

2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS 2018)

**Vancouver, British Columbia, Canada
21-25 May 2018**

Pages 1-588



**IEEE Catalog Number: CFP18023-POD
ISBN: 978-1-5386-4369-3**

**Copyright © 2018 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: | CFP18023-POD |
| ISBN (Print-On-Demand): | 978-1-5386-4369-3 |
| ISBN (Online): | 978-1-5386-4368-6 |
| 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

2018 IEEE International Parallel and Distributed Processing Symposium IPDPS 2018

Table of Contents

| | |
|--|--|
| Message from the General Chair .xx..... | |
| Message from the Program Chair and Vice Chair .xxii..... | |
| Technical Program .xxiv..... | |
| Conference Organization .xxvii..... | |
| Outside Reviewers .xxxvi..... | |

Keynote 1

| | |
|---|--|
| The Algorithmics of Write Optimization .1..... | |
| <i>Michael A. Bender (Stony Brook University)</i> | |

Session 1: Graph Algorithms 1

| | |
|--|--|
| MIDAS: Multilinear Detection at Scale .2..... | |
| <i>Saliya Ekanayake (Virginia Tech), Jose Cadena (Virginia Tech), Udayanga Wickramasinghe (Indiana University), and Anil Vullikanti (Virginia Tech)</i> | |
| Optimizing Parallel Graph Connectivity Computation via Subgraph Sampling .12..... | |
| <i>Michael Sutton (The Hebrew University of Jerusalem), Tal Ben-Nun (ETH Zurich), and Amnon Barak (The Hebrew University of Jerusalem)</i> | |
| Parallel Algorithms Through Approximation: B-Edge Cover .22..... | |
| <i>S M Ferdous (Purdue University), Arif Khan (Pacific Northwest National Lab), and Alex Pothen (Computer Science Purdue University)</i> | |
| A Parallel Algorithm for Bayesian Network Inference Using Arithmetic Circuits .34..... | |
| <i>Md Vasimuddin (Indian Institute of Technology Bombay), Sriram P. Chockalingam (Georgia Institute of Technology), and Srinivas Aluru (Georgia Institute of Technology)</i> | |

Session 2: Large-Scale Applications 1

| | |
|--|--|
| Cataloging the Visible Universe Through Bayesian Inference at Petascale .44..... | |
| <i>Jeffrey Regier (University of California), Kiran Pamnany (Intel Corporation), Keno Fischer, Andreas Noack, Maximilian Lam, Jarrett Revels, Steve Howard, Ryan Giordano, David Schlegel, Jon McAuliffe, Rollin Thomas, and Prabhat _ (Lawrence Berkeley National Laboratory)</i> | |

| | |
|--|--|
| Efficient, Parallel At-Scale Correlation Analysis for Atom Probe Tomography on Hybrid Architectures .54..... | |
| | <i>Hao Lu (Oak Ridge National Laboratory), Sudip K. Seal (Oak Ridge National Laboratory), Gregory Muzyn (University of Tennessee), Wei Guo (Oak Ridge National Laboratory), and Jonathan D. Poplawsky (Oak Ridge National Laboratory)</i> |
| A Fast and Massively-Parallel Inverse Solver for Multiple-Scattering Tomographic Image Reconstruction .64..... | |
| | <i>Mert Hidayetoglu (University of Illinois at Urbana-Champaign), Carl Pearson (University of Illinois at Urbana-Champaign), Izzat El Hajj (University of Illinois at Urbana-Champaign), Levent Gurel (University of Illinois at Urbana-Champaign), Weng Cho Chew (University of Illinois at Urbana-Champaign), and Wen-Mei Hwu (University of Illinois at Urbana-Champaign)</i> |
| Real-Time Massively Distributed Multi-object Adaptive Optics Simulations for the European Extremely Large Telescope .75..... | |
| | <i>Hatem Ltaief (KAUST), Ali Charara (KAUST), Damien Gratadour (Observatoire de Paris), Nicolas Doucet (Observatoire de Paris), Bilel Hadri (KAUST), Eric Gendron (Observatoire de Paris), Saber Feki (KAUST), and David Keyes (KAUST)</i> |

Session 3: Performance / QoS / Resilience

| | |
|---|---|
| Performance Isolation of Data-Intensive Scale-out Applications in a Multi-tenant Cloud .85..... | |
| | <i>Palden Lama (University of Texas at San Antonio), Shaoqi Wang (University of Colorado), Xiaobo Zhou (University of Colorado), and Dazhao Cheng (University of North Carolina at Charlotte)</i> |
| QoS Support for Scientific Workflows Using Software-Defined Storage Resource Enclaves .95..... | |
| | <i>Suman Karki (Washington State University Vancouver), Bao Nguyen (Washington State University Vancouver), and Xuechen Zhang (Washington State University Vancouver)</i> |
| Scalable Data Resilience for In-Memory Data Staging .105..... | |
| | <i>Shaohua Duan (Rutgers Discovery Informatics Institute), Pradeep Subedi (Rutgers Discovery Informatics Institute), Keita Teranishi (Sandia National Laboratory), Philip Davis (Rutgers Discovery Informatics Institute), Hemanth Kolla (Sandia National Laboratory), Marc Gamell (Intel Corporation), and Manish Parashar (Rutgers Discovery Informatics Institute)</i> |
| Performance and Scalability of Lightweight Multi-kernel Based Operating Systems .116..... | |
| | <i>Balazs Gerofi (RIKEN Advanced Institute For Computational Science), Rolf Riesen (Intel Corporation), Masamichi Takagi (RIKEN Advanced Institute For Computational Science), Taisuke Boku (University of Tsukuba), Kengo Nakajima (University of Tokyo), Yutaka Ishikawa (RIKEN Advanced Institute For Computational Science), and Robert W. Wisniewski (Intel Corporation)</i> |

