

2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS 2019)

**Rio de Janeiro, Brazil
20 – 24 May 2019**

Pages 1-556



**IEEE Catalog Number: CFP19023-POD
ISBN: 978-1-7281-1247-3**

**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:	CFP19023-POD
ISBN (Print-On-Demand):	978-1-7281-1247-3
ISBN (Online):	978-1-7281-1246-6
ISSN:	1530-2075

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 International Parallel and Distributed Processing Symposium (IPDPS) **IPDPS 2019**

Table of Contents

Message from the General Chair .xix.....	
Message from the Program Chair and Vice-chair .xxii.....	
Message from the Steering Chair .xxiii.....	
IPDPS 2019 Technical Program .xxiv.....	
IPDPS 2019 Organization .xxvii.....	
IPDPS 2019 Reviewers .xxxv.....	

Keynote 1

Coding the Continuum .1.....	
<i>Ian Foster (Argonne National Laboratory; University of Chicago)</i>	

Session 1: Graph Algorithms

LACC: A Linear-Algebraic Algorithm for Finding Connected Components in Distributed Memory .2.....	
<i>Ariful Azad (Indiana University) and Aydin Buluc (Lawrence Berkeley National Lab)</i>	
Shared-Memory Exact Minimum Cuts .13.....	
<i>Monika Henzinger (University of Vienna), Alexander Noe (University of Vienna), and Christian Schulz (University of Vienna)</i>	
Distributed Weighted All Pairs Shortest Paths Through Pipelining .23.....	
<i>Udit Agarwal (UT Austin) and Vijaya Ramachandran (UT Austin)</i>	
Local Distributed Algorithms in Highly Dynamic Networks .33.....	
<i>Philipp Bamberger (University of Freiburg), Fabian Kuhn (University of Freiburg), and Yannic Maus (University of Freiburg)</i>	

Session 2: HPC Systems

Effects and Benefits of Node Sharing Strategies in HPC Batch Systems .43.....	
<i>Alvaro Frank (Johannes Gutenberg University Mainz), Tim Süß (Johannes Gutenberg University Mainz), and André Brinkmann (Johannes Gutenberg University Mainz)</i>	

Design Space Exploration of Next-Generation HPC Machines .54.....	
	<i>Constantino Gómez (Barcelona Supercomputing Center), Francesc Martínez (Barcelona Supercomputing Center), Adrià Armejach (Barcelona Supercomputing Center), Miquel Moretó (Barcelona Supercomputing Center), Filippo Mantovani (Barcelona Supercomputing Center), and Marc Casas (Barcelona Supercomputing Center)</i>
A Modular Benchmarking Infrastructure for High-Performance and Reproducible Deep Learning .66.....	
	<i>Tal Ben-Nun (ETH Zurich), Maciej Besta (ETH Zurich), Simon Huber (ETH Zurich), Alexandros Nikolaos Ziogas (ETH Zurich), Daniel Peter (ETH Zurich), and Torsten Hoefler (ETH Zurich)</i>
Double-Precision FPUs in High-Performance Computing: An Embarrassment of Riches? .78.....	
	<i>Jens Domke (Tokyo Institute of Technology), Kazuaki Matsumura (Tokyo Institute of Technology), Mohamed Wahib (AIST-TokyoTech Real World Big-Data Computation Open Innovation Laboratory), Haoyu Zhang (Tokyo Institute of Technology), Keita Yashima (Tokyo Institute of Technology), Toshiki Tsuchikawa (Tokyo Institute of Technology), Yohei Tsuji (Tokyo Institute of Technology), Artur Podobas (Tokyo Institute of Technology), and Satoshi Matsuoka (RIKEN Center for Computational Science (R-CCS), RIKEN, Japan)</i>

Session 3: Numerical Algorithms

Communication-Avoiding Cholesky-QR2 for Rectangular Matrices .89.....	
	<i>Edgar Solomonik (University of Illinois at Urbana-Champaign)</i>
Asynchronous Multigrid Methods .101.....	
	<i>Jordi Wolfson-Pou (Georgia Institute of Technology) and Edmond Chow (Georgia Institute of Technology)</i>
Fast Batched Matrix Multiplication for Small Sizes Using Half-Precision Arithmetic on GPUs .111.....	
	<i>Ahmad Abdelfattah (University of Tennessee), Stanimire Tomov (University of Tennessee), and Jack Dongarra (University of Tennessee, ORNL, University of Manchester)</i>
Load-Balanced Sparse MTTKRP on GPUs .123.....	
	<i>Israt Nisa (The Ohio State University), Jiajia Li (Pacific Northwest National Laboratory), Aravind Sukumaran-Rajam (The Ohio State University), Richard Vuduc (Georgia Institute of Technology), and P. (Saday) Sadayappan (The Ohio State University)</i>

