

# **2019 IEEE International Conference on Cluster Computing (CLUSTER 2019)**

**Albuquerque, New Mexico, USA  
23 – 26 September 2019**



**IEEE Catalog Number: CFP19235-POD  
ISBN: 978-1-7281-4735-2**

**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:	CFP19235-POD
ISBN (Print-On-Demand):	978-1-7281-4735-2
ISBN (Online):	978-1-7281-4734-5
ISSN:	1552-5244

**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

# 2019 IEEE International Conference on Cluster Computing (CLUSTER)

## Table of Contents

- Message from the General Co-Chairs
- Message from the Program Co-Chairs
- HPCMASPA 2019 Workshop Welcome
- CLUSTER 2019 Committees
- Keynotes

## IEEE Cluster 2019

### Top 3 Papers of Cluster 2019:

- "Evaluating Burst Buffer Placement in HPC Systems" .....1  
*Harsh Khetawat (NCSU); Christopher Zimmer (ORNL); Frank Mueller (NCSU); Scott Atchley and Sudharshan Vazhkudai (ORNL); and Misbah Mubarak (ANL)*
- "Algorithm-Based Fault Tolerance for Parallel Stencil Computations" .....12  
*Aurélien Cavelan and Florina M. Ciorba (University of Basel, Switzerland)*
- "STASH: Fast Hierarchical Aggregation Queries for Effective Visual Spatiotemporal Explorations" .....23  
*Saptashwa Mitra, Paahuni Khandelwal, Shrideep Pallickara, and Sangmi Lee Pallickara (Colorado State University)*

### Deep Learning 1:

- "Efficient User-Level Storage Disaggregation for Deep Learning" .....34  
*Yue Zhu, Weikuan Yu, and Bing Jiao (Florida State University); Kathryn Mohror and Adam Moody (Lawrence Livermore National Laboratory); and Fahim Chowdhury (Florida State University)*
- "FluentPS: A Parameter Server Design with Low-frequency Synchronization for Distributed Deep Learning" .....46  
*Xin Yao, Xueyu Wu, and Cho-Li Wang (HKU)*
- "Performance Characterization of DNN Training using TensorFlow and PyTorch on Modern Clusters" .....58  
*Arpan Jain, Ammar Ahmad Awan, Quentin Anthony, Hari Subramoni, and Dhabaleswar K. Panda (The Ohio State University)*

### Parallel Applications Using Alternate Models:

- "Leveraging Task-Based Polar Decomposition Using PARSEC on Massively Parallel Systems" .....69

*Dalal Sukkari (KAUST); Mathieu Faverge (INRIA); and Hatem Ltaief and David Keyes (KAUST)*

- **"Engineering a Distributed Histogram Sort"** .....81  
*Roger Kowalewski, Pascal Jungblut, and Karl Fuerlinger (LMU Munich)*
- **"Asynchronous Task-Based Execution of the Reverse Time Migration for the Oil and Gas Industry"** .....92  
*Amani Alonazi and Hatem Ltaief (KAUST); Issam Said (NVIDIA); Samuel Thibault (University of Bordeaux, LaBRI INRIA Bordeaux Sud-Ouest); and David Keyes (KAUST)*

## Deep Learning 2:

- **"A Quantitative Study of Deep Learning Training on Heterogeneous Supercomputers"** .....103  
*Jingoo Han (Virginia Tech); Luna Xu (IBM Research); M. Mustafa Rafique (Rochester Institute of Technology); Ali R. Butt (Virginia Tech); and Seung-Hwan Lim (Oak Ridge National Laboratory)*
- **"Parallelizing Training of Deep Generative Models on Massive Scientific Datasets"** .....115  
*Sam Ade Jacobs, Brian Van Essen, Tim Moon, Jae Seung Yeom, David Hysom, Brian Spears, Rushil Anirudh, Jayaraman Thiagarajan, Shusen Liu, Jim Gaffney, Peer-Timo Bremer, Tom Benson, Peter Robinson, and Luc Peterson (Lawrence Livermore National Lab)*
- **"Quantifying the Impact of Memory Errors in Deep Learning"** .....125  
*Zhao Zhang, Lei Huang, Ruizhu Huang, and Weijia Xu (TACC); and Daniel S. Katz (University of Illinois)*

