

# **2021 IEEE 28th International Conference on High Performance Computing, Data, and Analytics (HiPC 2021)**

**Virtual Conference  
17-18 December 2021**



**IEEE Catalog Number: CFP21176-POD  
ISBN: 978-1-6654-1017-5**

**Copyright © 2021 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:	CFP21176-POD
ISBN (Print-On-Demand):	978-1-6654-1017-5
ISBN (Online):	978-1-6654-1016-8
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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2021 IEEE 28th International Conference on High Performance Computing, Data, and Analytics (HiPC) **HiPC 2021**

## Table of Contents

Message from the HiPC 2021 General Co-chairs .....	xii
Message from the HiPC 2021 Program Chairs .....	xiv
HiPC 2021 Organization .....	xvi
HiPC 2021 Steering Committee .....	xvii
HiPC 2021 Technical Program Committee .....	xxiii
Keynote 1: Jingren Zhou .....	xxi
Keynote 2: Adam Belay .....	xxii
Keynote 3: Michela Taufer .....	xxiii
HiPC 2021 Technical Program .....	xxiv

### Technical Session 1: Scalable Algorithms and Systems for Data Science - part 1

Parallel Actors and Learners: A Framework for Generating Scalable RL Implementations .....	1
<i>Chi Zhang (University of Southern California, USA), Sanmukh Rao Kuppannagari (University of Southern California, USA), and Viktor K Prasanna (University of Southern California, USA)</i>	
DEISA: Dask-Enabled In Situ Analytics .....	11
<i>Amal Gueroudji (Université Paris-Saclay, UVSQ, CNRS, CEA, France), Julien Bigot (Université Paris-Saclay, UVSQ, CNRS, CEA, France), and Bruno Raffin (Univ. Grenoble Alpes, Inria, CNRS, France)</i>	
A Model of Graph Transactional Coverage Patterns with Applications to Drug Discovery .....	21
<i>A. Srinivas Reddy (IIIT Hyderabad, India), P. Krishna Reddy (IIIT Hyderabad, India), Anirban Mondal (Ashoka University, India), and U. Deva Priyakumar (IIIT Hyderabad, India)</i>	
Faster Parallel Training of Word Embeddings .....	31
<i>Eliza Wszola (ETH Zurich, Switzerland), Martin Jaggi (School of Computer and Communication Sciences, EPFL, Switzerland), and Markus Püschel (ETH Zurich, Switzerland)</i>	

CMAP-LAP: Configurable Massively Parallel Solver for Lattice Problems .....	42
<i>Nariaki Tateiwa (Kyushu University, Japan), Yuji Shinano (Zuse Institute Berlin (ZIB), Germany), Keiichiro Yamamura (Kyushu University, Japan), Akihiro Yoshida (Kyushu University, Japan), Shizuo Kaji (Kyushu University, Japan), Masaya Yasuda (Rikkyo University, Japan), and Katsuki Fujisawa (Kyushu University, Japan)</i>	
MulConn: User-Transparent I/O Subsystem for High-Performance Parallel File Systems .....	53
<i>Hwajung Kim (Seoul National University, Republic of Korea), Jiwoo Bang (Seoul National University, Republic of Korea), Dong Kyu Sung (Seoul National University, Republic of Korea), Hyeonsang Eom (Seoul National University, Republic of Korea), Heon Y. Yeom (Seoul National University, Republic of Korea), and Hanul Sung (Sangmyung University, Republic of Korea)</i>	