Session 4: Memory Designs and Optimizations

| | |
|--|--|
| Architectural Support for Unlimited Memory Versioning and Renaming .126..... | |
| <i>Eran Gilad (Technion - Israel Institute of Technology), Tehila Mayzels (Technion - Israel Institute of Technology), Elazar Raab (Technion - Israel Institute of Technology), Mark Oskin (University of Washington), and Yoav Etsion (Technion - Israel Institute of Technology)</i> | |
| CTA-Aware Prefetching and Scheduling for GPU .137..... | |
| <i>Gunjae Koo (University of Southern California), Hyeran Jeon (San Jose State University), Zhenhong Liu (University of Illinois at Urbana-Champaign), Nam Sung Kim (University of Illinois at Urbana-Champaign), and Murali Annavaram (University of Southern California)</i> | |
| CIAO: Cache Interference-Aware Throughput-Oriented Architecture and Scheduling for GPUs .149..... | |
| <i>Jie Zhang (Yonsei University), Shuwen Gao (Intel), Nam Sung Kim (University of Illinois Urbana-Champaign), and Myoungsoo Jung (Yonsei University)</i> | |
| Millipede: Die-Stacked Memory Optimizations for Big Data Machine Learning Analytics .160..... | |
| <i>Nitin _ (NVIDIA), Mithuna Thottethodi (Purdue University), and T. N. Vijaykumar (Purdue University)</i> | |

Session 5: Scheduling

| | |
|---|--|
| Scheduling Monotone Moldable Jobs in Linear Time .172..... | |
| <i>Klaus Jansen (University of Kiel) and Felix Land (University of Kiel)</i> | |
| The Power to Schedule a Parallel Program .182..... | |
| <i>Kunal Agrawal (Washington University in Saint Louis) and Seth Gilbert (National University of Singapore)</i> | |
| Scheduling Parallel Tasks under Multiple Resources: List Scheduling vs. Pack Scheduling .194..... | |
| <i>Hongyang Sun (Vanderbilt University), Redouane Elghazi (ENS Lyon), Ana Gainaru (Vanderbilt University), Guillaume Aupy (INRIA), and Padma Raghavan (Vanderbilt University)</i> | |
| Parallel Scheduling of DAGs under Memory Constraints .204..... | |
| <i>Loris Marchal (CNRS), Hanna Nagy (Technical University of Cluj-Napoca), Bertrand Simon (CNRS), and Frédéric Vivien (CNRS)</i> | |

Session 6: Learning

| | |
|--|--|
| Evaluating Active Learning with Cost and Memory Awareness .214..... | |
| <i>Dmitry Duplyakin (University of Utah), Jed Brown (University of Colorado), and Donna Calhoun (Boise State University)</i> | |
| Semantics-Preserving Parallelization of Stochastic Gradient Descent .224..... | |
| <i>Saeed Maleki (Microsoft Research), Madanlal Musuvathi (Microsoft Research), and Todd Mytkowicz (Microsoft Research)</i> | |

| | |
|--|--|
| Efficient Gradient Boosted Decision Tree Training on GPUs .234..... | |
| | <i>Zeyi Wen (National University of Singapore), Bingsheng He (National University of Singapore), Ramamohanarao Kotagiri (The University of Melbourne), Shengliang Lu (National University of Singapore), and Jiashuai Shi (South China University of Technology)</i> |
| BitFlow: Exploiting Vector Parallelism for Binary Neural Networks on CPU .244..... | |
| | <i>Yuwei Hu (TuSimple Inc), Jidong Zhai (Tsinghua University), Dinghua Li (TuSimple Inc), Yifan Gong (TuSimple Inc), Yuhao Zhu (University of Rochester), Wei Liu (TuSimple Inc), Lei Su (TuSimple Inc), and Jiangming Jin (TuSimple Inc)</i> |