## Workflows:

- **"NORNS: Extending Slurm to Support Data-Driven Workflows through Asynchronous Data Staging"** .....137  
*Alberto Miranda (Barcelona Supercomputing Center); Adrian Jackson (EPCC, The University of Edinburgh); Tommaso Tocci (Barcelona Supercomputing Center); Iakovos Panourgias (EPCC, The University of Edinburgh); and Ramon Nou (Barcelona Supercomputing Center)*
- **"Leveraging Machine Learning for Anticipatory Data Delivery in Extreme Scale In-situ Workflows"** .....149  
*Pradeep Subedi, Philip E. Davis, and Manish Parashar (Rutgers University)*
- **"Harmony: An Approach for Geo-distributed Processing of Big-Data Applications"** .....160  
*Han Zhang (National University of Singapore); Lavanya Ramapantulu (International Institute of Information Technology); and Yong Meng Teo (National University of Singapore)*

## Clustering:

- **"MuDBSCAN: An Exact Scalable DBSCAN Algorithm for Big Data Exploiting Spatial Locality"** .....171  
*Aditya Sarma, Poonam Goyal, Sonal Kumari, Anand Wani, Jagat Sesh Challa, Saiyedul Islam, and Navneet Goyal (Birla Institute of Technology & Science, Pilani)*

### Machine Learning:

- **"HarpGBDT: Optimizing Gradient Boosting Decision Tree for Parallel Efficiency"** .....182  
*Bo Peng, Langshi Chen, Jiayu Li, Miao Jiang, and Selahattin Akkas (Indiana University); Egor Smirnov, Ruslan Israfilov, Sergey Khekhnev, and Andrey Nikolaev (Intel Corporation); and Judy Qiu (Indiana University)*
- **"Training Google Neural Machine Translation on an Intel CPU Cluster"** .....193  
*Dhiraj Kalamkar, Kunal Banerjee, Sudarshan Srinivasan, Srinivas Sridharan, Evangelos Georganas, Mikhail Smorkalov, Cong Xu, and Alexander Heinecke (Intel)*

### Data Centers and Clouds:

- **"MBECN: Enabling ECN with Micro-burst Traffic in Multi-queue Data Center"** .....203  
*Kexi Kang, Jinghui Zhang, Jiahui Jin, Dian Shen, and Junzhou Luo (Southeast University); Wenxin Li (Hong Kong University of Science and Technology); and Zhiang Wu (Nanjing University of Finance and Economics)*
- **"Large-Scale Analysis of the Docker Hub Dataset"** .....215  
*Nannan Zhao (Virginia Tech); Vasily Tarasov (IBM Research—Almaden); Hadeel Albahar (Virginia Tech); Ali Anwar, Lukas Rupperecht, Dimitrios Skourtis, and Amit S. Warke (IBM Research—Almaden); Mohamed Mohamed (Apple); and Ali R. Butt (Virginia Tech)*
- **"DP Greedy: A Two-Phase Caching Algorithm for Mobile Cloud Services"** .....225  
*Dong Huang, Xiaopeng Fan, and Yang Wang (Shenzhen Institutes of Advanced Technology); Shuibing He (Zhejiang University); and Chengzhong Xu (University of Macau)*

### Message Passing:

- **"MPI Sessions: Evaluation of an Implementation in Open MPI"** .....235  
*Nathan Hjelm (Google Inc.); Howard Pritchard and Samuel K. Gutiérrez (Los Alamos National Laboratory); Daniel J. Holmes (EPCC, The University of Edinburgh); Ralph Castain (Intel); and Anthony Skjellum (University of Tennessee at Chattanooga)*
- **"Give MPI Threading a Fair Chance: A Study of Multithreaded MPI Designs"** ..246  
*Thananon Patinyasakdikul, David Eberius, and George Bosilca (University of Tennessee) and Nathan Hjelm (University of New Mexico)*
- **"Fast and Faithful Performance Prediction of MPI Applications: the HPL Case Study"** .....257  
*Tom Cornebize (Univ Grenoble Alpes, French Institute for Research in Computer Science and Automation (INRIA)); Arnaud Legrand (National Center for Scientific*

Research (CNRS), French Institute for Research in Computer Science and Automation (INRIA)); and Franz Christian Heinrich (French Institute for Research in Computer Science and Automation (INRIA))