## Technical Session 2: HPC Algorithms

Monte Carlo Tree Search for Task Mapping onto Heterogeneous Platforms .....	63
<i>Ta-Yang Wang (University of Southern California, Los Angeles), William Chang (University of Southern California, Los Angeles), Ajitesh Srivastava (University of Southern California, Los Angeles), Rajgopal Kannan (US Army Research Lab, Los Angeles), and Viktor Prasanna (University of Southern California, Los Angeles)</i>	
Shared-Memory Implementation of the Karp-Sipser Kernelization Process .....	71
<i>Johannes Langguth (Simula Research Laboratory and University of Bergen, Norway), Ioannis Panagiotas (Sorbonne Université, CNRS, LIP6, France), and Bora Uçar (CNRS and LIP, Université de Lyon - Inria - ENS Lyon, France)</i>	
How to Avoid Zero-Spacing in Fractionally-Strided Convolution? A Hardware-Algorithm Co-Design Methodology .....	81
<i>Yuan Meng (University of Southern California), Sanmukh Kuppannagari (University of Southern California), Rajgopal Kannan (US Army Research Lab), and Viktor Prasanna (University of Southern California)</i>	
PPBT: A High Performance Parallel Search Tree .....	91
<i>Jiawen Guan (ShanghaiTech University, China) and Rui Fan (ShanghaiTech University, China)</i>	
Deciding Non-Compressible Blocks in Sparse Direct Solvers Using Incomplete Factorization .....	101
<i>Eragul Korkmaz (Univ. Bordeaux Talence, France), Mathieu Fauerge (Univ. Bordeaux Talence, France), Pierre Ramet (Univ. Bordeaux Talence, France), and Grégoire Pichon (Univ Lyon, EnsL, UCBL, CNRS, Inria, LIP, France)</i>	

## Technical Session 3: HPC Applications

Efficient Parallel Algorithms for Computing Percolation Centrality .....	111
<i>Athreya Chandramouli (International Institute of Information Technology, Hyderabad), Sayantan Jana (International Institute of Information Technology, Hyderabad), and Kishore Kothapalli (International Institute of Information Technology, Hyderabad)</i>	

Accelerating JPEG Decompression on GPUs .....	121
<i>André Weissenberger (Johannes Gutenberg University, Germany) and Bertil Schmidt (Johannes Gutenberg University, Germany)</i>	
Towards Zero-Waste Recovery and Zero-Overhead Checkpointing in Ensemble Data Assimilation .....	131
<i>Kai Keller (Barcelona Supercomputing Center (BSC-CNS), Spain), Adrian Cristal Kestelman (Barcelona Supercomputing Center (BSC-CNS), Spain), and Leonardo Bautista Gomez (Barcelona Supercomputing Center (BSC-CNS), Spain)</i>	
Predictive Analysis of Large-Scale Coupled CFD Simulations with the CPX Mini-App .....	141
<i>A. Powell (University of Warwick, UK), K. Choudry (University of Warwick, UK), A. Prabhakar (University of Warwick, UK), I.Z. Reguly (Pazmany Peter Catholic University, Hungary), D. Amirante (University of Surrey, UK), S.A. Jarvis (University of Birmingham, UK), and G.R. Mudalige (University of Warwick, UK)</i>	
The 16,384-Node Parallelism of 3D-CNN Training on an Arm CPU Based Supercomputer .....	152
<i>Akihiro Tabuchi (Fujitsu Limited, Japan), Koichi Shirahata (Fujitsu Limited, Japan), Masafumi Yamazaki (Fujitsu Limited, Japan), Akihiko Kasagi (Fujitsu Limited, Japan), Takumi Honda (Fujitsu Limited, Japan), Kouji Kurihara (Fujitsu Limited, Japan), Kentaro Kawakami (Fujitsu Limited, Japan), Tsuguchika Tabaru (Fujitsu Limited, Japan), Naoto Fukumoto (Fujitsu Limited, Japan), Akiyoshi Kuroda (RIKEN Center for Computational Science, Japan), Takaaki Fukai (RIKEN Center for Computational Science, Japan), and Kento Sato (RIKEN Center for Computational Science, Japan)</i>	