Session 4: Scheduling and Load Balancing I

Practically Efficient Scheduler for Minimizing Average Flow Time of Parallel Jobs .134.....	
	<i>Kunal Agrawal (Washington University in St. Louis), I-Ting Angelina Lee (Washington University in St. Louis), Jing Li (New Jersey Institute of Technology), Kefu Lu (Washington University in St. Louis), and Benjamin Moseley (Carnegie Mellon University)</i>
Scheduling on (Un-)Related Machines with Setup Times .145.....	
	<i>Klaus Jansen (Kiel University), Marten Maack (Kiel University), and Alexander Mucker (Paderborn University)</i>

A Scalable Clustering-Based Task Scheduler for Homogeneous Processors Using DAG Partitioning .155.....	
	<i>M. Yusuf Özkaya (Georgia Institute of Technology), Anne Benoit (ENS Lyon, France & Georgia Institute of Technology), Bora Uçar (CNRS, ENS Lyon, France & Georgia Institute of Technology), Julien Herrmann (Inria & Georgia Institute of Technology), and Ümit V. Çatalyürek (Georgia Institute of Technology)</i>
Reservation Strategies for Stochastic Jobs .166.....	
	<i>Guillaume Aupy (Inria & Labri, Univ. of Bordeaux), Ana Gainaru (Vanderbilt University), Valentin Honoré (Inria & Labri, Univ. of Bordeaux), Padma Raghavan (Vanderbilt University), Yves Robert (Laboratoire LIP, ENS Lyon, France & University of Tennessee), and Hongyang Sun (Vanderbilt University)</i>

Session 5: Accelerating Neural Networks

Exploiting Flow Graph of System of ODEs to Accelerate the Simulation of Biologically-Detailed Neural Networks .176.....	
	<i>Bruno Ricardo da Cunha Magalhaes (École polytechnique fédérale de Lausanne), Thomas Sterling (Indiana University), Felix Schuermann (École polytechnique fédérale de Lausanne), and Michael Hines (Department of Neuroscience, Yale University)</i>
Runtime Concurrency Control and Operation Scheduling for High Performance Neural Network Training .188	
	<i>Jiawen Liu (University of California, Merced), Dong Li (University of California, Merced), Gokcen Kestor (Pacific Northwest National Laboratory), and Jeffrey Vetter (Oak Ridge National Laboratory)</i>
Dynamic Memory Management for GPU-Based Training of Deep Neural Networks .200.....	
	<i>Shriram S B (Indian Institute of Technology Bombay), Anshuj Garg (Indian Institute of Technology Bombay), and Purushottam Kulkarni (Indian Institute of Technology Bombay)</i>
Improving Strong-Scaling of CNN Training by Exploiting Finer-Grained Parallelism .210.....	
	<i>Nikoli Dryden (University of Illinois at Urbana-Champaign), Naoya Maruyama (Lawrence Livermore National Laboratory), Tom Benson (Lawrence Livermore National Laboratory), Tim Moon (Lawrence Livermore National Laboratory), Marc Snir (University of Illinois at Urbana-Champaign), and Brian Van Essen (Lawrence Livermore National Laboratory)</i>

Session 6: GPU Computing I

Excavating the Potential of GPU for Accelerating Graph Traversal .221.....	
	<i>Pengyu Wang (Shanghai Jiao Tong University), Lu Zhang (Shanghai Jiao Tong University), Chao Li (Shanghai Jiao Tong University), and Minyi Guo (Shanghai Jiao Tong University)</i>

ParILUT - A Parallel Threshold ILU for GPUs .231.....
Hartwig Anzt (Karlsruhe Institute of Technology / University of Tennessee), Tobias Ribizel (Karlsruhe Institute of Technology), Goran Flegar (University of Jaume I), Edmond Chow (Georgia Institute of Technology), and Jack Dongarra (University of Tennessee / Oak Ridge National Lab / University of Manchester)

C-GDR: High-Performance Container-Aware GPU Direct MPI Communication Schemes on RDMA Networks 242
Jie Zhang (The Ohio State University), Xiaoyi Lu (The Ohio State University), Ching-Hsiang Chu (The Ohio State University), and Dhableswar K. (DK) Panda (The Ohio State University)

Slate: Enabling Workload-Aware Efficient Multiprocessing for Modern GPGPUs .252.....
Tyler Allen (Clemson University), Xizhou Feng (Clemson University), and Rong Ge (Clemson University)

