

2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW 2016)

**Chicago, Illinois, USA
23-27 May 2016**

Pages 1-928



**IEEE Catalog Number: CFP1651J-POD
ISBN: 978-1-5090-3683-7**

**Copyright © 2016 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 publication is a representation of what appears in the IEEE Digital Libraries. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP1651J-POD
ISBN (Print-On-Demand):	978-1-5090-3683-7
ISBN (Online):	978-1-5090-3682-0

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

2016 IEEE International Parallel and Distributed Processing Symposium Workshops

IPDPSW 2016

Table of Contents

Message from the General Chair	xxix
Message from the Workshops Chair.....	xxxi

Workshop 1-HCW - Heterogeneity in Computing Workshop

HCW Introduction	1
<i>Denis Trystram and Erik Saule</i>	
Message from the HCW Steering Committee Chair	3
<i>Behrooz Shirazi</i>	
Message from the HCW General Chair	4
<i>Denis Trystram</i>	
Message from the HCW Program Committee Chair	5
<i>Erik Saule</i>	
HCW 2016 Keynote Talk	6
<i>Mahmut Kandemir</i>	

Session 1: Heterogeneity in the Cloud

Towards a Green, QoS-Enabled Heterogeneous Cloud Infrastructure	7
<i>Julio Proaño, Carmen Carrión, and M. Blanca Caminero</i>	
Predicting Job Completion Time in Heterogeneous MapReduce Environments	17
<i>Rekha Singhal and Abhishek Verma</i>	
Minimizing Rental Cost for Multiple Recipe Applications in the Cloud	28
<i>Fouad Hanna, Loris Marchal, Jean-Marc Nicod, Laurent Philippe, Veronika Rehn-Sonigo, and Hala Sabbah</i>	

Session 2: Heterogeneity in Single Node Systems

Providing Fairness in Heterogeneous Multicores with a Predictive, Adaptive Scheduler	38
<i>Saeid Barati and Hank Hoffmann</i>	
clCaffe: OpenCL Accelerated Caffe for Convolutional Neural Networks	50
<i>Jeremy Bottleson, SungYe Kim, Jeff Andrews, Preeti Bindu, Deepak N. Murthy, and Jingyi Jin</i>	
Parallel Graph Partitioning on a CPU-GPU Architecture	58
<i>Bahareh Goodarzi, Martin Burtscher, and Dhrubajyoti Goswami</i>	

Session 3: Heterogeneity and Energy

Dynamic Resource Management for Parallel Tasks in an Oversubscribed Energy-Constrained Heterogeneous Environment	67
<i>Dylan Machovec, Bhavesh Khemka, Sudeep Pasricha, Anthony A. Maciejewski, Howard Jay Siegel, Gregory A. Koenig, Michael Wright, Marcia Hilton, Rejendra Rambharos, and Neena Imam</i>	
Analyzing the Energy Efficiency of the Fast Multipole Method Using a DVFS-Aware Energy Model	79
<i>Jee W. Choi and Richard W. Vuduc</i>	
Evaluation of Emerging Energy-Efficient Heterogeneous Computing Platforms for Biomolecular and Cellular Simulation Workloads	89
<i>John E. Stone, Michael J. Hallock, James C. Phillips, Joseph R. Peterson, Zaida Luthey-Schulten, and Klaus Schulten</i>	

Workshop 2-Raw - Reconfigurable Architectures Workshop

RAW Introduction and Committees	101
<i>Marco D. Santambrogio, Ramachandran Vaidyanathan, Diana Goehringer, and Steve Wilton</i>	
RAW 2016 Keynotes	103
<i>Peter Hofstee, Patrick Lysaght, and Dirk van den Heuvel</i>	

Session 1: Application Mapping and Design Space Exploration

Clustering and Mapping Algorithm for Application Distribution on a Scalable FPGA Cluster	105
<i>Lester Kalms and Diana Göhringer</i>	
A Fast and Accurate Cost Model for FPGA Design Space Exploration in HPC Applications	114
<i>Syed Waqar Nabi and Wim Vanderbauwhede</i>	