## Technical Session 4: HPC Architecture and System Software

iPUG for Multiple Graphcore IPUs: Optimizing Performance and Scalability of Parallel Breadth-First Search .....	162
<i>Luk Burchard (Simula Research Laboratory, Norway), Xing Cai (Norway University of Oslo, Norway), and Johannes Langguth (Simula Research Laboratory, Norway BI Norwegian Business School, Norway)</i>	
Empirical Analysis of Architectural Primitives for NVRAM Consistency .....	172
<i>Arun Kp (Indian Institute of Technology Kanpur, India), Debadatta Mishra (Indian Institute of Technology Kanpur, India), and Biswabandan Panda (Indian Institute of Technology Bombay, India)</i>	
JACC: An OpenACC Runtime Framework with Kernel-Level and Multi-GPU Parallelization .....	182
<i>Kazuaki Matsumura (Barcelona Supercomputing Center (BSC)), Simon Garcia De Gonzalo (Barcelona Supercomputing Center (BSC)), and Antonio J. Peña (Barcelona Supercomputing Center (BSC))</i>	

## Technical Session 5: HPC Algorithms and Architecture

Anti-Section Transitive Closure .....	192
<i>Oded Green (NVIDIA), Zhihui Du (New Jersey Institute of Technology), Sanyamee Patel (New Jersey Institute of Technology), Zehui Xie (Stevens Institute of Technology), Hang Liu (Stevens Institute of Technology), and David A. Bader (New Jersey Institute of Technology)</i>	

Column-Segmented Sparse Matrix-Matrix Multiplication on Multicore CPUs .....	202
<i>Xiaojing An (Georgia Institute of Technology, USA) and Ümit V. Çatalyürek (Georgia Institute of Technology, USA; Amazon Web Services)</i>	
Multi-stage Memory Efficient Strassen’s Matrix Multiplication on GPU .....	212
<i>Arjun Gopala Krishnan (Concordia University, Canada) and Dhrubajyoti Goswami (Concordia University, Canada)</i>	
Optimizing k-Path Selection for Randomized Interconnection Networks .....	222
<i>Md Nahid Newaz (Oakland University, USA) and Md Atiqul Mollah (Oakland University, USA)</i>	
Dynamic Voltage and Frequency Scaling to Improve Energy-Efficiency of Hardware Accelerators .....	232
<i>Siqin Liu (Ohio University Athens) and Avinash Karanth (Ohio University Athens)</i>	

## Technical Session 6: HPC System Software

Adaptive Placement of Data Analysis Tasks for Staging Based In-Situ Processing .....	242
<i>Zhe Wang (Rutgers University), Pradeep Subedi (University of Utah), Matthieu Dorier (Argonne National Laboratory), Philip E. Davis (Rutgers University), and Manish Parashar (University of Utah)</i>	
HEALS: A Parallel eALS Recommendation System on CPU/GPU Heterogeneous Platforms .....	252
<i>Qihan Wang (William &amp; Mary, USA), Wei Niu (William &amp; Mary, USA), Li Chen (iLambda, Inc), Ruoming Jin (Kent State University, USA), and Bin Ren (William &amp; Mary, USA)</i>	
Shrinking Sample Search Algorithm for Automatic Tuning of GPU Kernels .....	262
<i>Xiang Li (Ohio State University, USA) and Gagan Agrawal (Augusta University, USA)</i>	
Towards Architecture-Aware Hierarchical Communication Trees on Modern HPC Systems .....	272
<i>Bharath Ramesh (The Ohio State University, USA), Jahanzeb Maqbool Hashmi (The Ohio State University, USA), Shulei Xu (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Mahdieh Ghazimirsaeed (The Ohio State University, USA), Mohammadreza Bayatpour (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	

## Technical Session 7: Scalable Algorithms and Systems for Data Science - part 2

DistMILE: A Distributed Multi-level Framework for Scalable Graph Embedding .....	282
<i>Yuntian He (The Ohio State University), Saket Gururkar (The Ohio State University), Pouya Kousha (The Ohio State University), Hari Subramoni (The Ohio State University), Dhabaleswar K. Panda (The Ohio State University), and Srinivasan Parthasarathy (The Ohio State University)</i>	
Model-Based Reinforcement Learning for Elastic Stream Processing in Edge Computing .....	292
<i>Jinlai Xu (University of Pittsburgh, USA) and Balaji Palanisamy (University of Pittsburgh, USA)</i>	