Session 7: Learning and Prediction Systems

A Deep Recurrent Neural Network Based Predictive Control Framework for Reliable Distributed Stream Data Processing .262.....
Jielong Xu (Syracuse University), Jian Tang (Syracuse University), Zhiyuan Xu (Syracuse University), Chengxiang Yin (Syracuse University), Kevin Kwiat (US Air Force Research Lab), and Charles Kamhoua (US Army Research Lab)

Architecting Racetrack Memory Preshift through Pattern-Based Prediction Mechanisms .273.....
Adrian Colaso (University of Cantabria), Pablo Prieto (University of Cantabria), Pablo Abad (University of Cantabria), Jose Angel Gregorio (University of Cantabria), and Valentin Puente (University of Cantabria)

DLHub: Model and Data Serving for Science .283.....
Ryan Chard (Argonne National Laboratory), Zhuozhao Li (University of Chicago), Kyle Chard (University of Chicago), Logan Ward (University of Chicago), Yadu Babuji (University of Chicago), Anna Woodard (University of Chicago), Steven Tuecke (University of Chicago), Ben Blaiszik (University of Chicago), Michael Franklin (University of Chicago), and Ian Foster (Argonne National Laboratory)

Identifying Latent Reduced Models to Precondition Lossy Compression .293.....
Huizhang Luo (New Jersey Institute of Technology), Dan Huang (New Jersey Institute of Technology), Qing Liu (New Jersey Institute of Technology), Zhenbo Qiao (New Jersey Institute of Technology), Hong Jiang (University of Texas at Arlington), Jing Bi (Beijing University of Technology, China), Haitao Yuan (Beijing Jiaotong University, China), Mengchu Zhou (New Jersey Institute of Technology), Jinzhen Wang (New Jersey Institute of Technology), and Zhenlu Qin (New Jersey Institute of Technology)

Session 8: Multicore Computing

- QoS-Driven Coordinated Management of Resources to Save Energy in Multi-core Systems .303.....
Mehrzad Nejat (Chalmers University of Technology), Miquel Pericas (Chalmers University of Technology), and Per Stenstrom (Chalmers University of Technology)
- Efficient Architecture-Aware Acceleration of BWA-MEM for Multicore Systems .314.....
Md Vasimuddin (Parallel Computing Lab, Intel Corporation, Bangalore, India), Sanchit Misra (Parallel Computing Lab, Intel Corporation, Bangalore, India), Heng Li (Dana-Farber Cancer Institute, Boston, USA; Harvard Medical School), and Srinivas Aluru (Georgia Institute of Technology)
- Power and Performance Tradeoffs for Visualization Algorithms .325.....
Stephanie Labasan (Lawrence Livermore National Laboratory, University of Oregon), Matthew Larsen (Lawrence Livermore National Laboratory), Hank Childs (University of Oregon), and Barry Rountree (Lawrence Livermore National Laboratory)
- Northup: Divide-and-Conquer Programming in Systems with Heterogeneous Memories and Processors .335....
Shuai Che (Microsoft) and Jieming Yin (Advanced Micro Devices, Inc.)

Plenary Session: Best Papers

- Distributed Approximate k-Core Decomposition and Min-Max Edge Orientation: Breaking the Diameter Barrier .345.....
T-H. Hubert Chan (The University of Hong Kong), Mauro Sozio (Télécom ParisTech University), and Bintao Sun (The University of Hong Kong)
- FALCON: Efficient Designs for Zero-Copy MPI Datatype Processing on Emerging Architectures .355.....
Jahanzeb Maqbool Hashmi (The Ohio State University), Sourav Chakraborty (The Ohio State University), Mohammadreza Bayatpour (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhabaleswar K. (DK) Panda (The Ohio State University)
- Two Elementary Instructions Make Compare-and-Swap .365.....
Pankaj Khanchandani (ETH Zurich) and Roger Wattenhofer (ETH Zurich)
- Robust Dynamic Resource Allocation via Probabilistic Task Pruning in Heterogeneous Computing Systems.375
James Gentry (University of Louisiana at Lafayette), Chavit Demninnart (University of Louisiana at Lafayette), and Mohsen Amini Salehi (University of Louisiana at Lafayette)

Keynote 2

- Two Roads to Parallelism: From Serial Code to Programming with STAPL .385.....
Lawrence Rauchwerger (Texas A&M University)