Session 7: Compilers and Libraries

| | |
|---|---|
| Lightweight MPI Communicators with Applications to Perfectly Balanced Quicksort .254..... | |
| | <i>Michael Axtmann (Karlsruhe Institute of Technology), Armin Wiebigke (Karlsruhe Institute of Technology), and Peter Sanders (Karlsruhe Institute of Technology)</i> |
| Improving Network Throughput with Global Communication Reordering .266..... | |
| | <i>Wim Lavrijsen (LBNL), Costin Iancu (LBNL), and Xing Pan (UNC)</i> |
| Highly Efficient Compensation-Based Parallelism for Wavefront Loops on GPUs .276..... | |
| | <i>Kaixi Hou (Virginia Tech), Hao Wang (Virginia Tech), Wu-chun Feng (Virginia Tech), Jeffrey S. Vetter (Oak Ridge National Lab), and Seyong Lee (Oak Ridge National Lab)</i> |
| Development and Application of a Hybrid Programming Environment on an ARM/DSP System for High Performance Computing .286..... | |
| | <i>Gaurav Mitra (Texas Instruments Inc.), Jonathan Bohmann (Southwest Research Institute), Ian Lintault (nCore HPC), and Alistair P. Rendell (Australian National University)</i> |

Session 8: Optimizations for Emerging Storage Systems

| | |
|--|--|
| GC-Aware Request Steering with Improved Performance and Reliability for SSD-Based RAIDs .296..... | |
| | <i>Suzhen Wu (Xiamen University), Weidong Zhu (Xiamen University), Guixin Liu (Xiamen University), Hong Jiang (University of Texas-Arlington), and Bo Mao (Xiamen University)</i> |
| A Set-aware Key-Value Store on Shingled Magnetic Recording Drives with Dynamic Band .306..... | |
| | <i>Ting Yao (Wuhan National Laboratory for Optoelectronics), Zhihu Tan (School of Computer Science and Technology), Jiguang Wan (Wuhan National Laboratory for Optoelectronics), Ping Huang (Temple University), Yiwen Zhang (Wuhan National Laboratory for Optoelectronics), Changsheng Xie (Wuhan National Laboratory for Optoelectronics), and Xubin He (Temple University)</i> |
| Software-Hardware Managed Last-level Cache Allocation Scheme for Large-Scale NVRAM-Based Multicores Executing Parallel Data Analytics Applications .316..... | |
| | <i>Masab Ahmad (University of Connecticut), Halit Dogan (University of Connecticut), Fabio Checconi (IBM Research), Xinyu Que (IBM Research), Daniele Buono (IBM Research), and Omer Khan (University of Connecticut)</i> |

MOCA: Memory Object Classification and Allocation in Heterogeneous Memory Systems .326.....
Aditya Narayan (Boston University), Tiansheng Zhang (Boston University), Shaizeen Aga (University of Michigan), Satish Narayanasamy (University of Michigan), and Ayse Coskun (Boston University)

Best Paper Nominees - Plenary

Communication-Free Massively Distributed Graph Generation .336.....
Daniel Funke (Karlsruhe Institute of Technology), Sebastian Lamm (Karlsruhe Institute of Technology), Peter Sanders (Karlsruhe Institute of Technology), Christian Schulz (Karlsruhe Institute of Technology), Darren Strash (Karlsruhe Institute of Technology), and Moritz von Looz (Karlsruhe Institute of Technology)

Understanding and Modeling Lossy Compression Schemes on HPC Scientific Data .348.....
Tao Lu (New Jersey Institute of Technology), Qing Liu (New Jersey Institute of Technology), Xubin He (Temple University), Huizhang Luo (New Jersey Institute of Technology), Eric Suchyta (Oak Ridge National Laboratory), Jong Choi (Oak Ridge National Laboratory), Norbert Podhorszki (Oak Ridge National Laboratory), Scott Klasky (Oak Ridge National Laboratory), Mathew Wolf (Oak Ridge National Laboratory), Tong Liu (Temple University), and Zhenbo Qiao (New Jersey Institute of Technology)

UBIS: Utilization-Aware Cluster Scheduling .358.....
Karthik Kambatla (Facebook Inc), Vamsee Yarlagadda (Cloudera Inc.), ĩñigo Goiri (Microsoft Research), and Ananth Grama (Purdue University)

Hardware Transactional Memory Meets Memory Persistency .368.....
Daniel Castro (Instituto Superior Técnico & INESC-ID), Paolo Romano (Instituto Superior Técnico & INESC-ID), and João Barreto (Instituto Superior Técnico & INESC-ID)

Keynote 2

Empowering Flexible and Scalable High Performance Architectures with Embedded Photonics .378.....
Keren Bergman (Columbia University)

Session 9: Numerical Algorithms

Large Bandwidth-Efficient FFTs on Multicore and Multi-socket Systems .379.....
Doru Thom Popovici (Carnegie Mellon University), Tze Meng Low (Carnegie Mellon University), and Franz Franchetti (Carnegie Mellon University)

Lattice H-Matrices on Distributed-Memory Systems .389.....
Akihiro Ida (The University of Tokyo)

| | |
|---|---|
| Evaluating the Performance and Cost of Accelerating Seismic Processing with CUDA, OpenCL, OpenACC, and OpenMP .399..... | |
| | <i>Tiago Lobato Gimenes (Institute of Computing - University of Campinas), Flávia Pisani (Institute of Computing - University of Campinas), and Edson Borin (Institute of Computing - University of Campinas)</i> |
| Avoiding Synchronization in First-Order Methods for Sparse Convex Optimization .409..... | |
| | <i>Aditya Devarakonda (University of California), Kimon Fountoulakis (University of California), James Demmel (University of California), and Michael W. Mahoney (University of California)</i> |

