2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW 2018)

Vancouver, British Columbia, Canada 21-25 May 2018

Pages 1-665



IEEE Catalog Number: ISBN: CFP1851J-POD 978-1-5386-5556-6

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:	CFP1851J-POD
ISBN (Print-On-Demand):	978-1-5386-5556-6
ISBN (Online):	978-1-5386-5555-9

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



2018 IEEE International Parallel and Distributed Processing Symposium Workshops IPDPSW 2018

Table of Contents

Message from the General Chair .xxviii	
Message from the Workshops Chair and Vice Chair xxx	

HCW: Heterogeneity in Computing Workshop

Introduction to HCW 2018 .1 Alexey Lastovetsky (University College Dublin, Ireland) and Sudeep Pasricha (Colorado State University)
Message from the HCW Steering Committee Chair .2 Behrooz Shirazi (Washington State University)
Message from the HCW General Chair .3. Alexey Lastovetsky (University College Dublin, Ireland)
Message from the HCW Program Committee Chair .4 Sudeep Pasricha (Colorado State University)
HCW 2018 Keynote Talk 1 .5 Manish Parashar (Rutgers)
HCW 2018 Keynote Talk 2 .6 Umit Catalyurek (Georgia Institute of Technology)

Session 1: Reconfigurable and Cloud Systems

User-Transparent Translation of Machine Instructions to Programmable Hardware .7 Leslie Barron (University of Toronto) and Tarek S. Abdelrahman (University of Toronto)
Budget-Aware Scheduling Algorithms for Scientific Workflows with Stochastic Task Weights on
Heterogeneous IaaS Cloud Platforms .15
Yves Caniou (ENS Lyon), Eddy Caron (ENS Lyon), Aurélie Kong Win Chang
(ENS Lyon), and Yves Robert (ENS Lyon & University of Tennessee
Knoxville)
Optimizing Parallel Reduction on OpenCL FPGA Platform – A Case Study of Frequent Pattern Compression 27
Zheming Jin (Argonne National Lab) and Hal Finkel (Argonne National
Lab)

Session 2: Workload Scheduling and Architecture Analysis

Approximation Algorithm for Scheduling Applications on Hybrid Multi-core Machines with Communications Delays 36.
Massinissa Ait Aba (CEA, LIST, Computing and Design Environment Laboratory), Lilia Zaourar (CEA, LIST, Computing and Design Environment Laboratory), and Alix Munier (LIP6-UPMC)
Exploration and Design of a Synchronous Message Passing Framework for a CPU-NPU Heterogeneous
Architecture .46.
Sean Pennefather (Rhodes University), Karen Bradshaw (Rhodes
University), and Barry Irwin (Rhodes University)
Large Scale Data Centers Simulation Based on Baseline Test Model .5.7
Fei Lei (Beihang University), Lei Yu (Beihang University), Bing Shao
(Beihang University), Fei Teng (Southwest Jiaotong University), and Bo
Zhou (Huawei Technologies CO.)
Application Performance on a Cluster-Booster System .69.
Anke Kreuzer (Jülich Supercomputing Centre), Norbert Eicker (Jülich
Supercomputing Centre), Jorge Amaya (Kaholieke Universiteit Leuven),
and Estela Suarez (Juelich Supercomputing Centre)

RAW: Reconfigurable Architectures Workshop

ntroduction to RAW 2018.79
Marco D. Santambrogio (Politecnico di Milano, Italy), Diana Goehringer (TU Dresden, Germany), Dirk Stroobandt (Ghent University, Belgium), and Ken Eguro (Microsoft Research, USA)
RAW 2018 Invited Talks .8.1
Jürgen Becker (University of Karlsruhe), Viktor K. Prasanna
(University of Southern California), Markus Weimer (Microsoft), Wayne
Luk (Imperial College London), Kaveh Aasaraai (Jump Trading LLC), and
Derek Chiou (Microsoft)

Session 1: Platforms and Memory

Transport-Triggered Soft Cores 83
Pekka Jääskeläinen (Tampere University of Technology, Finland), Aleksi
Tervo (Tampere University of Technology, Finland), Guillermo Payá Vayá
(Leibniz Universität Hannover), Timo Viitanen (Tampere University of
Technology, Finland), Nicolai Behmann (Leibniz Universität Hannover),
Jarmo Takala (Tampere University of Technology, Finland), and Holger
Blume (Leibniz Universität Hannover)
OXiGen: A Tool for Automatic Acceleration of C Functions Into Dataflow FPGA-Based Kernels .9.1
Francesco Devenelli (Dolitecrico di Milano) Marco Dabozzi

Francesco Peverelli (Politecnico di Milano), Marco Rabozzi (Politecnico di Milano), Emanuele Del Sozzo (Politecnico di Milano), and Marco D. Santambrogio (Politecnico di Milano)

Session 2: Applications

An FPGA-Based Acceleration Methodology and Performance Model for Iterative Stencils .1.15 Enrico Reggiani (Politecnico di Milano), Giuseppe Natale (Politecnico di Milano), Carlo Moroni (E-lysis), and Marco D. Santambrogio (Politecnico di Milano)
High-Performance High-Order Stencil Computation on FPGAs Using OpenCL .123 Hamid Reza Zohouri (Tokyo Institute of Technology), Artur Podobas (Tokyo Institute of Technology), and Satoshi Matsuoka (Tokyo Institute of Technology)
TiReX: Tiled Regular eXpression Matching Architecture .131. Alessandro Comodi (Politecnico di Milano), Davide Conficconi (Politecnico di Milano), Alberto Scolari (Politecnico di Milano), and Marco D. Santambrogio (Politecnico di Milano)
Hardware Implementation of POSITs and Their Application in FPGAs .138 Artur Podobas (Tokyo Institute of Technology) and Satoshi Matsuoka (Tokyo Institute of Technology)

Session 3: Machine Learning 1

Robustness of Surface EMG Classifiers with Fixed-Point Decomposition on Reconfigurable Architecture .146. Luca Cerina (Politecnico Di Milano), Giuseppe Franco (BioRobotic Institute of Scuola Superiore Sant'Anna (SSSA) Pisa), Pierandrea Cancian (Politecnico Di Milano), and Marco D. Santambrogio (Politecnico Di Milano)
Hardware/Software Codesign for Convolutional Neural Networks Exploiting Dynamic Partial
Reconfiguration on PYNQ .154......
Florian Kästner (Ruhr-University Bochum), Benedikt Janßen (Ruhr-University Bochum), Frederik Kautz (Ruhr-University Bochum), Michael Hübner (Ruhr-University Bochum), and Giulio Corradi (Xilinx GmbH) Streaming Architecture for Large-Scale Quantized Neural Networks on an FPGA-Based Dataflow Platform .162 Chaim Baskin (Technion - Israel Institute of Technology), Natan Liss (Technion - Israel Institute of Technology), Evgenii Zheltonozhskii (Technion - Israel Institute of Technology), Alex M. Bronstein (Technion - Israel Institute of Technology), and Avi Mendelson (Technion - Israel Institute of Technology)