Session 9: Cloud Computing

- Z-Dedup: A Case for Deduplicating Compressed Contents in Cloud .386.....
Zhichao Yan (University of Texas at Arlington), Hong Jiang (University of Texas at Arlington), Yujuan Tan (Chongqing University), Stan Skelton (NetApp), and Hao Luo (Twitter)
- An Architecture and Stochastic Method for Database Container Placement in the Edge-Fog-Cloud Continuum .396.....
Petar Kochovski (University of Ljubljana), Rizos Sakellariou (University of Manchester), Marko Bajec (University of Ljubljana), Pavel Drobintsev (Peter the Great St.Petersburg Polytechnic University), and Vlado Stankovski (University of Ljubljana)
- Online Live VM Migration Algorithms to Minimize Total Migration Time and Downtime .406.....
Nikos Tziritas (Shenzhen Institutes of Advanced Technology), Thanasis Loukopoulos (University of Thessaly), Samee Khan (North Dakota State University), Cheng-Zhong Xu (University of Macau), and Albert Zomaya (University of Sydney)
- Semantics-Aware Virtual Machine Image Management in IaaS Clouds .418.....
Nishant Saurabh (University of Innsbruck, University of Klagenfurt), Julian Remmers (University of Innsbruck), Dragi Kimovski (University of Klagenfurt), Radu Prodan (University of Klagenfurt, University of Innsbruck), and Jorge G. Barbosa (LIACC, Faculdade de Engenharia da Universidade do Porto)

Session 10: Graph Algorithms II

- Composing Optimization Techniques for Vertex-Centric Graph Processing via Communication Channels .428.
Yongzhe Zhang (SOKENDAI (The Graduate University for Advanced Studies), National Institute of Informatics) and Zhenjiang Hu (SOKENDAI (The Graduate University for Advanced Studies), National Institute of Informatics, University of Tokyo)
- CuSP: A Customizable Streaming Edge Partitioner for Distributed Graph Analytics .439.....
Loc Hoang (The University of Texas at Austin), Roshan Dathathri (The University of Texas at Austin), Gurbinder Gill (The University of Texas at Austin), and Keshav Pingali (The University of Texas at Austin)
- Accelerating Sequence Alignment to Graphs .451.....
Chirag Jain (Georgia Institute of Technology), Sanchit Misra (Intel Corporation), Haowen Zhang (Georgia Institute of Technology), Alexander Dilthey (University Hospital of Dusseldorf), and Srinivas Aluru (Georgia Institute of Technology)
- Accurate, Efficient and Scalable Graph Embedding .462.....
Hanqing Zeng (University of Southern California), Hongkuan Zhou (University of Southern California), Ajitesh Srivastava (University of Southern California), Rajgopal Kannan (US Army Research Lab-West), and Viktor Prasanna (University of Southern California)

Session 11: Linear Algebra

- Matrix Powers Kernels for Thick-Restart Lanczos with Explicit External Deflation .472.....
Ichitaro Yamazaki (University of Tennessee), Zhaojun Bai (University of California, Davis), Ding Lu (University of California, Davis), and Jack Dongarra (University of Tennessee, Knoxville)
- Revisiting the I/O-Complexity of Fast Matrix Multiplication with Recomputations .482.....
Roy Nissim (The Hebrew University of Jerusalem) and Oded Schwartz (The Hebrew University of Jerusalem)
- Computation of Matrix Chain Products on Parallel Machines .491.....
Elad Weiss (The Hebrew University of Jerusalem, Israel) and Oded Schwartz (The Hebrew University of Jerusalem, Israel)
- Overlapping Communications with Other Communications and Its Application to Distributed Dense Matrix Computations .501.....
Hua Huang (Georgia Institute of Technology) and Edmond Chow (Georgia Institute of Technology)

Session 12: Storage Systems

- Data Jockey: Automatic Data Management for HPC Multi-tiered Storage Systems .511.....
Woong Shin (Oak Ridge National Laboratory), Christopher D. Brumgard (Oak Ridge National Laboratory), Bing Xie (Oak Ridge National Laboratory), Sudharshan S. Vazhkudai (Oak Ridge National Laboratory), Devarshi Ghoshal (Lawrence Berkeley National Laboratory), Sarp Oral (Oak Ridge National Laboratory), and Lavanya Ramakrishnan (Lawrence Berkeley National Laboratory)
- NCQ-Aware I/O Scheduling for Conventional Solid State Drives .523.....
Hao Fan (Huazhong University of Science and Technology), Song Wu (Huazhong University of Science and Technology), Shadi Ibrahim (Inria, IMT Atlantique, LS2N), Ximing Chen (Huazhong University of Science and Technology), Hai Jin (Huazhong University of Science and Technology), Jiang Xiao (Huazhong University of Science and Technology), and Haibing Guan (Shanghai Key Laboratory of Scalable Computing and Systems, Shanghai Jiao Tong University)
- Optimizing the Parity Check Matrix for Efficient Decoding of RS-Based Cloud Storage Systems .533.....
Junqing Gu (Shanghai Jiao Tong University), Chentao Wu (Shanghai Jiao Tong University), Xin Xie (Shanghai Jiao Tong University), Han Qiu (Shanghai Jiao Tong University), Jie Li (Shanghai Jiao Tong University), Minyi Guo (Shanghai Jiao Tong University), Xubin He (Temple University), Yuanyuan Dong (Alibaba Group), and Yafei Zhao (Alibaba Group)
- D3: Deterministic Data Distribution for Efficient Data Reconstruction in Erasure-Coded Distributed Storage Systems .545.....
Zhipeng Li (University of Science and Technology of China), Min Lv (University of Science and Technology of China), Yinlong Xu (University of Science and Technology of China), Yongkun Li (University of Science and Technology of China), and Liangliang Xu (University of Science and Technology of China)