Session 10: GPU Hashing and Searching

| | |
|---|---|
| A Dynamic Hash Table for the GPU .419..... | |
| | <i>Saman Ashkiani (University of California), Martin Farach-Colton (Rutgers University), and John D. Owens (University of California)</i> |
| GPU LSM: A Dynamic Dictionary Data Structure for the GPU .430..... | |
| | <i>Saman Ashkiani (University of California), Shengren Li (University of California), Martin Farach-Colton (Rutgers University), Nina Amenta (University of California), and John D. Owens (University of California)</i> |
| WarpDrive: Massively Parallel Hashing on Multi-GPU Nodes .441..... | |
| | <i>Daniel Jünger (Johannes Gutenberg University), Christian Hundt (Johannes Gutenberg University), and Bertil Schmidt (Johannes Gutenberg University)</i> |
| Quotient Filters: Approximate Membership Queries on the GPU .451..... | |
| | <i>Afton Geil (University of California), Martin Farach-Colton (Rutgers University), and John Owens (University of California)</i> |

Session 11: Domain-Specific, Runtime and Autotuning

| | |
|---|--|
| BabelFlow: An Embedded Domain Specific Language for Parallel Analysis and Visualization .463..... | |
| | <i>Steve Petruzza (University of Utah), Sean Treichler (Stanford University), Valerio Pascucci (University of Utah), and Peer-Timo Bremer (Lawrence Livermore National Lab)</i> |
| Online Tuning of Parallelism Degree in Parallel Nesting Transactional Memory .474..... | |
| | <i>Jingna Zeng (Kungliga Tekniska högskolan), Paolo Romano (Instituto Superior Técnico), João Barreto (Instituto Superior Técnico), Luís Rodrigues (Instituto Superior Técnico), and Seif Haridi (Kungliga Tekniska högskolan/RISE SICS)</i> |
| Work-Stealing, Locality-Aware Actor Scheduling .484..... | |
| | <i>Saman Barghi (University of Waterloo) and Martin Karsten (University of Waterloo)</i> |

Indigo: A Domain-Specific Language for Fast, Portable Image Reconstruction .495.....
Michael Driscoll (UC Berkeley and LBNL), Benjamin Brock (UC Berkeley and LBNL), Frank Ong (UC Berkeley), Jonathan Tamir (UC Berkeley), Hsiou-Yuan Liu (UC Berkeley), Michael Lustig (UC Berkeley), Armando Fox (UC Berkeley), and Katherine Yelick (UC Berkeley and LBNL)

Session 12: Resource Management

Swallow: Joint Online Scheduling and Coflow Compression in Datacenter Networks .505.....
Qihua Zhou (Nanjing University of Posts and Telecommunications), Peng Li (The University of Aizu), Kun Wang (Nanjing University of Posts and Telecommunications), Deze Zeng (China University of Geoscience), Song Guo (The Hong Kong Polytechnic University), and Minyi Guo (Shanghai Jiao Tong University)

Auto-tuning Streamed Applications on Intel Xeon Phi .515.....
Peng Zhang (National University of Defense Technology), Jianbin Fang (National University of Defense Technology), Tao Tang (National University of Defense Technology), Canqun Yang (National University of Defense Technology), and Zheng Wang (Lancaster University)

Analyzing Resource Trade-offs in Hardware Overprovisioned Supercomputers .526.....
Ryuichi Sakamoto (The University of Tokyo), Tapasya Patki (Lawrence Livermore National Laboratory), Thang Cao (The University of Tokyo), Masaaki Kondo (The University of Tokyo), Koji Inoue (Kyushu University), Masatsugu Ueda (Kyushu University), Daniel Ellsworth (Colorado College), Barry Rountree (Lawrence Livermore National Laboratory), and Martin Schulz (Technical University of Munich)

Harnessing the Power of Many: Extensible Toolkit for Scalable Ensemble Applications .536.....
Vivek Balasubramanian (Rutgers University), Matteo Turilli (Rutgers University), Weiming Hu (Penn State University), Matthieu Lefebvre (Princeton University), Wenjie Lei (Princeton University), Ryan Modrak (Princeton University), Guido Cervone (Penn State University), Jeroen Tromp (Princeton University), and Shantenu Jha (Brookhaven National Laboratory)

Session 13: Tensors

A Fill Estimation Algorithm for Sparse Matrices and Tensors in Blocked Formats .546.....
Peter Ahrens (Massachusetts Institute of Technology), Helen Xu (Massachusetts Institute of Technology), and Nicholas Schiefer (Massachusetts Institute of Technology)

Communication Lower Bounds for Matricized Tensor Times Khatri-Rao Product .557.....
Grey Ballard (Wake Forest University), Nicholas Knight (New York University), and Kathryn Rouse (Wake Forest University)

Blocking Optimization Techniques for Sparse Tensor Computation .568.....
Jee Choi (IBM T. J. Watson Research), Xing Liu (Intel), Shaden Smith (University of Minnesota), and Tyler Simon (University of Maryland)