Session 4: Machine Learning 2

A Framework with Cloud Integration for CNN Acceleration on FPGA Devices .170 Niccolò Raspa (Politecnico di Milano), Giuseppe Natale (Politecnico di Milano), Marco Bacis (Politecnico di Milano), and Marco D. Santambrogio (Politecnico di Milano)
Kibo: An Open-Source Fixed-Point Tool-kit for Training and Inference in FPGA-Based Deep Learning
Networks .1.78
Daniel Holanda Noronha (University of British Columbia), Philip H.W.
Leong (University of Sydney), and Steven J.E. Wilton (University of
British Columbia)
A Reconfigurable Accelerator for Morphological Operations .186
Menbere Kina Tekleyohannes (University of Kaiserslautern), Christian
Weis (University of Kaiserslautern), Norbert Wehn (University of
Kaiserslautern), Martin Klein (Wipotec Wiege-und Positioniersysteme
GmbH), and Michael Siegrist (Wipotec Wiege-und Positioniersysteme
GmbH)

Session 5: Short Papers 1

MP-STREAM: A Memory Performance Benchmark for Design Space Exploration on Heterogeneous HPC Devices 194
Syed Waqar Nabi (University of Glasgow) and Wim Vanderbauwhede (University of Glasgow)
FIDA: A Framework to Automatically Integrate FPGA Kernels Within Data-Science Applications .198 Luca Stornaiuolo (Politecnico di Milano), Alberto Parravicini (Politecnico di Milano), Donatella Sciuto (Politecnico di Milano), and Marco D. Santambrogio (Politecnico di Milano)
 High-Level Reliability Evaluation of Reconfiguration-Based Fault Tolerance Techniques .202 Tien Thanh Nguyen (CEA/CTREG/DPLOIRE), Mathieu Thevenin (CEA Saclay), Anthony Mouraud (CEA/CTREG/DPLOIRE), Gwenole Corre (CEA Saclay), Olivier Pasquier (IETR, Universit de Nantes), and Sebastien Pillement (IETR, Universit de Nantes)
Dynamic Reconfiguration for Real-Time Automotive Embedded Systems in Fail-Operational Context .206 Florian Oszwald (BMW Group), Jürgen Becker (Karlsruhe Institute of Technology), Philipp Obergfell (BMW Group), and Matthias Traub (BMW

Group)

Session 6: Short Papers 2

FPGA Implementation of Pattern Matching for Industrial Control Systems .210 Peter Rouget (Laboratory of Informatics, Robotics and Microelectronics of Montpellier (LIRMM)), Benoit Badrignans (SECLAB), Pascal Benoit (Laboratory of Informatics, Robotics and Microelectronics of Montpellier (LIRMM)), and Lionel Torres (Laboratory of Informatics, Robotics and Microelectronics of Montpellier (LIRMM))
A Parallel, Energy Efficient Hardware Architecture for the merAligner on FPGA Using Chisel HCL .2.14 Lorenzo Di Tucci (Politecnico di Milano), Davide Conficconi (Politecnico di Milano), Alessandro Comodi (Politecnico di Milano), Steven Hofmeyr (Lawrence Berkeley National Laboratory), David Donofrio (Lawrence Berkeley National Laboratory), and Marco D. Santambrogio (Politecnico di Milano)
Redundant Binary to Two's Complement Converter on FPGAs Through Fabric Aware Scan Based Encoding Approach for Fault Localization Support .2.18 <i>Ayan Palchaudhuri (Indian Institute of Technology Kharagpur) and</i> <i>Anindya Sundar Dhar (Indian Institute of Technology Kharagpur)</i>
An Application-Specific Memory Management Unit for FPGA-SoCs .222 Matthias Goebel (Technische Universitaet Berlin), Ilja Behnke (Technische Universitaet Berlin), Ahmed Elhossini (Al Azhar University), and Ben Juurlink (Technische Universitaet Berlin)

HiCOMB: High Performance Computational Biology

Introduction to HiCOMB 2018 .226. Srinivas Aluru (Georgia Institute of Technology), David Bader (Georgia Institute of Technology), and Paul Medvedev (Penn State)
HiCOMB Keynote 1 227 James Taylor (Johns Hopkins University)
HICOMB Keynote 2 .228 Onur Mutlu (ETH Zurich)
GraphNER: Using Corpus Level Similarities and Graph Propagation for Named Entity Recognition .229 Golnar Sheikhshab (Simon Fraser University), Elizabeth Starks (Canada's Michael Smith Genome Sciences Centre), Aly Karsan (Canada's Michael Smith Genome Sciences Centre), Readman Chiu (Canada's Michael Smith Genome Sciences Centre), Anoop Sarkar (Simon Fraser University), and Inanc Birol (Canada's Michael Smith Genome Sciences Centre)
Modifying HMMER3 to Run Efficiently on the Cori Supercomputer Using OpenMP Tasking .239 William Arndt (Lawrence Berkeley National Laboratory)
Rerooting Trees Increases Opportunities for Concurrent Computation and Results in Markedly Improved Performance for Phylogenetic Inference .247 Daniel L. Ayres (University of Maryland) and Michael P. Cummings (University of Maryland)
Sequence Alignment Through the Looking Glass .257. Raja Appuswamy (EURECOM), Jacques Fellay (EPFL), and Nimisha Chaturvedi (EPFL)

GABB: Graph Algorithms Building Blocks

Introduction to GABB 2018	267
Tim Mattson (Intel)	

Keynote Session

Graph Algorithms in the Language of Linear Algebra: How Did We Get Here, and Where Do We Go Next	t? .268
John R. Gilbert (University of California Santa Barbara)	

Spectral Graph Drawing: Building Blocks and Performance Analysis .269..... Shad Kirmani (eBay Inc) and Kamesh Madduri (Pennsylvania State University)