Session 13: Applications I

SunwayLB: Enabling Extreme-Scale Lattice Boltzmann Method Based Computing Fluid Dynamics Simulations on Sunway TaihuLight .557.....

Zhao Liu (Tsinghua University), XueSen Chu (China Ship Scientific Research Center), Xiaojing Lv (National Supercomputer Center in Wuxi), Hongsong Meng (National Supercomputer Center in Wuxi), Shupeng Shi (National Supercomputer Center in Wuxi), Wenji Han (China Ship Scientific Research Center), Jingheng Xu (Tsinghua University), Haohuan Fu (Tsinghua University), and Guangwen Yang (Tsinghua University)

Containers in HPC: A Scalability and Portability Study in Production Biological Simulations .567.....

Oleksandr Rudyy (Barcelona Supercomputing Center), Marta Garcia-Gasulla (Barcelona Supercomputing Center), Filippo Mantovani (Barcelona Supercomputing Center), Alfonso Santiago (Barcelona Supercomputing Center), Raül Sirvent (Barcelona Supercomputing Center), and Mariano Vázquez (Barcelona Supercomputing Center, ELEM Biotech)

PaKman: Scalable Assembly of Large Genomes on Distributed Memory Machines .578.....

Priyanka Ghosh (Washington State University), Sriram Krishnamoorthy (Pacific Northwest National Laboratory), and Ananth Kalyanaraman (Washington State University)

Language Modeling at Scale .590.....

Md Mostofa Ali Patwary (Baidu), Milind Chabbi (Baidu), Heewoo Jun (Baidu), Jiaji Huang (Baidu), Greg Diamos (Baidu), and Kenneth Church (Baidu)

Session 14: File Systems

DYRS: Bandwidth-Aware Disk-to-Memory Migration of Cold Data in Big-Data File Systems .600.....

Simbarashe Dzinamarira (Rice University), Florin Dinu (University of Sydney), and T. S. Eugene Ng (Rice University)

iez: Resource Contention Aware Load Balancing for Large-Scale Parallel File Systems .610.....

Bharti Wadhwa (Virginia Tech), Arnab K. Paul (Virginia Tech), Sarah Neuwirth (University of Heidelberg), Feiyi Wang (Oak Ridge National Laboratory), Sarp Oral (Oak Ridge National Laboratory), Ali R. Butt (Virginia Tech), Jon Bernard (Virginia Tech), and Kirk W. Cameron (Virginia Tech)

SimFS: A Simulation Data Virtualizing File System Interface .621.....

Salvatore Di Girolamo (ETH Zurich), Pirmin Schmid (ETH Zurich), Thomas C. Schulthess (ETH Zurich), and Torsten Hoefler (ETH Zurich)

Sizing and Partitioning Strategies for Burst-Buffers to Reduce IO Contention .631.....

Guillaume Aupy (Inria and Université de Bordeaux), Olivier Beaumont (Inria and Université de Bordeaux), and Lionel Eyraud-Dubois (Inria and Université de Bordeaux)