TTLG - An Efficient Tensor Transposition Library for GPUs .578.....
*Jyothi Vedurada (IIT Madras), Arjun Suresh (Ohio State University),
 Aravind Sukumaran Rajam (Ohio State University), Jinsung Kim (Ohio
 State University), Changwan Hong (Ohio State University), Ajay Panyala
 (PNNL), Sriram Krishnamoorthy (PNNL), V. Krishna Nandivada (IIT
 Madras), Rohit Kumar Srivastava (Ohio State University), and P.
 Sadayappan (Ohio State University)*

Session 14: Large Scale Applications 2

Do Developers Understand IEEE Floating Point? .589.....
*Peter Dinda (Northwestern University) and Conor Hetland (Northwestern
 University)*

sDPF-RSA: Utilizing Floating-point Computing Power of GPUs for Massive Digital Signature
 Computations .599.....
*Jiankuo Dong (Chinese Academy of Sciences, Beijing), Fangyu Zheng
 (Chinese Academy of Sciences, Beijing), Niall Emmart (University of
 Massachusetts), Jingqiang Lin (Chinese Academy of Sciences, Beijing),
 and Charles Weems (University of Massachusetts)*

Rethinking large-scale Economic Modeling for Efficiency: Optimizations for GPU and Xeon Phi Clusters .610.
*Simon Scheidegger (University of Zurich), Dmitry Mikushin (University
 of Zurich), Felix Kubler (University of Zurich), and Olaf Schenk
 (Universit a della Svizzera italiana, Switzerland)*

A Fast Scalable Implicit Solver with Concentrated Computation for Nonlinear Time-Evolution Problems
 on Low-Order Unstructured Finite Elements .620.....
*Tsuyoshi Ichimura (The University of Tokyo and RIKEN), Kohei Fujita
 (The University of Tokyo and RIKEN), Masashi Horikoshi (Intel K.K.),
 Larry Meadows (Intel Corporation), Kengo Nakajima (The University of
 Tokyo and RIKEN), Takuma Yamaguchi (The University of Tokyo), Kentaro
 Koyama (Fujitsu Limited), Hikaru Inoue (Fujitsu Limited), Akira Naruse
 (NVIDIA Corporation), Keisuke Katsushima (The University of Tokyo),
 Muneo Hori (The University of Tokyo and RIKEN), and Lalith Maddegedara
 (The University of Tokyo and RIKEN)*

Session 15: Data Operations

Characterizing Scheduling Delay for Low-Latency Data Analytics Workloads .630.....
*Wei Chen (University of Colorado), Aidi Pi (University of Colorado),
 Shaoqi Wang (University of Colorado), and Xiaobo Zhou (University of
 Colorado)*

Runtime Scheduling Policies for Distributed Graph Algorithms .640.....
*Jesun Sahariar Firoz (Pacific Northwest National Lab and Indiana
 University Bloomington), Marcin Zalewski (Pacific Northwest National
 Lab), Andrew Lumsdaine (Pacific Northwest National Lab and University
 of Washington), and Martina Barnas (Indiana University Bloomington)*

Communication Efficient Checking of Big Data Operations .650.....
*Lorenz Hübschle-Schneider (Karlsruhe Institute of Technology) and
 Peter Sanders (Karlsruhe Institute of Technology)*

What Size Should Your Buffers to Disks be? .660.....
Guillaume Aupy (Inria), Olivier Beaumont (Inria), and Lionel Eyraud-Dubois (Inria)

Session 16: Power and Temperature

THOR: THERmal-aware Optimizations for extending ReRAM Lifetime .670.....
Majed Valad Beigi (Northwestern University) and Gokhan Memik (Northwestern University)

CoolPIM: Thermal-Aware Source Throttling for Efficient PIM Instruction Offloading .680.....
Lifeng Nai (Google), Ramyad Hadidi (Georgia Institute of Technology), He Xiao (Georgia Institute of Technology), Hyojong Kim (Georgia Institute of Technology), Jaewoong Sim (Intel Labs), and Hyesoon Kim (Georgia Institute of Technology)

GreenSprint: Effective Computational Sprinting in Green Data Centers .690.....
Haoran Cai (Huazhong University of Science and Technology), Xu Zhou (Sangfor Technologies Co.), Qiang Cao (Huazhong University of Science and Technology), Hong Jiang (University of Texas at Arlington), Feng Sheng (Huazhong University of Science and Technology), Xiandong Qi (Hong Kong University of Science and Technology), Jie Yao (Huazhong University of Science and Technology), Changsheng Xie (Huazhong University of Science and Technology), Liang Xiao (Huazhong University of Science and Technology), and Liang Gu (Sangfor Technologies Co.)