Latency, Power, and Security Optimization in Distributed Reconfigurable Embedded Systems	124
<i>Hyunsuk Nam and Roman Lysecky</i>	

Session 2: Applications

A Reconfigurable Fixed-Point Architecture for Adaptive Beamforming	132
<i>Daniel Llamocca and Daniel N. Aloisio</i>	
Parameterizable FPGA-Based Kalman Filter Coprocessor Using Piecewise	
Affine Modeling	139
<i>Aaron Mills, Phillip H. Jones, and Joseph Zambreno</i>	
High Throughput Large Scale Sorting on a CPU-FPGA Heterogeneous	
Platform	148
<i>Chi Zhang, Ren Chen, and Viktor Prasanna</i>	
An FPGA Architecture to Accelerate the Burrows Wheeler Transform by Using	
a Linear Sorter	156
<i>Juan Andrés Pérez-Celis, José Martínez-Carranza, Alicia Morales-Reyes, Claudia Feregrino-Uribe, and René Cumplido</i>	

Session 3: Processor Architectures

A 16-Bit Reconfigurable Encryption Processor for π-Cipher	162
<i>Mohamed El-Hadeby, Hristina Mihajlovska, Danilo Gligoroski, Amit Kulkarni, Dirk Stroobandt, and Kevin Skadron</i>	
Dynamic Self-Reconfiguration of a MIPS-Based Soft-Processor Architecture	172
<i>S. Nolting, G. Payá-Vayá, F. Giesemann, H. Blume, S. Niemann, and C. Müller-Schloer</i>	
An Application-Specific Instruction Set Processor for Power Quality Monitoring	181
<i>Steffen Vaas, Marc Reichenbach, and Dietmar Fey</i>	

Session 4: Scheduler and Runtime Systems

Resource-Efficient Scheduling for Partially-Reconfigurable FPGA-Based	
Systems	189
<i>Andrea Purgato, Davide Tantillo, Marco Rabozzi, Donatella Sciuto, and Marco D. Santambrogio</i>	
Scheduler for Inhomogeneous and Irregular CGRAs with Support for Complex	
Control Flow	198
<i>Tajas Ruschke, Lukas Johannes Jung, Dennis Wolf, and Christian Hochberger</i>	
LinROS: A Linux-Based Runtime System for Reconfigurable MPSoCs	208
<i>Jens Rettkowski, Philipp Wehner, Evgeni Cutishev, and Diana Göhringer</i>	

Session 5: High Level Synthesis and Object-Oriented Programming

On the Automation of High Level Synthesis of Convolutional Neural Networks	217
<i>Emanuele Del Sozzo, Andrea Solazzo, Antonio Miele, and Marco D. Santambrogio</i>	
Scala-Based Domain-Specific Language for Creating Accelerator-Based SoCs	225
<i>Gianluca C. Durelli, Fabrizio Spada, Christian Pilato, and Marco D. Santambrogio</i>	
OOGen: An Automated Generation Tool for Custom MPSoC Architectures	
Based on Object-Oriented Programming Methods	233
<i>Hongyuan Ding, Sen Ma, Miaoqing Huang, and David Andrews</i>	