Session 1: Generating Graphs with Known Properties

Parallel Generation of Large-Scale Random Graphs .2.7.8 Anil Kumar S. Vullikanti (Virginia Tech)
Design, Generation, and Validation of Extreme Scale Power-Law Graphs 279 Jeremy Kepner (Massachusetts Institute of Technology), Siddharth Samsi (Massachusetts Institute of Technology), William Arcand (Massachusetts Institute of Technology), David Bestor (Massachusetts Institute of Technology), Bill Bergeron (Massachusetts Institute of Technology), Tim Davis (Texas A&M), Vijay Gadepally (Massachusetts Institute of Technology), Michael Houle (Massachusetts Institute of Technology), Matthew Hubbell (Massachusetts Institute of Technology), Matthew Hubbell (Massachusetts Institute of Technology), Hayden Jananthan (Massachusetts Institute of Technology), Vanderbilt University), Michael Jones (Massachusetts Institute of Technology), Anna Klein (Massachusetts Institute of Technology), Peter Michaleas (Massachusetts Institute of Technology), Roger Pearce (Lawrence Livermore National Laboratory), Lauren Milechin (Massachusetts Institute of Technology), Julie Mullen (Massachusetts Institute of Technology), Andrew Prout (Massachusetts Institute of Technology), Antonio Rosa (Massachusetts Institute of Technology), Geoff Sanders (Lawrence Livermore National Laboratory), Charles Yee (Massachusetts Institute of Technology), and Albert Reuther (Massachusetts Institute of Technology)
On Large-Scale Graph Generation with Validation of Diverse Triangle Statistics at Edges and Vertices.287

On Large-Scale Graph Generation with Validation of Diverse Triangle Statistics at Edges and Vertices.287..... Geoffrey Sanders (Lawrence Livermore National Laboratory), Roger Pearce (Lawrence Livermore National Laboratory), Timothy La Fond (Lawrence Livermore National Laboratory), and Jeremy Kepner (MIT Lincoln Laboratory)

Session 2: GraphBLAS Implementations

Patterns of GraphBLAS Algorithms: Tales from the Trenches .297 Scott McMillan (Carnegie Mellon University)
Implementing the GraphBLAS C API 298 José E. Moreira (IBM Thomas J. Watson Research Center), Manoj Kumar (IBM Thomas J. Watson Research Center), and William P. Horn (IBM Thomas J. Watson Research Center)
PyGB: GraphBLAS DSL in Python with Dynamic Compilation Into Efficient C++ .3.10 Jesse Chamberlin (University of Washington), Marcin Zalewski (Pacific Northwest National Laboratory), Scott McMillan (Carnegie Mellon University), and Andrew Lumsdaine (University of Washington)

Session 3: Graph Building Blocks Community Meeting

A Survey of Modern Analysis on Graphs: Open Problems	.320
Chris Long (US Department of Defense)	

EduPar: NSF/TCPP W. on Parallel and Distributed Computing Education

Introduction to EduPar 2018 .321 Martina Barnas (Indiana University), Sushil K. Prasad (Georgia State University), and Satish Puri (Marquette University)
EduPar 2018 Keynote .323

Alexandru Iosup (Vrije Universiteit Amsterdam)

EduPar Session 1

 ParallelAR: An Augmented Reality App and Instructional Approach for Learning Parallel Programming Scheduling Concepts .324. Marin Abernethy (The University of Auckland), Oliver Sinnen (The University of Auckland), Joel Adams (Calvin College), Giuseppe De Ruvo (The University of Auckland), and Nasser Giacaman (The University of Auckland), Auckland)
Learning from Optimizing Matrix-Matrix Multiplication .332 Devangi N. Parikh (The University of Texas at Austin), Jianyu Huang (The University of Texas at Austin), Margaret E. Myers (The University of Texas at Austin), and Robert A. van de Geijn (The University of Texas at Austin)
An Entertaining Approach to Parallel Programming Education .340 Emanuel Buzek (Charles University) and Martin Kruliš (Charles University)
Predicting Success in Undergraduate Parallel Programming via Probabilistic Causality Analysis .347 Sunny Raj (University of Central Florida) and Sumit Kumar Jha

(University of Central Florida)

EduPar Session 2

A Comprehensive Course on Big Data for Undergraduate Students .353 Jawwad Ahmed Shamsi (National University of Computer and Emerging Sciences), Syed Zain ul Hassan (National University of Computer and Emerging Sciences), Narmeen Bawany (National University of Computer and Emerging Sciences), and Nausheen Shoaib (National University of Computer and Emerging Sciences)	
Experiences on Teaching Parallel and Distributed Computing for Undergraduates .361 Erik Saule (UNC Charlotte)	
Teaching Parallel Programming with Active Learning .369. Mohammad Amin Kuhail (University of Missouri-Kansas City), Spencer Cook (University of Missouri-Kansas City), Joshua W. Neustrom (University of Missouri-Kansas City), and Praveen Rao (University of Missouri-Kansas City)	
Teaching Big Data and Cloud Computing: A Modular Approach .3.77 Debzani Deb (Winston-Salem State University), Sebastian Cousins (Winston-Salem State University), and Muztaba Fuad (Winston-Salem State University)	

HIPS: High Level Programming Models and Supportive Environments

troduction to HIPS 2018 .384
Karl Fuerlinger (Ludwig-Maximilians-Universität (LMU) München) and
Philip C. Roth (Oak Ridge National Laboratory)
IPS 2018 Keynote 386
Christian Trott (Sandia National Laboratorias)
Christian Tron (Sanata National Laboratories)

Session 1: Tool Support for Parallel Programming Environments

Visualization of Multi-layer I/O Performance in Vampir .387 Hartmut Mix (Technische Universitaet Dresden), Christian Herold (Technische Universitaet Dresden), and Matthias Weber (Technische Universitaet Dresden)
An Operational Semantic Basis for Building an OpenMP Data Race Checker .395 Simone Atzeni (University of Utah) and Ganesh Gopalakrishnan (University of Utah)
Unobtrusive Support for Asynchronous GUI Operations with Java Annotations .405 Mostafa Mehrabi (University of Auckland), Nasser Giacaman (University of Auckland), and Oliver Sinnen (University of Auckland)

Session 2: Distributed Memory and Task-Based Programming

Non-intrusively Avoiding Scaling Problems in and out of MPI Collectives .4.15...... Hongbo Li (University of California, Riverside), Zizhong Chen (University of California, Riverside), Rajiv Gupta (University of California, Riverside), and Min Xie (National University of Defense Technology) Modular Programming of Synchronization and Communication Among Tasks in Parallel Programs .425...... Bernie van Veen (Open University of the Netherlands) and Sung-Shik Jongmans (Open University of the Netherlands / Imperial College London)

Scalable Collectives for Distributed Asynchronous Many-Task Runtimes .436..... Matthew Whitlock (Oklahoma State University), Hemanth Kolla (Sandia National Laboratories), Sean Treichler (Stanford University), Philippe Pébay (NexGen Analytics), and Janine C. Bennett (Sandia National Laboratories)

HPBDC: High-Performance Big Data, Deep Learning, and Cloud Computing

Regular Paper Session 1: High-Performance Data Processing Systems

Improving I/O Performance Through Colocating Interrelated Input Data and Near-Optimal Load Balancing.448 Felix Seibert (Zuse Institute Berlin (ZIB)), Mathias Peters (Humboldt-Universität zu Berlin), and Florian Schintke (Zuse Institute Berlin (ZIB))