Layout-Aware Hardware-Assisted Designs for Derived Data Types in MPI .....	302
<i>Kaushik Kandadi Suresh (The Ohio State University, USA), Bharath Ramesh (The Ohio State University, USA), Chen Chun Chen (The Ohio State University, USA), Seyedeh Mahdieh Ghazimirsaeed (The Ohio State University, USA), Mohammadreza Bayatpour (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	
Parallel Algorithms for Efficient Computation of High-Order Line Graphs of Hypergraphs .....	312
<i>Xu T. Liu (Washington State University; Pacific Northwest National Lab), Jesun Firoz (Pacific Northwest National Lab), Andrew Lumsdaine (Pacific Northwest National Lab: University of Washington, USA), Cliff Joslyn (Pacific Northwest National Lab), Sinan Aksoy (Pacific Northwest National Lab), Brenda Praggastis (Pacific Northwest National Lab), and Assefaw H. Gebremedhin (Washington State University)</i>	

## Technical Session 8: Scalable Algorithms and Systems for Data Science - part 3

Asynchronous I/O Strategy for Large-Scale Deep Learning Applications .....	322
<i>Sunwoo Lee (Northwestern University), Qiao Kang (Northwestern University), Kewei Wang (Northwestern University), Jan Balewski (National Energy Research Scientific Computing Center), Alex Sim (Lawrence Berkeley National Laboratory), Ankit Agrawal (Northwestern University), Alok Choudhary (Northwestern University), Peter Nugent (Lawrence Berkeley National Laboratory), Kesheng Wu (Lawrence Berkeley National Laboratory), and Wei-keng Liao (Northwestern University)</i>	
SYMBIOMON: A High Performance, Composable Monitoring Service .....	332
<i>Srinivasan Ramesh (University of Oregon), Robert Ross (Argonne National Laboratory), Matthieu Dorier (Argonne National Laboratory), Allen Malony (University of Oregon), Philip Carns (Argonne National Laboratory), and Kevin Huck (University of Oregon)</i>	
Load-Balancing Parallel I/O of Compressed Hierarchical Layouts .....	343
<i>Ke Fan (University of Alabama at Birmingham, USA), Duong Hoang (University of Utah, USA), Steve Petruzza (Utah State University, USA), Thomas Gilray (University of Alabama at Birmingham, USA), Valerio Pascucci (University of Utah, USA), and Sidharth Kumar (University of Alabama at Birmingham, USA)</i>	
CUDA-DClust+: Revisiting Early GPU-Accelerated DBSCAN Clustering Designs .....	354
<i>Madhav Poudel (Northern Arizona University, USA) and Michael Gowanlock (Northern Arizona University, USA)</i>	

## HiPC 2021 Short Papers

Static Graphs for Coding Productivity in OpenACC .....	364
<i>Leonel Toledo (Barcelona Supercomputing Center (BSC) - Fundació i2CAT, Spain), Pedro Valero-Lara (Oak Ridge National Laboratory), Jeffrey Vetter (Oak Ridge National Laboratory), and Antonio J. Peña (Barcelona Supercomputing Center (BSC))</i>	