Joint Server and Network Energy Saving in Data Centers for Latency-Sensitive Applications .700.....
Liang Zhou (University of California Riverside), Chih-Hsun Chou (University of California Riverside), Laxmi N. Bhuyan (University of California Riverside), K. K. Ramakrishnan (University of California Riverside), and Daniel Wong (University of California Riverside)

Keynote 3

The Day After Tomorrow: The Looming Post-Exascale Crisis .710.....
Bruce Hendrickson (Lawrence Livermore National Lab)

Session 17: Graph Algorithms 2

Implicit Decomposition for Write-Efficient Connectivity Algorithms .711.....
Naama Ben-David (Carnegie Mellon University), Guy Blelloch (Carnegie Mellon University), Jeremy Fineman (Georgetown University), Phillip Gibbons (Carnegie Mellon University), Yan Gu (Carnegie Mellon University), Charles McGuffey (Carnegie Mellon University), and Julian Shun (Carnegie Mellon University)

Distributed Symmetry Breaking in Graphs with Bounded Diversity .723.....
Leonid Barenboim (Open University of Israel) and Tzalik Maimon (Open University of Israel)

Complete Visitability for Autonomous Robots on Graphs .733.....
Aisha Aljohani (Kent State University), Pavan Poudel (Kent State University), and Gokarna Sharma (Kent State University)

Local Mixing Time: Distributed Computation and Applications .743.....
Anisur Rahaman Molla (National Institute of Science Education and Research (NISER) Bhubaneswar) and Gopal Pandurangan (University of Houston)

Session 18: Performance Modeling and Analysis

Roofline Guided Design and Analysis of a Multi-stencil CFD Solver for Multicore Performance .753.....
Bahareh Mostafazadeh (University of California Irvine), Ferran Marti (University of California Irvine), Feng Liu (University of California Irvine), and Aparna Chandramowlishwaran (University of California Irvine)

Taming the "Monster": Overcoming Program Optimization Challenges on SW26010 Through Precise Performance Modeling .763.....
Shizhen Xu (Tsinghua University), Yuanchao Xu (Tsinghua University), Wei Xue (Tsinghua University), Xipeng Shen (North Carolina State University), Fang Zheng (National Research Center of Parallel Computer Engineering and Technology), Xiaomeng Huang (Tsinghua University), and Guangwen Yang (Tsinghua University)

Performance and Accuracy Trade-offs of HPC Application Modeling and Simulation .774.....
Zhou Tong (Florida State University), Xin Yuan (Florida State University), Scott Pakin (Los Alamos National Lab), and Michael Lang (Los Alamos National Lab)

PADDLE: Performance Analysis Using a Data-Driven Learning Environment .784.....
Jayaraman J. Thiagarajan (Lawrence Livermore National Laboratory), Rushil Anirudh (Lawrence Livermore National Laboratory), Bhavya Kailkhura (Lawrence Livermore National Laboratory), Nikhil Jain (Lawrence Livermore National Laboratory), Tanzima Islam (Western Washington University), Abhinav Bhatele (Lawrence Livermore National Laboratory), Jae-Seung Yeom (Lawrence Livermore National Laboratory), and Todd Gamblin (Lawrence Livermore National Laboratory)

Session 19: Memory and Data Access

Efficient Solving of Scan Primitive on Multi-GPU Systems .794.....
Adrián Perez Diéguez (University of A Coruña), Margarita Amor (University of A Coruña), Ramón Doallo (University of A Coruña), Akira Nukada (Tokyo Institute of Technology), and Satoshi Matsuoka (Tokyo Institute of Technology)

Quantifying the Performance and Energy-Efficiency Impact of Hardware Transactional Memory on Scientific Applications on Large-Scale NUMA Systems .804.....
Jinsu Park (UNIST) and Woongki Baek (UNIST)

| | |
|---|---|
| GPU-Accelerated Large-Scale Genome Assembly .814 | |
| | <i>Sayan Goswami (Louisiana State University), Kisung Lee (Louisiana State University), Shayan Shams (Louisiana State University), and Seung-Jong Park (Louisiana State University)</i> |
| GPU Data Access on Complex Geometries for D3Q19 Lattice Boltzmann Method .825 | |
| | <i>Gregory Herschlag (Duke University), Seyong Lee (Oak Ridge National Laboratory), Jeffrey S. Vetter (Oak Ridge National Laboratory), and Amanda Randles (Duke University)</i> |

Session 20: Exception Handling & Error Detection

| | |
|---|--|
| SLIMFAST: Reducing Metadata Redundancy in Sound and Complete Dynamic Data Race Detection .835 | |
| | <i>Yuanfeng Peng (University of Pennsylvania), Christian DeLozier (University of Pennsylvania), Ariel Eizenberg (University of Pennsylvania), William Mansky (Princeton University), and Joseph Devietti (University of Pennsylvania)</i> |
| SWORD: A Bounded Memory-Overhead Detector of OpenMP Data Races in Production Runs .845 | |
| | <i>Simone Atzeni (University of Utah), Ganesh Gopalakrishnan (University of Utah), Zvonimir Rakamaric (University of Utah), Ignacio Laguna (Lawrence Livermore National Laboratory), Gregory L. Lee (Lawrence Livermore National Laboratory), and Dong H. Ahn (Lawrence Livermore National Laboratory)</i> |
| Unobtrusive Asynchronous Exception Handling with Standard Java Try/Catch Blocks .855 | |
| | <i>Mostafa Mehrabi (PhD candidate at the university of Auckland), Nasser Giacaman (Lecturer at the university of Auckland), and Oliver Sinnen (Senior lecturer at the university of Auckland)</i> |
| COMPI: Concolic Testing for MPI Applications .865 | |
| | <i>Hongbo Li (University of California, Riverside), Sihuan Li (University of California, Riverside), Zachary Benavides (University of California, Riverside), Zizhong Chen (University of California, Riverside), and Rajiv Gupta (University of California, Riverside)</i> |