How Well do CPU, GPU and Hybrid Graph Processing Frameworks Perform? .458..... Tanuj Kr Aasawat (University of British Columbia), Tahsin Reza (University of British Columbia), and Matei Ripeanu (University of British Columbia)

EASIS: An Optimized Information Service for High Performance Computing Environment .467..... Can Wu (University of Chinese Academy of Sciences), Xiaoning Wang (Chinese Academy of Sciences), Haili Xiao (Chinese Academy of Sciences), Rongqiang Cao (Chinese Academy of Sciences), Yining Zhao (Chinese Academy of Sciences), and Xuebin Chi (University of Chinese Academy of Sciences)

Regular Paper Session 2: High-Performance Data Processing Applications

and Peigang Zou (Shenzhen Institute of Advanced Technology)

atSNPInfrastructure, a Case Study for Searching Billions of Records While Providing Significant Cost Savings over Cloud Providers 497..... Christopher Harrison (University of Wisconsin - Madison), Sundüz Keles (University of Wisconsin - Madison), Rebecca Hudson (University of Wisconsin - Madison), Sunyoung Shin (University of Wisconsin -Madison), and Inês Dutra (Universidade do Porto)

Short Paper Session 1: Data Processing on HPC and Cloud Environments

Improvement of the Log Pattern Extracting Algorithm Using Text Similarity .507..... Yining Zhao (Chinese Academy of Sciences), Xiaodong Wang (University of Chinese Academy of Sciences), Haili Xiao (Chinese Academy of Sciences), and Xuebin Chi (University of Chinese Academy of Sciences)

The Performance Analysis of Cache Architecture Based on Alluxio over Virtualized Infrastructure .5.15...... Xu Chang (University of the Chinese Academy of Sciences) and Li Zha (University of the Chinese Academy of Sciences)

AsHES: Accelerators and Hybrid Exascale Systems

Introduction to AsHES 2018 .520. Sunita Chandrasekaran (University of Delaware), Antonio J. Peña (Barcelona Supercomputing Center), and Min Si (Argonne National Laboratory)
AsHES 2018 Keynote .521
Michael Wolfe (NVIDIA)

Session 1: Runtime Scheduling and Performance Analytics

NVIDIA Tensor Core Programmability, Performance & Precision .522 Stefano Markidis (KTH Royal Institute of Technology), Steven Wei Der Chien (KTH Royal Institute of Technology), Erwin Laure (KTH Royal Institute of Technology), Ivy Bo Peng (Oak Ridge National Laboratory), and Jeffrey S. Vetter (Oak Ridge National Laboratory)
Optimizing an Atomics-Based Reduction Kernel on OpenCL FPGA Platform .532 Zheming Jin (Argonne National Lab) and Hal Finkel (Argonne National Lab)
Leveraging Data-Flow Task Parallelism for Locality-Aware Dynamic Scheduling on Heterogeneous Platforms .540
Osman Seckin Simsek (The University of Manchester), Andi Drebes (The
University of Manchester), and Antoniu Pop (The University of
Manchester)

Session 2: Algorithms and Applications

Tacho: Memory-Scalable Task Parallel Sparse Cholesky Factorization 550 Kyungjoo Kim (Sandia National Laboratories), H. Carter Edwards (Sandia National Laboratories), and Sivasankaran Rajamanickam (Sandia National Laboratories)
Sorting Large Datasets with Heterogeneous CPU/GPU Architectures .560 Michael Gowanlock (Northern Arizona University) and Ben Karsin (University of Hawaii at Manoa)
Improving Performance of Genomic Aligners on Intel Xeon Phi-Based Architectures .5.70 Shaolong Chen (Universitat Autònoma de Barcelona) and Miquel A. (Universitat Autònoma de Barcelona)

Session 3: Emerging Accelerator Architectures

An Initial Characterization of the Emu Chick .579
Eric Hein (Georgia Institute of Technology), Tom Conte (Georgia
Institute of Technology), Jeffrey Young (Georgia Institute of
Technology), Srinivas Eswar (Georgia Institute of Technology), Jiajia
Li (Georgia Institute of Technology), Patrick Lavin (Georgia Institute
of Technology), Richard Vuduc (Georgia Institute of Technology), and
Jason Riedy (Georgia Institute of Technology)
Exploring the Vision Processing Unit as Co-Processor for Inference .589
Exploring the Vision Processing Unit as Co-Processor for Inference .589 Sergio Rivas-Gomez (KTH Royal Institute of Technology), Antonio J.
Exploring the Vision Processing Unit as Co-Processor for Inference .589 Sergio Rivas-Gomez (KTH Royal Institute of Technology), Antonio J. Pena (Barcelona Supercomputing Center (BSC)), David Moloney (Intel
Exploring the Vision Processing Unit as Co-Processor for Inference .589 Sergio Rivas-Gomez (KTH Royal Institute of Technology), Antonio J. Pena (Barcelona Supercomputing Center (BSC)), David Moloney (Intel Ireland Ltd.), Erwin Laure (KTH Royal Institute of Technology), and
Exploring the Vision Processing Unit as Co-Processor for Inference .589 Sergio Rivas-Gomez (KTH Royal Institute of Technology), Antonio J. Pena (Barcelona Supercomputing Center (BSC)), David Moloney (Intel Ireland Ltd.), Erwin Laure (KTH Royal Institute of Technology), and Stefano Markidis (KTH Royal Institute of Technology)

PDCO: Parallel / Distributed Computing and Optimization

Introduction to PDCO 2018 599.
Grégoire Danoy (University of Luxembourg), Didier El Baz (LAAS-CNRS,
France), Vincent Boyer (University of Nuevo Leon, Mexico), and Bernabé
Dorronsoro (Universidad de Cádiz, Špain)

Session 1: Scheduling, Parallel Genetic Algorithms, Genetic Programming

On Integrating Population-Based Metaheuristics with Cooperative Parallelism .601 Jheisson Lopez (University of Antioquia/UNGS), Danny Munera (University of Antioquia), Daniel Diaz (University Paris-1 / CRI), and Salvador Abreu (University of Evora / LISP)
A Competitive Approach for Bi-Level Co-Evolution .609.
Emmanuel Kieffer (SnT Interdisciplinary Centre), Grégoire Danoy
(FSTC-CSC-ILIAS), Pascal Bouvry (FSTC-CSC-ILIAS), and Anass Nagih
(LCOMS Research Unit)