Performance of Local Push Algorithms for Personalized PageRank on Multi-core Platforms .....	370
<i>Madhav Aggarwal (National Institute of Technology, India), Bingyi Zhang (University of Southern California, USA), and Viktor Prasanna (University of Southern California, USA)</i>	
BEE Orchestrator: Running Complex Scientific Workflows on Multiple Systems .....	376
<i>Jacob Tronge (Kent State University, USA), Patricia Grubel (Los Alamos National Laboratory, USA), Timothy Randles (Los Alamos National Laboratory, USA), Quincy Wofford (Los Alamos National Laboratory, USA), Rusty Davis (Los Alamos National Laboratory, USA), Steven Anaya (Los Alamos National Laboratory, USA), and Qiang Guan (Kent State University, USA)</i>	
OpenACC Multi-GPU Approach for WSM6 Microphysics .....	382
<i>Hércules Cardoso da Silva (Federal University of MS, Brazil), Marco A. Stefanos (Federal University of MS, Brazil), and Vinícius Capistrano (Federal University of MS, Brazil)</i>	
Large-Message Nonblocking MPI_Iallgather and MPI_Ibcast Offload via BlueField-2 DPU .....	388
<i>Nick Sarkauskas (The Ohio State University), Mohammadreza Bayatpour (The Ohio State University), Tu Tran (The Ohio State University), Bharath Ramesh (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhabaleswar K. Panda (The Ohio State University)</i>	
Optimizing Multi-range Based Error-Bounded Lossy Compression for Scientific Datasets .....	394
<i>Yuanjian Liu (University of Chicago, USA), Sheng Di (Argonne National Laboratory, USA), Kai Zhao (University of California, USA), Sian Jin (Washington State University, USA), Cheng Wang (Argonne National Laboratory, USA), Kyle Chard (University of Chicago, USA), Dingwen Tao (Washington State University, USA), Ian Foster (University of Chicago, USA), and Frank Cappello (Argonne National Laboratory, USA)</i>	
An In-Depth I/O Pattern Analysis in HPC Systems .....	400
<i>Jiwoo Bang (Seoul National University, Korea), Chungyong Kim (Seoul National University, Korea), Kesheng Wu (Lawrence Berkeley Nat'l Laboratory, USA), Alex Sim (Lawrence Berkeley Nat'l Laboratory, USA), Suren Byna (Lawrence Berkeley Nat'l Laboratory, USA), Hanul Sung (Sangmyung University, Korea), and Hyeonsang Eom (Seoul National University, Korea)</i>	
FaaSter: Accelerated Functions-as-a-Service with Heterogeneous GPUs .....	406
<i>Anshuj Garg (Indian Institute of Technology Bombay), Purushottam Kulkarni (Indian Institute of Technology Bombay), Umesh Bellur (Indian Institute of Technology Bombay), and Sriram Yenamandra (Georgia Institute of Technology)</i>	
RSP-Hist: Approximate Histograms for Big Data Exploration on Hadoop Clusters .....	412
<i>Salman Salloum (Shenzhen University, China) and Joshua Zhexue Huang (Shenzhen University, China)</i>	
A Programming API Implementation for Secure Data Analytics Applications with Homomorphic Encryption on GPUs .....	418
<i>Shuangsheng Lou (The Ohio State University) and Gagan Agrawal (Augusta University)</i>	



A Fused Inference Design for Pattern-Based Sparse CNN on Edge Devices .....	424
<i>Jia Guo (Ohio State University, USA), Radu Teodorescu (Ohio State University, USA), and Gagan Agrawal (Augusta University, USA)</i>	
Cloud-Based Urgent Computing for Forest Fire Spread Prediction Under Data Uncertainties .....	430
<i>Edigley Fraga (Universitat Autònoma de Barcelona, Spain), Ana Cortés (Universitat Autònoma de Barcelona, Spain), Tomàs Margalef (Universitat Autònoma de Barcelona, Spain), and Porfideo Hernández (Universitat Autònoma de Barcelona, Spain)</i>	
Exploring Thread Coarsening on FPGA .....	436
<i>Mostafa Eghbali Zarch (North Carolina State University, USA), Reece Neff (North Carolina State University, USA), and Michela Becchi (North Carolina State University, USA)</i>	
PILOT: a Runtime System to Manage Multi-tenant GPU Unified Memory Footprint .....	442
<i>John Ravi (NC State University), Tri Nguyen (NC State University), Huiyang Zhou (NC State University), and Michela Becchi (NC State University)</i>	
A Computational Technique for Parallel Solution of Diagonally Dominant Banded Linear Systems .....	448
<i>S. Chandra Sekhara Rao (Indian Institute of Technology Delhi, India) and Rabia Kamra (Indian Institute of Technology Delhi, India)</i>	
<b>Author Index</b> .....	<b>455</b>