Short Papers

A Hardware/Software Co-Design Approach for Control Applications with Static Real-Time Reallocation	241
<i>Benedikt Janßen, Moataz Naserdin, and Michael Hübner</i>	
On How to Improve FPGA-Based Systems Design Productivity via SDAccel	247
<i>Giulia Guidi, Enrico Reggiani, Lorenzo Di Tucci, Gianluca Durelli, Michaela Blott, and Marco D. Santambrogio</i>	
A Rapid Prototyping Method to Reduce the Design Time in Commercial High-Level Synthesis Tools	253
<i>Jones Y. Mori, André Werner, Florian Fricke, and Michael Hüebner</i>	
ARTNoCs: An Evaluation Framework for Hardware Architectures of Real-Time NoCs	259
<i>Salma Hesham, Diana Göhringer, and Mohamed Abd El Ghany</i>	
A Fully Parameterized Virtual Coarse Grained Reconfigurable Array for High Performance Computing Applications	265
<i>Amit Kulkarni, Elias Vasteenkiste, Dirk Stroobandt, Andreas Brokalakis, and Antonios Nikitakis</i>	
Assessing Multi-task Placement Algorithms in RCUs	271
<i>Anita Tino and Kaamran Raahemifar</i>	
Efficient Hardware Debugging Using Parameterized FPGA Reconfiguration	277
<i>Alexandra Kourfali and Dirk Stroobandt</i>	
Enabling Dynamic Reconfiguration of Numerical Methods for the Robotic Motion Control Task	283
<i>Fynn Schwiegelshohn, Florian Kästner, and Michael Hübner</i>	
Hardware Architectures for Frequent Itemset Mining Based on Equivalence Classes Partitioning	289
<i>Martin Letras, Raudel Hernández-León, and Rene Cumplido</i>	
Parallel Protein Identification Using an FPGA-Based Solution	295
<i>Fabiola Casasopa, Gea Bianchi, Gianluca C. Durelli, and Marco D. Santambrogio</i>	

Face Recognition Using Local Binary Patterns Histograms (LBPH) on an FPGA-Based System on Chip (SoC)	300
<i>Nikolaos Stekas and Dirk van den Heuvel</i>	

Workshop 3-HIPS - High-Level Parallel Programming Models and Supportive Environments

HIPS Introduction and Committees	305
<i>David Boehme and Xu Liu</i>	
HIPS 2016 Keynote	307
<i>Tim Mattson</i>	

Session 1: Debugging and Optimization

Detecting Anomalies in Concurrent Programs Based on Dynamic Control Flow Changes	308
<i>Faheem Ullah and Thomas R. Gross</i>	
Controlling the Memory Subscription of Distributed Applications with a Task-Based Runtime System	318
<i>Marc Sergent, David Goudin, Samuel Thibault, and Olivier Aumage</i>	
Reducing Redundant Search in Parallel Graph Mining Using Exceptions	328
<i>Shingo Okuno, Tasuku Hiraishi, Hiroshi Nakashima, Masahiro Yasugi, and Jun Sese</i>	

Session 2: Heterogeneous Computing

Evaluating OpenMP 4.0's Effectiveness as a Heterogeneous Parallel Programming Model	338
<i>Matt Martineau, Simon McIntosh-Smith, and Wayne Gaudin</i>	
Employing Compression Solutions under OpenACC	348
<i>Ebad Salehi, Ahmad Lashgar, and Amirali Baniasadi</i>	
CAFe: Coarray Fortran Extensions for Heterogeneous Computing	357
<i>Craig Rasmussen, Matthew Sottile, Soren Rasmussen, Dan Nagle, and William Dumas</i>	

Session 3: Parallel Algorithms and Systems

Embedding Concurrent Generators	366
<i>Peter Mills and Clinton Jeffery</i>	
The Case for Binary Rewriting at Runtime for Efficient Implementation of High-Level Programming Models in HPC	376
<i>Josef Weidendorfer and Jens Breitbart</i>	

PTRAM: A Parallel Topology-and Routing-Aware Mapping Framework for Large-Scale HPC Systems	386
<i>Seyed H. Mirsadeghi and Ahmad Afsahi</i>	

A Comparison of High-Level Programming Choices for Incomplete Sparse Factorization Across Different Architectures	397
<i>Joshua Dennis Booth, Kyungjoo Kim, and Sivasankaran Rajamanickam</i>	

Workshop 4-HiCOMB - High Performance Computational Biology

HiCOMB Introduction and Committees	407
<i>Srinivas Aluru, David A. Bader, Ananth Kalyanaraman, and Jaroslaw Zola</i>	