A GPU Parallel Approximation Algorithm for Scheduling Parallel Identical Machines to Minimize Makespan .619 Yuanzhe Li (Wayne State University), Laleh Ghalami (Wayne State University), Loren Schwiebert (Wayne State University), and Daniel Grosu (Wayne State University)
A Survey on Parallel Genetic Algorithms for Shop Scheduling Problems .629 Jia Luo (LAAS-CNRS, Universite de Toulouse, CNRS) and Didier El Baz (LAAS-CNRS, Universite de Toulouse, CNRS)
Session 2: Parallel Distributed Computing Systems and Optimization, Applications
Scalable b-Matching on GPUs .637 Md. Naim (University of Bergen) and Fredrik Manne (University of Bergen)
Automated Analysis of Task-Parallel Execution Behavior Via Artificial Neural Networks .647 Richard Neill (University of Manchester), Andi Drebes (University of Manchester), and Antoniu Pop (University of Manchester)
Data Stream Processing at Network Edges .657 Thanasis Loukopoulos (University of Thessaly), Nikos Tziritas (University of Thessaly), Maria Koziri (University of Thessaly), George Stamoulis (University of Thessaly), Samee U. Khan (National Science Foundation), Cheng-Zhong Xu (Chinese Academy of Sciences), and Albert Y. Zomaya (University of Sydney)
 WA-RRNS: Reliable Data Storage System Based on Multi-cloud .666 Andrei Tchernykh (CICESE Research Center, Ensenada), Mikhail Babenko (North-Caucasus Federal University Stavropol), Vanessa Miranda-López (CICESE Research Center, Ensenada), Alexander Yu. Drozdov (Moscow Institute of Physics and Technology, Moscow), and Arutyun Avetisyan (Institute for System Programming RAS, Moscow)

HPPAC: High-Performance, Power-Aware Computing

Introduction to HPPAC 2018 .67.4 Shuaiwen Leon Song (Pacific Northwest National Lab), Natalie Bates (Energy Efficient HPC Working Group), and Ang Li (Pacific Northwest National Lab)
HPPAC 2018 Keynote .6.7.5. Gregory Koenig (ACM member)
DEEP-Mon: Dynamic and Energy Efficient Power Monitoring for Container-Based Infrastructures .676 Rolando Brondolin (Politecnico di Milano), Tommaso Sardelli (Politecnico di Milano), and Marco D. Santambrogio (Politecnico di Milano)

Energy and Power Aware Job Scheduling and Resource Management: Global Survey — Initial Analysis .685... Matthias Maiterth (Intel Corp. / LMU Munich), Gregory Koenig (EEHPCWG), Kevin Pedretti (Sandia National Laboratories), Siddhartha Jana (Intel), Natalie Bates (EEHPCWG), Andrea Borghesi (University of Bologna), Dave Montoya (Los Alamos National Laboratory), Andrea Bartolini (University of Bologna), and Milos Puzovic (Hartree Centre)

Mitigating Critical Path Decompression Latency in Compressed L1 Data Caches Via Prefetching <u>.694</u>..... Sean Rea (Lakehead University) and Ehsan Atoofian (Lakehead University)

Quality Assessment of GPU Power Profiling Mechanisms .702..... Satyabrata Sen (Oak Ridge National Laboratory), Neena Imam (Oak Ridge National Laboratory), and Chung-Hsing Hsu (Oak Ridge National Laboratory)

System Monitoring with lo2s: Power and Runtime Impact of C-State Transitions .7.12..... Thomas Ilsche (Technische Universität Dresden), Robert Schöne (Technische Universität Dresden), Philipp Joram (Technische Universität Dresden), Mario Bielert (Technische Universität Dresden), and Andreas Gocht (Technische Universität Dresden)

Power and Performance Tradeoff of a Floating-Point Intensive Kernel on OpenCL FPGA Platform .7.16...... Zheming Jin (Argonne National Lab) and Hal Finkel (Argonne National Lab)

Making a Case for Green High-Performance Visualization Via Embedded Graphics Processors .721..... Vignesh Adhinarayanan (Virginia Tech), Bishwajit Dutta (Virginia Tech), and Wu-chun Feng (Virginia Tech)

A Comparison of Power Management Mechanisms: P-States vs. Node-Level Power Cap Control .725...... Kevin Pedretti (Sandia National Laboratories), Ryan E. Grant (Sandia National Laboratories), James H. Laros III (Sandia National Laboratories), Michael Levenhagen (Sandia National Laboratories), Stephen L. Olivier (Sandia National Laboratories), Lee Ward (Sandia National Laboratories), and Andrew J. Younge (Sandia National Laboratories)

APDCM: Advances in Parallel and Distributed Computational Models

Introduction to APDCM 2018 .730 Oscar H. Ibarra (University of California, Santa Barbara), Koji Nakano (Hiroshima University), Akihiro Fujiwara (Kyushu Institute of Technology), and Susumu Matsumae (Saga University)
APDCM 2018 Keynote .732 Yuji Shinano (ZIB)

Session 1: Parallel Computing Models

Survey: Computational Models for Asymmetric Read and Write Costs .733..... Yan Gu (Carnegie Mellon University)

Implementation of Multioperations in Thick Control Flow Processors .744
Martti Forsell (VTT), Jussi Roivainen (VTT), Ville Leppänen
(University of Turku), and Jesper Larsson Träff (Vienna University of
Technology)
A Block Streaming Model for Irregular Applications .753
Anup Zope (Mississippi State University) and Edward Luke (Mississippi
State University)
An Optimal Parallel Algorithm for Computing the Summed Area Table on the GPU .763
Yutaro Emoto (Hiroshima University), Shunji Funasaka (Hiroshima
University), Hiroki Tokura (Hiroshima University), Takumi Honda
(Hiroshima University), Koji Nakano (Hiroshima University), and
Yasuaki Ito (Hiroshima University)

Session 2: Concurrency Models

Barrier Synchronization: Simplified, Generalized, and Solved Without Mutual Exclusion .7.3 Alex Aravind (University of Northern British Columbia)
An Analysis of Multilevel Checkpoint Performance Models .783 Daniel Dauwe (Colorado State University), Sudeep Pasricha (Colorado State University), Anthony A. Maciejewski (Colorado State University), and Howard Jay Siegel (Colorado State University)
Combining Checkpointing and Replication for Reliable Execution of Linear Workflows .793 Anne Benoit (ENS Lyon), Aurélien Cavelan (University of Basel), Florina M. Ciorba (University of Basel), Valentin Le Fèvre (ENS Lyon), and Yves Robert (ENS Lyon)
Optimal Cooperative Checkpointing for Shared High-Performance Computing Platforms .803 Thomas Herault (UTK), Yves Robert (ENS Lyon), Aurélien Bouteiller (UTK), Dorian Arnold (Emory Univ.), Kurt Ferreira (Sandia), George Bosilca (UTK), and Jack Dongarra (UTK)

Session 3: Distributed Computing and Wireless Networks

A Population Protocol for Uniform k-Partition Under Global Fairness .8.13 Hiroto Yasumi (National Institute of Technology), Naoki Kitamura (Nagoya Institute of Technology), Fukuhito Ooshita (Nara Institute of Science and Technology), Taisuke Izumi (Nagoya Institute of Technology), and Michiko Inoue (Nara Institute of Science and Technology)
On the Cost of Cloud-Assistance in Tree-Structured P2P Live Streaming .820 Satoshi Fujita (Hiroshima University)
Mutual Visibility for Robots with Lights Tolerating Light Faults .829 Gokarna Sharma (Kent State University)

Joint Cooperative Protocols and Distributed Beamforming Design with Efficient Secondary User Selection for Multi-hop Cognitive Radio Networks .837.....

