport Networks .171.....	<i>Chirayu Anant Haryan (Indian Institute of Technology Tirupati, India), G. Ramakrishna (Indian Institute of Technology Tirupati, India), Rupesh Nasre (Indian Institute of Technology Madras, India), and Allam Dinesh Reddy (Zippr Private Limited, Hyderabad, India)</i>
2D Static Resource Allocation for Compressed Linear Algebra and Communication Constraints .181	<i>Olivier Beaumont (Inria Bordeaux Sud-Ouest and University of Bordeaux), Lionel Eyraud-Dubois (Inria Bordeaux Sud-Ouest and University of Bordeaux), and Mathieu V�erit�e (Inria Bordeaux Sud-Ouest and University of Bordeaux)</i>
Algorithms for Preemptive Co-Scheduling of Kernels on GPUs .192.....	<i>Lionel Eyraud-Dubois (Inria University of Bordeaux, France) and Cristiana Bentes (State University of Rio de Janeiro, Brazil)</i>

Session 4: Runtime Systems

Understanding HPC Application I/O Behavior Using System Level Statistics .202.....	<i>Arnab K. Paul (Virginia Tech, USA), Olaf Faaland (Lawrence Livermore National Laboratory, USA), Adam Moody (Lawrence Livermore National Laboratory, USA), Elsa Gonsiorowski (Lawrence Livermore National Laboratory, USA), Kathryn Mohror (Lawrence Livermore National Laboratory, USA), and Ali R. Butt (Virginia Tech, USA)</i>
--	--

AMCilk: A Framework for Multiprogrammed Parallel Workloads .212.....	
	<i>Zhe Wang (Washington University in St. Louis), Chen Xu (Washington University in St. Louis), Kunal Agrawal (Washington University in St. Louis), and Jing Li (New Jersey Institute of Technology)</i>
Extending SLURM for Dynamic Resource-Aware Adaptive Batch Scheduling .223.....	
	<i>Mohak Chadha (Technical University of Munich), Jophin John (Technical University of Munich), and Michael Gerndt (Technical University of Munich)</i>
On the Marriage of Asynchronous Many Task Runtimes and Big Data: A Glance .233.....	
	<i>Joshua Suetterlein (Pacific Northwest National Laboratory, USA), Joseph Manzano (Pacific Northwest National Laboratory, USA), Andres Marquez (Pacific Northwest National Laboratory, USA), and Guang Gao (University of Delaware, USA)</i>
Exposing Data Locality in HPC-Based Systems by using the HDFS Backend .243.....	
	<i>Jose Rivadeneira (University Carlos III of Madrid), Felix Garcia-Carballeira (university Carlos III of Madrid), Jesus Carretero (University Carlos III of Madrid), and Javier Garcia-Blas (University Carlos III of Madrid)</i>
PufferFish: NUMA-Aware Work-Stealing Library using Elastic Tasks .251.....	
	<i>Vivek Kumar (IIIT-Delhi)</i>
Design and Study of Elastic Recovery in HPC Applications .261.....	
	<i>Kai Keller (Barcelona Supercomputing Center (BSC-CNS)), Konstantinos Parasyris (Lawrence Livermore National Laboratory (LLNL)), and Leonardo Bautista-Gomez (Barcelona Supercomputing Center (BSC-CNS))</i>

Session 5: System Software and Architecture

Accelerating Force-Directed Graph Layout with Processing-in-Memory Architecture .271.....	
	<i>Ruihao Li (The University of Texas at Austin, USA), Shuang Song (Facebook, USA), Qinzhe Wu (The University of Texas at Austin, USA), and Lizy K. John (The University of Texas at Austin, USA)</i>
Nonblocking Persistent Software Transactional Memory .283.....	
	<i>H. Alan Beadle (University of Rochester), Wentao Cai (University of Rochester), Haosen Wen (University of Rochester), and Michael L. Scott (University of Rochester)</i>
GPU-FPtuner: Mixed-Precision Auto-Tuning for Floating-Point Applications on GPU .294.....	
	<i>Ruidong Gu (North Carolina State University) and Michela Becchi (North Carolina State University)</i>
Batched Small Tensor-Matrix Multiplications on GPUs .305.....	
	<i>Keke Zhai (University of Florida), Tania Banerjee (University of Florida), Adeesha Wijayasiri (University of Moratuwa), and Sanjay Ranka (University of Florida)</i>
Temporal Based Intelligent LRU Cache Construction .315.....	
	<i>Pavan Nittur (Samsung Research Institute Bangalore), Anuradha Kanukotla (Samsung Research Institute Bangalore), and Narendra Mutyala (Samsung Research Institute Bangalore)</i>

Boosting LSTM Performance Through Dynamic Precision Selection 323.....
*Franyell Silfa (Polytechnic University of Catalonia), Jose Maria Arnau
(Polytechnic University of Catalonia), and Antonio González
(Polytechnic University of Catalonia)*

Author Index 335.....