## Cluster Communication:

- **"X-RDMA: Effective RDMA Middleware in Large-scale Production Environments"** .....268  
Teng Ma (Tsinghua University, Alibaba); Tao Ma, Zhuo Song, Jingxuan Li, and Huaixin Chang (Alibaba); Kang Chen (Tsinghua University); Hai Jiang (Arkansas State University); and Yongwei Wu (Tsinghua University)
- **"Propagation and Decay of Injected One-Off Delays on Clusters: A Case Study"** .....280  
Ayesha Afzal, Georg Hager, and Gerhard Wellein (Friedrich-Alexander University Erlangen-Nürnberg)
- **"An Empirical Study of Cryptographic Libraries for MPI Communications"** .....290  
Abu Naser, Mohsen Gavahi, Cong Wu, Viet Tung Hoang, Zhi Wang, and Xin Yuan (Florida State University)

## Efficient Storage:

- **"RE-Store: Reliable and Efficient KV-Store with Erasure Coding and Replication"** .....301  
Yuzhe Li, Jiang Zhou, and Weiping Wang (Institute of Information Engineering, Chinese Academy of Sciences); and Yong Chen (Texas Tech University)
- **"Compact Filter Structures for Fast Data Partitioning"** .....313  
Qing Zheng, Charles Cranor, Ankush Jain, Gregory Ganger, Garth Gibson, and George Amvrosiadis (Carnegie Mellon University); and Bradley Settlemyer and Gary Grider (Los Alamos National Lab)
- **"Building Reliable High-Performance Storage Systems: An Empirical and Analytical Study"** .....325  
Zhi Qiao (University of North Texas; USRC, LANL); Song Fu (University of North Texas); and Hsing-Bung Chen and Bradley Settlemyer (Los Alamos National Laboratory)

## Compression:

- **"Analyzing the Impact of Lossy Compressor Variability on Checkpointing Scientific Simulations"** .....335  
Pavlo D. Triantafyllides, Tasmia Reza, and Jon C. Calhoun (Clemson University)
- **"Improving Performance of Data Dumping with Lossy Compression for Scientific Simulation"** .....340  
Xin Liang (UC, Riverside); Sheng Di (Argonne National Laboratory); Dingwen Tao (the University of Alabama); Sihuan Li (UC, Riverside); Bogdan Nicolae (Argonne National Laboratory); Zizhong Chen (UC, Riverside); and Franck Cappello (Argonne National Laboratory)

- **"Efficient Distributed Graph Analytics using Triply Compressed Sparse Format"** .....351  
*Mohammad Hasanzadeh Mofrad and Rami Melhem (University of Pittsburgh); and Yousuf Ahmad and Mohammad Hammoud (Carnegie Mellon University in Qatar)*

### Resource Allocation:

- **"SMQoS: Improving Utilization and Power Efficiency with QoS Awareness on GPUs"** .....362  
*Qingxiao Sun, Yi Liu, Hailong Yang, Zhongzhi Luan, and Depei Qian (Beihang University)*
- **"Mitigating Inter-Job Interference via Process-Level Quality-of-Service"** .....367  
*Lee Savoie and David Lowenthal (University of Arizona); Bronis de Supinski and Kathryn Mohror (Lawrence Livermore National Laboratory); and Nikhil Jain (Nvidia)*
- **"Kube-Knots: Resource Harvesting through Dynamic Container Orchestration in GPU-based Datacenters"** .....372  
*Prashanth Thinakaran and Jashwant Raj Gunasekaran (Penn State); Bikash Sharma (Facebook); and Mahmut Kandemir and Chita Das (Penn State)*
- **"Scheduling Independent Stochastic Tasks on Heterogeneous Cloud Platforms"** .....385  
*Yiqin Gao (ENS Lyon); Louis-Claude Canon (Univ. Franche Comté); Yves Robert (ENS Lyon, Univ. Tenn. Knoxville); and Frédéric Vivien (ENS Lyon)*

### Applications:

- **"Multi-physics simulations of particle tracking in arterial geometries with a scalable moving window algorithm"** .....396  
*Gregory J. Herschlag (Duke University); John Gounley (Oak Ridge National Laboratory); Sayan Roychowdhury (Duke University); Erik W. Draeger (Lawrence Livermore National Laboratory); and Amanda Randles (Duke University)*
- **"Fast and Scalable Implementations of Influence Maximization Algorithms"** .....407  
*Marco Minutoli (Pacific Northwest National Laboratory, Washington State University); Mahantesh Halappanavar (Pacific Northwest National Laboratory); Ananth Kalyanaraman (Washington State University); and Arun Sathamur, Ryan McClure, and Jason McDermott (Pacific Northwest National Laboratory)*
- **"Scalable, High-Order Continuity Across Block Boundaries of Functional Approximations Computed in Parallel"** .....419  
*Iulian Grindeanu, Tom Peterka, Vijay Mahadevan, and Youssef S. Nashed (Argonne National Laboratory)*