Wei Chen (Tennessee State University), Liang Hong (Tennessee State University), Sudeep Bhattarai (Virginia Tech), Tony Sanchez (Tennessee State University), Ebholo Ijieh (Tennessee State University), Stacie Severyn (U.S. Air Force Research Laboratory), and Leonard Lightfoot (U.S. Air Force Research Laboratory)

A Novel Handover Control Strategy Combined with Multi-hop Routing in LEO Satellite Networks .845..... Chaofan Duan (National University of Defense Technology), Jing Feng (National University of Defense Technology), Haotian Chang (National University of Defense Technology), Bin Song (National University of Defense Technology), and Zhikang Xu (National University of Defense Technology)

ParLearning: Parallel and Distributed Computing for Large-Scale Machine Learning and Big Data Analytics

Introduction to ParLearning 2018 852. Henri Bal (Vrije Universiteit, The Netherlands), Arindam Pal (TCS Research and Innovation, India), Azalia Mirhoseini (Google Brain, USA), and Thomas Parnell (IBM Research – Zurich, Switzerland)
ParLearning 2018 Invited Talk 1 .854 Abhinav Vishnu (AMD Research)
ParLearning 2018 Invited Talk 2 .855 Azalia Mirhoseini (Google Brain)
ParLearning 2018 Invited Talk 3 .856 Thomas Parnell (IBM Research, Zürich)
Near-Optimal Straggler Mitigation for Distributed Gradient Methods .857 Songze Li (University of Southern California), Seyed Mohammadreza Mousavi Kalan (University of Southern California), A. Salman Avestimehr (University of Southern California), and Mahdi Soltanolkotabi (University of Southern California)
Streaming Tiles: Flexible Implementation of Convolution Neural Networks Inference on Manycore Architectures .867 Nesma M. Rezk (Halmstad University), Madhura Purnaprajna (Amrita Vishwa Vidyapeetham), and Zain Ul-Abdin (Halmstad University)
Parallel Huge Matrix Multiplication on a Cluster with GPGPU Accelerators .8.7.7 Seungyo Ryu (Korea University) and Dongseung Kim (Korea University)
A Study of Clustering Techniques and Hierarchical Matrix Formats for Kernel Ridge Regression .883 Elizaveta Rebrova (University of Michigan), Gustavo Chávez (Lawrence Berkeley National Laboratory), Yang Liu (Lawrence Berkeley National Laboratory), Pieter Ghysels (Lawrence Berkeley National Laboratory), and Xiaoye Sherry Li (Lawrence Berkeley National Laboratory)

CHIUW: Chapel Implementers and Users Workshop

ntroduction to CHIUW 2018 .893. Michael Ferguson (Cray Inc.), Nikhil Padmanabhan (Yale University), and Brad Chamberlain (Cray Inc.)	•••
CHIUW 2018 Keynote .895 Katherine Yelick (University of California at Berkeley)	••
Session 1: Introduction	

Session 2: Applications of Chapel

Parallel Sparse Tensor Decomposition in Chapel .896 Thomas B. Rolinger (University of Maryland), Tyler A. Simon (University of Maryland), and Christopher D. Krieger (University of Maryland)	
Iterator-Based Optimization of Imperfectly-Nested Loops .906 Daniel Feshbach (Haverford College), Mary Glaser (Haverford College), Michelle Strout (University of Arizona), and David G. Wonnacott (Haverford College)	
Investigating Data Layout Transformations in Chapel .9.15 Apan Qasem (Texas State University), Ashwin M. Aji (AMD Research), and Michael L. Chu (AMD Research)	 !

Session 3: Chapel Design and Evolution

RCUArray: An RCU-Like Parallel-Safe Distributed Resizable Array .925..... Louis Jenkins (Bloomsburg University)

Session 4: Chapel Performance

Session 5: Tools

Purity: An Integrated, Fine-Grain, Data-Centric, Communication Profiler for the Chapel Language .934...... Richard B. Johnson (University of Maryland) and Jeffrey J. Hollingsworth (University of Maryland)

PDSEC: Parallel and Distributed Scientific and Engineering Comp

Introduction to PDSEC 2018 and Keynotes .943.... Peter Strazdins (The Australian National University), Keita Teranishi (Sandia National Laboratories), Raphaël Couturier (Univ. Bourgogne Franche-Comté, France), Joseph Antony (National Computational Infrastructure, Australia), Thomas Rauber (University of Bayreuth, Germany), Gudula Rünger (Chemnitz University of Technology, Germany), and Laurence T. Yang (St. Francis Xavier University, Canada)

Session 1: Best Paper

DM-HEOM: A Portable and Scalable Solver-Framework for the Hierarchical Equations of Motion .947	
Matthias Noack (Zuse Institute Berlin), Alexander Reinefeld (Zuse	
Institute Berlin), Tobias Kramer (Zuse Institute Berlin), and Thomas	
Steinke (Zuse Institute Berlin)	

Session 2: Parallel Computational Techniques

Optimization of Reordering Procedures in HOTRG for Distributed Parallel Computing .957 Haruka Yamada (University of Tsukuba), Akira Imakura (University of Tsukuba), Toshiyuki Imamura (RIKEN AICS), and Tetsuya Sakurai (University of Tsukuba)
Energy and Performance Improvement of Parallel ODE Solvers by Application-Specific Program Transformations .967 Thomas Rauber (University Bayreuth) and Gudula Rünger (Chemnitz University of Technology)
The Scalability of Embedded Structured Grids and Unstructured Grids in Large Scale Ice Sheet Modeling on Distributed Memory Parallel Computers 9.7 Phillip M. Dickens (University of Maine), Christopher Dufour (University of Maine), and James Fastook (University of Maine)
 TNT: A Solver for Large Dense Least-Squares Problems that Takes Conjugate Gradient from Bad in Theory, to Good in Practice .987 Joseph M. Myre (University of St. Thomas, Saint Paul), Erich Frahm (University of Minnesota, Twin Cities), David J. Lilja (University of Minnesota, Twin Cities), and Martin O. Saar (ETH-Zurich)