Session 15: GPU Computing II

- On Optimizing Complex Stencils on GPUs .641.....
Prashant Rawat (The Ohio State University), Miheer Vaidya (The Ohio State University), Aravind Sukumaran-Rajam (The Ohio State University), Atanas Rountev (The Ohio State University), Louis-Noel Pouchet (Colorado State University), and P. Sadayappan (The Ohio State University)
- Themis: Predicting and Reining in Application-Level Slowdown on Spatial Multitasking GPUs .653.....
Wenyi Zhao (Shanghai Jiao Tong University), Quan Chen (Shanghai Jiao Tong University), Hao Lin (Alibaba Group), Jianfeng Zhang (Alibaba Group), Jingwen Leng (Shanghai Jiao Tong University), Chao Li (Shanghai Jiao Tong University), Wenli Zheng (Shanghai Jiao Tong University), Li Li (Shanghai Jiao Tong University), and Minyi Guo (Shanghai Jiao Tong University)
- Exploiting Adaptive Data Compression to Improve Performance and Energy-Efficiency of Compute Workloads in Multi-GPU Systems .664.....
Mohammad Khavari Tavana (Northeastern University), Yifan Sun (Northeastern University), Nicolas Bohm Agostini (Northeastern University), and David Kaeli (Northeastern University)
- Dual Pattern Compression Using Data-Preprocessing for Large-Scale GPU Architectures .675.....
Kyung Hoon Kim (Texas A&M University), Priyank Devpura (Intel Corporation), Abhishek Nayyar (Texas A&M University), Andrew Doolittle (Texas A&M University), Kihwan Yum (Texas A&M University), and Eun Jung Kim (Texas A&M University)

Session 16: Scheduling and Load Balancing II

- Adapting Batch Scheduling to Workload Characteristics: What Can We Expect From Online Learning? .686...
Arnaud Legrand (Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, Grenoble, France), Denis Trustram (Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, Grenoble, France), and Salah Zrigui (Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, Grenoble, France)
- Aladdin: Optimized Maximum Flow Management for Shared Production Clusters .696.....
Heng Wu (Institute of Software, Chinese Academy of Sciences), Wenbo Zhang (Institute of Software, Chinese Academy of Sciences), Yuanjia Xu (Institute of Software, Chinese Academy of Sciences), Hao Xiang (Institute of Software, Chinese Academy of Sciences), Tao Huang (Institute of Software, Chinese Academy of Sciences), Haiyang Ding (Alibaba Group), and Zheng Zhang (Alibaba Group)
- mmWave Wireless Backhaul Scheduling of Stochastic Packet Arrivals .708.....
Pawel Garncarek (University of Wroclaw), Tomasz Jurdziski (University of Wroclaw), Dariusz Kowalski (University of Liverpool; SWPS University of Social Sciences and Humanities, Warsaw, Poland), and Miguel Mosteiro (Pace University)
- Tight & Simple Load Balancing .718.....
Petra Berenbrink (Universität Hamburg), Tom Friedetzky (Durham University), Dominik Kaaser (Universität Hamburg), and Peter Kling (Universität Hamburg)

Keynote 3

The Path to Delivering Programable Exascale Systems .727.....
Luiz DeRose (Cray Inc.)

Session 17: Managing Data

An Error-Reflective Consistency Model for Distributed Data Stores .728.....
Philip Dexter (SUNY Binghamton), Kenneth Chiu (SUNY Binghamton), and Bedri Sendir (IBM Research)

A High-Performance Distributed Relational Database System for Scalable OLAP Processing .738.....
Jason Arnold (IIT), Boris Glavic (IIT), and Ioan Raicu (IIT)

An Approach for Parallel Loading and Pre-Processing of Unstructured Meshes Stored in Spatially Scattered Fashion .749.....
Ondej Meca (IT4Innovations, VŠB - Technical University of Ostrava, Ostrava, Czech Republic), Lubomír ťha (IT4Innovations, VŠB - Technical University of Ostrava, Ostrava, Czech Republic), and Tomáš Brzobohatý (IT4Innovations, VŠB - Technical University of Ostrava, Ostrava, Czech Republic)

Session 18: Message Passing

Exploring MPI Communication Models for Graph Applications Using Graph Matching as a Case Study .761...
Sayan Ghosh (Washington State University), Mahantesh Halappanavar (Pacific Northwest National Laboratory), Ananth Kalyanaraman (Washington State University), Arif Khan (Pacific Northwest National Laboratory), and Assefaw Gebremedhin (Washington State University)

BigSpa: An Efficient Interprocedural Static Analysis Engine in the Cloud .771.....
Zhiqiang Zuo (Nanjing University), Rong Gu (Nanjing University), Xi Jiang (Nanjing University), Zhaokang Wang (Nanjing University), Yihua Huang (Nanjing University), Linzhang Wang (Nanjing University), and Xuandong Li (Nanjing University)

An Efficient Collaborative Communication Mechanism for MPI Neighborhood Collectives .781.....
S. Mahdieh Ghazimirsaeed (Queen's University), Seyed H. Mirsadeghi (Queen's University), and Ahmad Afsahi (Queen's University)