Session I

The Divisible Load Balance Problem with Shared Cost and Its Application to Phylogenetic Inference	408
<i>Constantin Scholl, Kassian Kobert, Tomáš Flouri, and Alexandros Stamatakis</i>	

Efficient Computation of Linkage Disequilibria as Dense Linear Algebra Operations	418
<i>Nikolaos Alachiotis, Thom Popovici, and Tze Meng Low</i>	

Improving Reaction Kernel Performance in Lattice Microbes: Particle-Wise Propensities and Run-Time Generated Code	428
<i>Michael J. Hallock and Zaida Luthey-Schulten</i>	

Session II

SparkScore: Leveraging Apache Spark for Distributed Genomic Inference	435
<i>Amir Bahmani, Alexander B. Sibley, Mahmoud Parsian, Kouros Owzar, and Frank Mueller</i>	

A Scalable Pipeline for Transcriptome Profiling Tasks with On-Demand Computing Clouds	443
<i>Shayan Shams, Nayong Kim, Xiandong Meng, Ming Tai Ha, Shantenu Jha, Zhong Wang, and Joohyun Kim</i>	

A Memory and Time Scalable Parallelization of the Reptile Error-Correction Code	453
<i>Vipin Sachdeva, Srinivas Aluru, and David A. Bader</i>	

Session III

Real-Time Agent-Based Modeling Simulation with in-Situ Visualization of Complex Biological Systems: A Case Study on Vocal Fold Inflammation and Healing	463
---	-----

*Nuttiiya Seekhao, Caroline Shung, Joseph Jaja, Luc Mongeau,
and Nicole Y. K. Li-Jessen*

A Novel Associative Memory Based Architecture for Sequence Alignment	473
--	-----

*M. Ali Mirzaei, Francesco Crescioli, Sébastien Viret, William Tromeur,
Giovanni Calderini, Giovanni Marchiori, Guillaume Baulieu, and Geoffrey Galbit*

Workshop 5-APDCM - Advances in Parallel and Distributed Computational Models

APDCM Introduction and Committees	479
---	-----

Oscar H. Ibarra, Koji Nakano, Akihiro Fujiwara, and Susumu Matsumae

Session 1: Graph Algorithms

Stable Matching Beyond Bipartite Graphs	480
---	-----

Jie Wu

Fine-Grained Task Migration for Graph Algorithms Using Processing in Memory	489
--	-----

Paula Aguilera, Dong Ping Zhang, Nam Sung Kim, and Nuwan Jayasena

Session 2: Wireless Networks and Distributed Computing

Cross-Layered Security Approach with Compromised Nodes Detection in Cooperative Sensor Networks	499
--	-----

Wei Chen, Liang Hong, Sachin Shetty, Dan Lo, and Reginald Cooper

Model Checking Techniques for State Space Reduction in MANET Protocol Verification	509
---	-----

Hideharu Kojima, Yuta Nagashima, and Tatsuhiko Tsuchiya

New Biology Inspired Anonymous Distributed Algorithms to Compute Dominating and Total Dominating Sets in Network Graphs	517
--	-----

Feng Luo and Pradip K Srimani

Session 3: Distributed Computing and Models

Performance of Causal Consistency Algorithms for Partially Replicated Systems	525
--	-----

Ta-Yuan Hsu and Ajay D. Kshemkalyani

Performance Analysis of an I/O-Intensive Workflow Executing on Google Cloud and Amazon Web Services	535
<i>Hassan Nawaz, Gideon Juve, Rafael Ferreira Da Silva, and Ewa Deelman</i>	
Performance Models for Split-Execution Computing Systems	545
<i>Travis S. Humble, Alexander J. McCaskey, Jonathan Schrock, Hadayat Seddiqi, Keith A. Britt, and Neena Imam</i>	
A Model for Entropy of Parallel Execution	555
<i>Ernesto Gomez, Keith E. Schubert, and Ritchie Cai</i>	