Session 3: Invited Talks/Systems

An Energy-Efficient Asymmetric Multi-Processor for HPC Virtualization .996..... Chung Leee (The Australian National University) and Peter Strazdins (The Australian National Universit)

Session 4: Accelerators

reliminary Port and Evaluation of the Uintah AMT Runtime on Sunway TaihuLight 1006
Zhang Yang (Institute of Applied Physics and Computatinal
Mathematics), Damodar Sahasrabudhe (Scientific Computing and Imaging
Institute), Alan Humphrey (Scientific Computing and Imaging
Institute), and Martin Berzins (Scientific Computing and Imaging
Institute)
proving CADNA Performance on GPUs .1016
Pacôme Eberhart (Sorbonne Université), Baptiste Landreau (Sorbonne
Université), Julien Brajard (Sorbonne Université), Pierre Fortin
(Sorbonne Université), and Fabienne Jézéquel (Sorbonne Université)
Droving CADNA Performance on GPUs .1016. Pacôme Eberhart (Sorbonne Université), Baptiste Landreau (Sorbonne Université), Julien Brajard (Sorbonne Université), Pierre Fortin (Sorbonne Université), and Fabienne Jézéquel (Sorbonne Université)

Evaluation of MD5Hash Kernel on OpenCL FPGA Platform .1026
Zheming Jin (Argonne National Lab) and Hal Finkel (Argonne National
Lab)
Performance Optimization of Fully Anisotropic Elastic Wave Propagation on 2nd Generation Intel® Xeon
Phi(TM) Processors .1033.
Albert Farres (Universitat Politècnica de Catalunya / Barcelona
Supercomputing Center), Claudia Rosas (Barcelona Supercomputing
Center), Mauricio Hanzich (Barcelona Supercomputing Center), Alejandro
Duran (Intel Corporation Iberia), and Charles Yount (Intel
Corporation)

JSSPP: Job Scheduling Strategies for Parallel Processing

iWAPT: International Workshop on Automatic Performance Tunings

Introduction to iWAPT 2018 .1047..... Osni Marques (LBNL, USA), Reiji Suda (The University of Tokyo), Jakub Kurzak (University of Tennessee), and Akihiro Fujii (Kogakuin University, Japan)

Session 1: Machine Learning

iWAPT 2018 Invited Speaker 1 .1048 Sarah Knepper (Intel Corporation)
Use of Code Structural Features for Machine Learning to Predict Effective Optimizations .1049

Yuki Kawarabatake (Tohoku University), Mulya Agung (Tohoku University), Kazuhiko Komatsu (Tohoku University), Ryusuke Egawa (Tohoku University), and Hiroyuki Takizawa (Tohoku University)

Effective Machine Learning Based Format Selection and Performance Modeling for SpMV on GPUs .1056... Israt Nisa (The Ohio State University), Charles Siegel (Pacific Northwest National Laboratory), Aravind Sukumaran Rajam (The Ohio State University), Abhinav Vishnu (Pacific Northwest National Laboratory), and P. Sadayappan (The Ohio State University)

Session 2: AT Techniques

Tensile: Auto-Tuning GEMM GPU Assembly for All Problem Sizes .1066..... David E. Tanner (Advanced Micro Devices)

GreedyTalents: An Energy-Aware Auto-Tuning Method for Many-Core Processor .1076
Timothy M. Platt (Clarkson University), Zhiliu Yang (Clarkson
University), and Chen Liu (Clarkson University)
Auto-Tuning for the Era of Relatively High Bandwidth Memory Architectures: A Discussion Based on an
FDM Application .1084
Takahiro Katagiri (Nagoya University)
Threaded Accurate Matrix-Matrix Multiplications with Sparse Matrix-Vector Multiplications .1093
Shuntaro Ichimura (Nagoya University, Japan), Takahiro Katagiri
(Nagoya University, Japan), Katsuhisa Ozaki (Shibaura Institute of
Technology, Japan), Takeshi Ogita (Tokyo Women's Christian
University), and Toru Nagai (Nagoya University, Japan)

Session 3: Linear Algebra

iWAPT 2018 Invited Speaker 2 .1103 David E. Tanner (Advanced Micro Devices, Inc.)
 Algebraic Multigrid Solver Using Coarse Grid Aggregation with Independent Aggregation .1.104 Naoya Nomura (The University of Tokyo), Akihiro Fujii (Kogakuin University), Teruo Tanaka (Kogakuin University), Osni Marques (Lawrence Berkeley National Laboratory), and Kengo Nakajima (The University of Tokyo)
A Case Study on Modeling the Performance of Dense Matrix Computation: Tridiagonalization in the EigenExa Eigensolver on the K Computer .1.113 Takeshi Fukaya (Hokkaido University / RIKEN Center for Computational Science), Toshiyuki Imamura (RIKEN Center for Computational Science), and Yusaku Yamamoto (The University of Electro-Communications)

Session 4: AT Methodologies

AutoTuneTMP: Auto-Tuning in C++ With Runtime Template Metaprogramming .1.123
David Pfander (University of Stuttgart), Malte Brunn (University of
Stuttgart), and Dirk Pflüger (University of Stuttgart)
Methodology for Adaptive Active Message Coalescing in Task Based Runtime Systems 1133

Methodology for Adaptive Active Message Coalescing in Task Based Runtime Systems 1.1.3.3..... Bibek Wagle (Louisiana State University), Samuel Kellar (Louisiana State University), Adrian Serio (Louisiana State University), and Hartmut Kaiser (Louisiana State University)

ParSocial: Parallel and Distributed Processing for Computational Social Systems

Introduction to ParSocial 2018 .1141
Eunice E. Santos (Illinois Institute of Technology) and John Korah
(Illinois Institute of Technology)
ParSocial 2018 Keynote 1142
V. S. Subrahmanian (Dartmouth College)

Using Activity Patterns to Place Electric Vehicle Charging Stations in Urban Regions .1.14.3..... Anamitra Pal (Arizona State University), Pavan Rangudu (Virginia Tech), S. S. Ravi (Virginia Tech and SUNY Albany), and Anil K. Vullikanti (Virginia Tech)

Handling Vertex Deletions in Memory Scalable Anytime Anywhere Algorithms for Large and Dynamic
Social Networks .1153..... *Eunice E. Santos (Illinois Institute of Technology), John Korah*(Illinois Institute of Technology), and Vairavan Murugappan (Illinois Institute of Technology)

Integrating Cyber Security and Data Science for Social Media: A Position Paper .1.163..... Bhavani Thuraisingham (The University of Texas at Dallas), Murat Kantarcioglu (The University of Texas at Dallas), and Latifur Khan (The University of Texas at Dallas)