Session 19: Managing Power and Energy

Understanding the Impact of Dynamic Power Capping on Application Progress .793.....
Srinivasan Ramesh (University of Oregon), Swann Perarnau (Argonne National Laboratory), Sridutt Bhalachandra (Argonne National Laboratory), Allen D. Malony (University of Oregon), and Pete Beckman (Argonne National Laboratory)

Modelling DVFS and UFS for Region-Based Energy Aware Tuning of HPC Applications .805.....	
	<i>Mohak Chadha (Technische Universität München) and Michael Gerndt (Technische Universität München)</i>
SprintCon: Controllable and Efficient Computational Sprinting for Data Center Servers .815.....	
	<i>Wenli Zheng (Shanghai Jiao Tong University), Xiaorui Wang (The Ohio State University), Yue Ma (Shanghai Jiao Tong University), Chao Li (Shanghai Jiao Tong University), Hao Lin (Alibaba Group), Bin Yao (Shanghai Jiao Tong University), Jianfeng Zhang (Alibaba Group), and Minyi Guo (Shanghai Jiao Tong University)</i>
Drowsy-DC: Data Center Power Management System .825.....	
	<i>Mathieu Bacou (IRIT, Université de Toulouse), Grégoire Todeschi (IRIT, Université de Toulouse), Alain Tchana (IRIT, Université de Toulouse), Daniel Hagimont (IRIT, Université de Toulouse), Baptiste Lepers (École Polytechnique Fédérale de Lausanne), and Willy Zwaenepoel (École Polytechnique Fédérale de Lausanne)</i>

Session 20: Networks

Distributed Dominating Set and Connected Dominating Set Construction Under the Dynamic SINR Model .835	
	<i>Dongxiao Yu (Shandong University), Yifei Zou (The University of Hong Kong), Yong Zhang (Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences), Feng Li (Shandong University), Jiguo Yu (Qilu University of Technology (Shandong Academy of Science), Shandong Computer Center (National Supercomputer Center in Jinan); School of Information Science and Engineering, Qufu Normal University), Yu Wu (Dongguan University of Technology), Xiuzhen Cheng (Shandong University), and Francis C.M. Lau (The University of Hong Kong)</i>
MULTISKIPGRAPH: A Self-Stabilizing Overlay Network that Maintains Monotonic Searchability .845.....	
	<i>Linghui Luo (Paderborn University), Christian Scheideler (Paderborn University), and Thim Strothmann (Paderborn University)</i>
Network Size Estimation in Small-World Networks Under Byzantine Faults .855.....	
	<i>Soumyottam Chatterjee (University of Houston), Gopal Pandurangan (University of Houston), and Peter Robinson (McMaster University)</i>
MD-GAN: Multi-Discriminator Generative Adversarial Networks for Distributed Datasets .866.....	
	<i>Corentin Hardy (Technicolor and INRIA Rennes), Erwan Le Merrer (INRIA Rennes), and Bruno Sericola (INRIA Rennes)</i>

Session 21: Dealing with Faults

MOARD: Modeling Application Resilience to Transient Faults on Data Objects .878.....	
	<i>Luanzheng Guo (University of California, Merced) and Dong Li (EECS, University of California, Merced)</i>
SAFIRE: Scalable and Accurate Fault Injection for Parallel Multithreaded Applications .890.....	
	<i>Giorgis Georgakoudis (Lawrence Livermore National Laboratory, USA; Queen's University Belfast), Ignacio Laguna (Lawrence Livermore National Laboratory, USA), Hans Vandierendonck (Queen's University Belfast), Dimitrios S. Nikolopoulos (Queen's University Belfast), and Martin Schulz (Technische Universität München)</i>

- Optimal Placement of In-memory Checkpoints Under Heterogeneous Failure Likelihoods .900.....
Zaeem Hussain (University of Pittsburgh), Taieb Znati (University of Pittsburgh), and Rami Melhem (University of Pittsburgh)
- VeloC: Towards High Performance Adaptive Asynchronous Checkpointing at Large Scale .911.....
Bogdan Nicolae (Argonne National Laboratory), Adam Moody (Lawrence Livermore National Laboratory), Elsa Gonsiorowski (Lawrence Livermore National Laboratory), Kathryn Mohror (Lawrence Livermore National Laboratory), and Franck Cappello (Argonne National Laboratory)