### Tools and Optimization:

- **"DiffTrace: Efficient Whole-Program Trace Analysis and Diffing for Debugging"** .....428  
*Saeed Taheri and Ian Briggs (University of Utah); Martin Burtscher (Texas State University); and Ganesh Gopalakrishnan (University of Utah)*

- **"FSMonitor: Scalable File System Monitoring for Arbitrary Storage Systems" ...440**  
*Arnab K. Paul (Virginia Tech); Ryan Chard (Argonne National Laboratory); Kyle Chard and Steven Tuecke (University of Chicago); Ali R. Butt (Virginia Tech); and Ian Foster (Argonne National Laboratory, University of Chicago)*
- **"On the Benefits of Anticipating Load Imbalance for Performance Optimization of Parallel Applications" .....451**  
*Anthony Boulmier (University of Geneva); Franck Raynaud (University of Geneva); Nabil Abdennadher (HES-SO); and Bastien Chopard (University of Geneva)*

## Posters:

- **"Accelerating Hyperdimensional Classifier on Multiple GPUs" .....460**  
*Zheming Jin and Hal Finkel (Argonne National Lab)*
- **"Cost-efficiency of Large-scale Electronic Structure Simulations with Intel Xeon Phi Processors" .....462**  
*Hoon Ryu, and Seungmin Lee (Korea Institute of Science and Technology Information)*
- **"Design Exploration of Multi-FPGAs for Accelerating Deep Learning" .....464**  
*Teng Wang, Lei Gong, Chao Wang, Xuehai Zhou, and Huaping Chen (University of Science and Technology of China)*
- **"mSMS: PGAS Runtime with Efficient Thread-based Communication for Global-view Programming" .....466**  
*Hiroko Midorikawa, Kenji Kitagawa, and Yugo Sakaguchi (Seikei University)*
- **"Automatic Power Saving Method by Energy Aware Job Scheduler" .....468**  
*Hiroaki Imade, Takahiro Kagami, Tomohiro Otawa, Kouichi Hirai, Yoshio Sakaguchi (Fujitsu Ltd.); and Naoyuki Fujita (Japan Aerospace Exploration Agency)*
- **"Workflows for Performance Predictable and Reproducible HPC Applications" .....470**  
*Keira Haskins, Quincy Wofford, and Patrick G. Bridges (University of New Mexico)*
- **"Improving Access to HDFS using NVMeoF" .....472**  
*Daegyung Han, and Beomseok Nam (Sungkyunkwan University)*
- **"LogGOPSC: A Parallel Computation Model Extending Network Contention into LogGOPS" .....474**  
*Baicheng Yan, Yi Zhou, Limin Xiao, Jiantong Huo, and Zhaokai Wang (Beihang University)*

## HPCMASPA 2019 Workshop

- **"Improving Resource Utilization in Data Centers using an LSTM-based Prediction Model" .....476**  
*Kundjanasith Thonglek, Kohei Ichikawa, and Keichi Takahashi (Nara Institute of Science and Technology); Chawanat Nakasan (Kanazawa University); and Hajimu Iida (Nara Institute of Science and Technology)*
- **"Proxy or Imposter? A Method and Case Study to Determine the Answer" .....484**  
*Omar Aaziz, Jeanine Cook, and Courtenay Vaughan (SNL); and David Richards (LLNL)*



- **"Standardized Environment for Monitoring Heterogeneous Architectures"** .....493  
*Connor Brown, Benjamin Schwaller, Nathan Gauntt, Benjamin Allan and Kevin Davis (SNL)*
- **"ClusterCockpit — A web application for job-specific performance monitoring"** .....498  
*Jan Eitzinger, Thomas Gruber, Ayesha Afzal, Thomas Zeiser, Gerhard Wellein (Erlangen Regional Computing Center)*
- **"Diagnostic Analysis: Directional Relation Graph"** .....505  
*Sandy Kaur, Eun Kyung Lee (IBM T.J. Watson Research Center)*
- **"Learning from Five-year Resource-Utilization Data of Titan System"** .....510  
*Feiyi Wang, Sarp Oral, Satyabrata Sen, and Neena Imam (ORNL)*
- **"Rapidly Measuring Loop Footprints"** .....516  
*Ozgur O. Kilic, Nathan R. Tallent, and Ryan D. Friese (PNNL)*