GraML: Graph Algorithms and Machine Learning

Introduction to GraML 2018 .1.166. Antonino Tumeo (PNNL), Mahantesh Halappanavar (PNNL), John Feo (PNNL), Assefaw Gebremedhin (Washington State University), and Abhinav Vishnu (AMD Research)
GraML 2018 Keynote .1.168 Nesreen Ahmed (Intel)
Classification and Anomaly Detection in Traffic Patterns of New York City Taxis: A Case Study in Compound Analytics .1.169 Ronald D. Hagan (BAE Systems), Charles A. Phillips (University of Tennessee), Michael A. Langston (University of Tennessee), and Bradley J. Rhodes (BAE Systems)
V2V: Vector Embedding of a Graph and Applications .1175 Trong Duc Nguyen (Iowa State University) and Srikanta Tirthapura (Iowa State University)
Network Similarity Prediction in Time-Evolving Graphs: A Machine Learning Approach .1184 Keyvan Sasani (Washington State University), Mohammad Hossein Namaki (Washington State University), and Assefaw H. Gebremedhin (Washington State University)
Neural Networks and Graph Algorithms with Next-Generation Processors .1.194 Kathleen E. Hamilton (Oak Ridge National Laboratory), Catherine D. Schuman (Oak Ridge National Laboratory), Steven R. Young (Oak Ridge National Laboratory), Neena Imam (Oak Ridge National Laboratory), and Travis S. Humble (Oak Ridge National Laboratory)

CEBDA: Convergence of Extreme Scale Computing and Big Data Analysis

Introduction to CEBDA 2018 .1204..... Shadi Ibrahim (Inria, France), Manish Parashar (Rutgers University), Anna Queralt (Barcelona Supercomputing Center, Spain), and Domenico Talia (University of Calabria, Italy)

CEBDA 2018 Keynote .1205 Franck Cappello (Argonne National Laboratory,)
Data-Locality Aware Dynamic Schedulers for Independent Tasks with Replicated Inputs .1206 Olivier Beaumont (INRIA), Thomas Lambert (University of Manchester), Loris Marchal (CNRS), and Bastien Thomas (ENS Rennes)
Transferring Data from High-Performance Simulations to Extreme Scale Analysis Applications in Deal Time, 1214
The map Marrie and (University of St. The map) Silvie Direct (Augenne)
Thomas Marrinan (University of St. Thomas), Sitvio Rizzi (Argonne
National Laboratory), Joseph Insley (Argonne National Laboratory),
Brian Toonen (Argonne National Laboratory), William Allcock (Argonne
National Laboratory), and Michael E. Papka (Argonne National
Laboratory)
Towards a TRansparent I/O Solution .1221
Fotios Nikolaidis (French Alternative Energies and Atomic Energy
Commission), Nick Kossifidis (Foundation Of Research and Technology
Hellas), Thomas Leibovici (French Alternative Energies and Atomic
Energy Commission), and Soraya Zertal (University of Versailles)

MPP: Parallel Programming Model: Special Edition on Edge/Fog/In-Situ Computing

Introduction to MPP 2018 .1229 Leandro A. J. Marzulo (Universidade do Estado do Rio de Janeiro), Felipe M. G. França (Universidade Federal do Rio de Janeiro), Cristiana Bentes (Universidade do Estado do Rio de Janeiro), and Gabriele Mencagli (University of Pisa, Italy)
MPP 2018 Keynote .123.1
Vladimir Alves (NGD Systems, USA) and Jae Young Do (Microsoft
Research, USA)
Invited Paper: How Future Buildings Could Redefine Distributed Computing .1232
Yanik Ngoko (Qarnot Computing), Nicolas Saintherant (Qarnot
Computing), Christophe Cerin (University of Paris 13), and Denis
Trystram (University of Grenoble Alpes)

Session 1: Applications

A Smart Disk for In-Situ Face Recognition .1241
Victor C. Ferreira (Federal University of Rio de Janeiro), Alexandre
S. Nery (University of Brasilia), and Felipe M.G. Franca (Federal
University of Rio de Janeiro)
A DVND Local Search Implemented on a Dataflow Architecture for the Minimum Latency Problem .1250
Rodolfo Pereira Araujo (Universidade do Estado do Rio de Janeiro),
Igor Machado Coelho (Universidade do Estado do Rio de Janeiro), and
Leandro A. J. Marzulo (Universidade do Estado do Rio de Janeiro)

Session 2: Platforms and Tools

CompStor: An In-storage Computation Platform for Scalable Distributed Processing .1260...... Mahdi Torabzadehkashi (University of California, Irvine), Siavash Rezaei (University of California, Irvine), Vladimir Alves (NGD Systems), and Nader Bagherzadeh (University of California, Irvine)

Fog-Assisted Translation: Towards Efficient Software Emulation on Heterogeneous IoT Devices .1268...... Vanderson Martins do Rosario (University of Campinas, UNICAMP), Flávia Pisani (University of Campinas, UNICAMP), Alexandre Rodrigues Gomes (University of Campinas, UNICAMP), and Edson Borin (University of Campinas, UNICAMP)

PMAW: Programming Models and Algorithms Workshop

Introduction to PMAW 2018 .1278..... Martin Kong (Brookhaven National Laboratory) and Zoran Budimlic (Rice University)

ROME: Runtime and Operating Systems for the Many-Core Era

Introduction to ROME 2018 .1279. Stefan Lankes (RWTH Aachen University, Germany), Carsten Clauss (ParTec Cluster Competence Center GmbH, Germany), and Jens Breitbart (Robert Bosch GmbH, Germany)
ROME 2018 Keynote .1281 Sang-Hoon Kim (Virginia Tech)
ROME 2018 Invited Talk .1282 Karl Fuerlinger (Ludwig-Maximilians-University Munich)
Memory Footprint of Locality Information on Many-Core Platforms .1283 Brice Goglin (Inria)
Diagnosing Performance Fluctuations of High-Throughput Software for Multi-core CPUs .129.3 Soramichi Akiyama (National Institute of Advanced Industrial Science and Technology (AIST)), Takahiro Hirofuchi (National Institute of Advanced Industrial Science and Technology (AIST)), and Ryousei Takano (National Institute of Advanced Industrial Science and Technology (AIST))
Parallelizing MPI Using Tasks for Hybrid Programming Models .1303 Surabhi Jain (Intel Corporation), Gengbin Zheng (Intel Corporation), Maria Garzaran (Intel Corporation), James H. Cownie (Intel Corporation), Taru Doodi (Intel Corporation), and Terry L. Wilmarth (Intel Corporation)
A Study of Network Quality of Service in Many-Core MPI Applications .1313 Lee Savoie (The University of Arizona), David K. Lowenthal (The University of Arizona), Bronis R. de Supinski (Lawrence Livermore National Laboratory), and Kathryn Mohror (Lawrence Livermore National Laboratory)

Author Index