Session 22: Optimizing Memory Behavior

- HART: A Concurrent Hash-Assisted Radix Tree for DRAM-PM Hybrid Memory Systems .921.....
Wen Pan (San Diego State University), Tao Xie (San Diego State University), and Xiaojia Song (San Diego State University)
- LLC-Guided Data Migration in Hybrid Memory Systems .932.....
Evangelos Vasilakis (Chalmers University of Technology), Vassilis Papaefstathiou (Foundation for Research and Technology – Hellas (FORTH)), Pedro Trancoso (Chalmers University of Technology), and Ioannis Sourdis (Chalmers University of Technology)
- Software-Based Buffering of Associative Operations on Random Memory Addresses .943.....
Matthias Hauck (Computer Engineering Group, Ruprecht-Karls University of Heidelberg / SAP SE), Marcus Paradies (DLR), and Holger Fröning (Computer Engineering Group, Ruprecht-Karls University of Heidelberg)
- Combining Prefetch Control and Cache Partitioning to Improve Multicore Performance .953.....
Gongjin Sun (University of California, Irvine), Junjie Shen (University of California, Irvine), and Alexander V. Veidenbaum (University of California, Irvine)

Session 23: Programming Languages

- UPC++: A High-Performance Communication Framework for Asynchronous Computation .963.....
John Bachan (Lawrence Berkeley National Laboratory), Scott B. Baden (Lawrence Berkeley National Laboratory), Steven Hofmeyr (Lawrence Berkeley National Laboratory), Mathias Jacquelin (Lawrence Berkeley National Laboratory), Amir Kamil (Lawrence Berkeley National Laboratory and University of Michigan), Dan Bonachea (Lawrence Berkeley National Laboratory), Paul H. Hargrove (Lawrence Berkeley National Laboratory), and Hadia Ahmed (Lawrence Berkeley National Laboratory)
- Cpp-Taskflow: Fast Task-Based Parallel Programming Using Modern C++ .974.....
Tsung-Wei Huang (University of Illinois at Urbana-Champaign), Chun-Xun Lin (University of Illinois at Urbana-Champaign), Guannan Guo (University of Illinois at Urbana-Champaign), and Martin Wong (University of Illinois at Urbana-Champaign)
- Portal: A High-Performance Language and Compiler for Parallel N-Body Problems .984.....
Laleh Aghababaie Beni (University of California, Irvine), Saikiran Ramanan (University of California, Irvine), and Aparna Chandramowlishwaran (University of California, Irvine)

SAC Goes Cluster: Fully Implicit Distributed Computing .996.....
Thomas Macht (University of Amsterdam) and Clemens Greck (University of Amsterdam)

Session 24: Accelerating Graph Processing

Incremental Graph Processing for On-line Analytics .1007.....
Scott Sallinen (University of British Columbia), Roger Pearce (Lawrence Livermore National Laboratory), and Matei Ripeanu (University of British Columbia)

Incrementalization of Vertex-Centric Programs .1019.....
Timothy A. K. Zakian (University of Oxford), Ludovic A. R. Capelli (The University of Edinburgh), and Zhenjiang Hu (National Institute of Informatics)

GraphTinker: A High Performance Data Structure for Dynamic Graph Processing .1030.....
Wole Jaiyeoba (University of Virginia) and Kevin Skadron (University of Virginia)

Session 25: Applications II

FastJoin: A Skewness-Aware Distributed Stream Join System .1042.....
Shunjie Zhou (Huazhong University of Science and Technology), Fan Zhang (Huazhong University of Science and Technology), Hanhua Chen (Huazhong University of Science and Technology), Hai Jin (Huazhong University of Science and Technology), and Bing Bing Zhou (The University of Sydney)

A Bin-Based Bitstream Partitioning Approach for Parallel CABAC Decoding in Next Generation Video Coding .1053.....
Philipp Habermann (Technische Universität Berlin), Chi Ching Chi (Spin Digital Video Technologies GmbH), Mauricio Alvarez-Mesa (Spin Digital Video Technologies GmbH), and Ben Juurlink (Technische Universität Berlin)

Stochastic Gradient Descent on Modern Hardware: Multi-core CPU or GPU? Synchronous or Asynchronous?
1063
Yujing Ma (University of California Merced), Florin Rusu (University of California Merced), and Martin Torres (University of California Merced)

Session 26: Security and Reliability

Always be Two Steps Ahead of Your Enemy .1073.....
Thorsten Götte (Paderborn University), Vipin Ravindran Vijayalakshmi (RWTH Aachen), and Christian Scheideler (Paderborn University)

Peace Through Superior Puzzling: An Asymmetric Sybil Defense .1083.....
Diksha Gupta (University of New Mexico), Jared Saia (University of New Mexico), and Maxwell Young (Mississippi State University)

Rethinking Support for Region Conflict Exceptions .1095.....
*Swarnendu Biswas (Indian Institute of Technology Kanpur), Rui Zhang
(Ohio State University), Michael D. Bond (Ohio State University), and
Brandon Lucia (Carnegie Mellon University)*

Author Index