Session 21: Graph Algorithms 3

| | |
|---|--|
| Experimental Design of Work Chunking for Graph Algorithms on High Bandwidth Memory Architectures .875 | |
| | <i>George M Slota (Rensselaer Polytechnic Institute) and Siva Rajamanickam (Sandia National Labs)</i> |
| Distributed Louvain Algorithm for Graph Community Detection .885 | |
| | <i>Sayan Ghosh (Washington State University), Mahantesh Halappanavar (Pacific Northwest National Laboratory), Antonino Tumeo (Pacific Northwest National Laboratory), Ananth Kalyanaraman (Washington State University), Hao Lu (Oak Ridge National Laboratory), Daniel Chavarrià-Miranda (Trovares), Arif Khan (Pacific Northwest National Laboratory), and Assefaw Gebremedhin (Washington State University)</i> |

Application Codesign of Near-Data Processing for Similarity Search .896.....
Vincent T. Lee (University of Washington), Amrita Mazumdar (University of Washington), Carlo C. del Mundo (University of Washington), Armin Alaghi (University of Washington), Luis Ceze (University of Washington), and Mark Oskin (University of Washington)

Session 22: Linear Solvers

A Communication-Avoiding 3D LU Factorization Algorithm for Sparse Matrices .908.....
Piyush Sao (Georgia Institute of Technology), Xiaoye Sherry Li (Lawrence Berkeley National Lab), and Richard Vuduc (Georgia Institute of Technology)

A New GPU Algorithm to Compute a Level Set-Based Analysis for the Parallel Solution of Sparse Triangular Systems .920.....
Ernesto Dufrechou (Instituto de Computación) and Pablo Ezzatti (Instituto de Computación)

Performance of Hierarchical-matrix BiCGStab Solver on GPU Clusters .930.....
Ichitaro Yamazaki (The University of Tennessee), Ahmad Abdelfattah (The University of Tennessee), Akihiro Ida (The University of Tokyo), Satoshi Ohshima (Kyushu University), Stanimire Tomov (The University of Tennessee), Rio Yokota (Tokyo Institute of Technology), and Jack Dongarra (The University of Tennessee)

Convergence Models and Surprising Results for the Asynchronous Jacobi Method .940.....
Edmond Chow (Georgia Institute of Technology) and Edmond Chow (Georgia Institute of Technology)

Session 23: Runtime Systems and Libraries

Overhead-Conscious Format Selection for SpMV-Based Applications .950.....
Yue Zhao (North Carolina State University), Weijie Zhou (North Carolina State University), Xipeng Shen (North Carolina State University), and Graham Yiu (IBM Toronto Software Lab)

Cudele: An API and Framework for Programmable Consistency and Durability in a Global Namespace .960...
Michael A. Sevilla (University of California, Santa Cruz), Ivo Jimenez (University of California, Santa Cruz), Noah Watkins (University of California, Santa Cruz), Jeff LeFevre (University of California, Santa Cruz), Peter Alvaro (University of California, Santa Cruz), Shel Finkelstein (University of California, Santa Cruz), Patrick Donnelly (Red Hat), and Carlos Maltzahn (University of California, Santa Cruz)

SELECT: A Distributed Publish/Subscribe Notification System for Online Social Networks .970.....
Nuno Apolónia (Universitat Politècnica de Catalunya), Stefanos Antaris (University of Cyprus), Sarunas Girdzijauskas (Royal Institute of Technology), George Pallis (University of Cyprus), and Marios Dikaiakos (University of Cyprus)

A Lightweight Communication Runtime for Distributed Graph Analytics .980.....
Hoang-Vu Dang (University of Illinois at Urbana-Champaign), Roshan Dathathri (The University of Texas at Austin), Gurbinder Gill (The University of Texas at Austin), Alex Brooks (University of Illinois at Urbana-Champaign), Nikoli Dryden (University of Illinois at Urbana-Champaign), Andrew Lenharth (The University of Texas at Austin), Loc Hoang (The University of Texas at Austin), Keshav Pingali (The University of Texas at Austin), and Marc Snir (University of Illinois at Urbana-Champaign)

Session 24: Networks and Communication

Intra-Cluster Coalescing to Reduce GPU NoC Pressure .990.....
Lu Wang (Ghent University), Xia Zhao (Ghent University), David Kaeli (Northeastern University), Zhiying Wang (National University of Defense Technology), and Lieven Eeckhout (Ghent University)

HybridPass: Hybrid Scheduling for Mixed Flows in Datacenter Networks .1000.....
Bo Peng (Shanghai Jiao Tong University), Jianguo Yao (Shanghai Jiao Tong University), Zhengwei Qi (Shanghai Jiao Tong University), and Haibing Guan (Shanghai Jiao Tong University)

Scalable Power-Efficient Kilo-Core Photonic-Wireless NoC Architectures .1010.....
Avinash Kodi (Ohio University), Kyle Shifflet (Ohio University), Savas Kaya (Ohio University), Soumyasanta Laha (Ohio University), and Ahmed Louri (George Washington University)

Designing Efficient Shared Address Space Reduction Collectives for Multi-/Many-cores .1020.....
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 Dhableswar K. Panda (The Ohio State University)

Session 25: Distributed Computing

Tiny Groups Tackle Byzantine Adversaries .1030.....
Mercy O. Jaiyeola (Mississippi State University), Kyle Patron (Palantir Technologies), Jared Saia (University of New Mexico), Maxwell Young (Mississippi State University), and Qian M. Zhou (Mississippi State University)

Skueue: A Scalable and Sequentially Consistent Distributed Queue .1040.....
Michael Feldmann (Paderborn University), Christian Scheideler (Paderborn University), and Alexander Setzer (Paderborn University)

Self-Stabilizing Supervised Publish-Subscribe Systems .1050.....
Michael Feldmann (Paderborn University), Christina Kolb (Paderborn University), Christian Scheideler (Paderborn University), and Tim Strothmann (Paderborn University)

Spartan: A Framework For Sparse Robust Addressable Networks .1060.....
John Augustine (Indian Institute of Technology Madras) and Sumathi Sivasubramaniam (Indian Institute of Technology Madras)

Session 26: Graph Algorithms 4

- Beyond Binary Search: Parallel In-Place Construction of Implicit Search Tree Layouts .1070.....
Kyle Berney (University of Hawaii at Manoa), Henri Casanova (University of Hawaii at Manoa), Alyssa Higuchi (University of Hawaii at Manoa), Ben Karsin (University of Hawaii at Manoa), and Nodari Sitchinava (University of Hawaii at Manoa)
- An Energy-Efficient Single-Source Shortest Path Algorithm .1080.....
Sara Karamati (Georgia Institute of Technology), Jeffrey Young (Georgia Institute of Technology), and Richard Vuduc (Georgia Institute of Technology)
- Scalable Breadth-First Search on a GPU Cluster .1090.....
Yuechao Pan (University of California, Davis), Roger Pearce (Lawrence Livermore National Laboratory), and John D. Owens (University of California, Davis)

Session 27: Communication Performance

- Chameleon: Online Clustering of MPI Program Traces .1102.....
Amir Bahmani (NorthCarolina State Unievrity) and Frank Mueller (NorthCarolina State Unievrity)
- Trade-Off Study of Localizing Communication and Balancing Network Traffic on a Dragonfly System .1113.
Xin Wang (Illinois Institute of Technology), Misbah Mubarak (Argonne National Laboratory), Xu Yang (Illinois Institute of Technology), Robert B. Ross (Argonne National Laboratory), and Zhiling Lan (Illinois Institute of Technology)
- Level-Spread: A New Job Allocation Policy for Dragonfly Networks .1123.....
Yijia Zhang (Boston University), Ozan Tuncer (Boston University), Fulya Kaplan (Boston University), Katzalin Olcoz (Universidad Complutense de Madrid), Vitus J. Leung (Sandia National Laboratories), and Ayse K. Coskun (Boston University)

Session 28: Storage & FileSystem

- A Migratory Heterogeneity-Aware Data Layout Scheme for Parallel File Systems .1133.....
Shuibing He (Wuhan University), Xian-He Sun (Illinois Institute of Technology), Yang Wang (Shenzhen Institute of Advanced Technology), and Chengzhong Xu (Shenzhen Institute of Advanced Technology)
- LALCA: Locality-Aware Lock Contention Avoidance for NVMe-Based Scale-out Storage System .1143.....
Myoungwon Oh (SK Telecom), Sejin Park (Keimyong University), Jugwan Eom (SK Telecom), Seungmin Kim (SK Telecom), Sangjae Kim (SK Telecom), Kang-won Lee (SK Telecom), and Heon Y Yeom (Seoul National University)

Mitigating Traffic-Based Side Channel Attacks in Bandwidth-Efficient Cloud Storage .1153.....
*Pengfei Zuo (Huazhong University of Science and Technology), Yu Hua,
 Cong Wang, Wen Xia (Huazhong University of Science and Technology),
 Shunde Cao (Huazhong University of Science and Technology), Yukun Zhou
 (Huazhong University of Science and Technology), and Yuanyuan Sun
 (Huazhong University of Science and Technology)*

Chameleon: An Adaptive Wear Balancer for Flash Clusters .1163.....
*Nannan Zhao (Virginia Tech), Ali Anware (Virginia Tech), Yue Cheng
 (George Mason University), Mohammed Salman (Virginia Tech), Daping Li
 (Huazhong University of Science and Technology), Jiguang Wan (Huazhong
 University of Science and Technology), Changsheng Xie (Huazhong
 University of Science and Technology), Xubin He (Temple University),
 Feiyi Wang (Oak Ridge National Laboratory), and Ali Butt (Virginia
 Tech)*

Author Index