Session 4: Parallel Computing

FFT on XMT: Case Study of a Bandwidth-Intensive Regular Algorithm on a Highly-Parallel Many Core	561
<i>James Edwards and Uzi Vishkin</i>	
Parallelization of Recursive Preorder Traversal Based on Building and Winding Call Stacks	570
<i>Makoto Nakayama, Kenichi Yamazaki, and Satoshi Tanaka</i>	
A GPU Based Maximum Common Subgraph Algorithm for Drug Discovery Applications	580
<i>P. B. Jayaraj, K. Rahamathulla, and G. Gopakumar</i>	
Bitwise Parallel Bulk Computation on the GPU, with Application to the CKY Parsing for Context-Free Grammars	589
<i>Toru Fujita, Koji Nakano, and Yasuaki Ito</i>	
An Efficient Implementation of LZW Decompression in the FPGA	599
<i>Xin Zhou, Yasuaki Ito, and Koji Nakano</i>	

Workshop 6-ASHES - Accelerators and Hybrid Exascale Systems

AsHES Introduction and Committees	608
<i>James Dinan</i>	
AsHES 2016 Keynote	610
<i>Wen-mei Hwu</i>	

Session 1: Programming Models and Tools

Heterogeneous Streaming	611
<i>Chris J. Newburn, Gaurav Bansal, Michael Wood, Luis Crivelli, Judit Planas, Alejandro Duran, Paulo Souza, Leonardo Borges, Piotr Luszczek, Stanimire Tomov, Jack Dongarra, Hartwig Anzt, Mark Gates, Azzam Haidar, Yulu Jia, Khairul Kabir, Ichitaro Yamazaki, and Jesus Labarta</i>	
HMC-Sim-2.0: A Simulation Platform for Exploring Custom Memory Cube Operations	621
<i>John D. Leidel and Yong Chen</i>	

Alpaka — An Abstraction Library for Parallel Kernel Acceleration	631
<i>Erik Zenker, Benjamin Worpitz, René Widera, Axel Huebl, Guido Juckeland, Andreas Knüpfer, Wolfgang E. Nagel, and Michael Bussmann</i>	
A Tool for Bottleneck Analysis and Performance Prediction for GPU-Accelerated Applications	641
<i>Souley Madougou, Ana Lucia Varbanescu, Cees De Laat, and Rob Van Nieuwpoort</i>	

Session 2: Algorithms and Applications

Hessenberg Reduction with Transient Error Resilience on GPU-Based Hybrid Architectures	653
<i>Yulu Jia, Piotr Luszczek, and Jack Dongarra</i>	
Optimization of Block Sparse Matrix-Vector Multiplication on Shared-Memory Parallel Architectures	663
<i>Ryan Eberhardt and Mark Hoemmen</i>	
Basker: A Threaded Sparse LU Factorization Utilizing Hierarchical Parallelism and Data Layouts	673
<i>Joshua Dennis Booth, Sivasankaran Rajamanickam, and Heidi Thornquist</i>	
Efficiency of General Krylov Methods on GPUs — An Experimental Study	683
<i>Hartwig Anzt, Jack Dongarra, Moritz Kreutzer, Gerhard Wellein, and Martin Köhler</i>	

Session 3: Workload Scheduling

Refactoring Conventional Task Schedulers to Exploit Asymmetric ARM big.LITTLE Architectures in Dense Linear Algebra	692
<i>Luis Costero, Francisco D. Igual, Katzalin Olcoz, Sandra Catalán, Rafael Rodríguez-Sánchez, and Enrique S. Quintana-Ortí</i>	
Heterogeneous CAF-Based Load Balancing on Intel Xeon Phi	702
<i>Valeria Cardellini, Alessandro Fanfarillo, and Salvatore Filippone</i>	
Topology-Aware GPU Selection on Multi-GPU Nodes	712
<i>Iman Faraji, Seyed H. Mirsadeghi, and Ahmad Afsahi</i>	

Workshop 7-PCO - Parallel Computing and Optimization

PCO Introduction and Committees	721
<i>Didier El Baz and Bora Ucar</i>	

Session I: Parallel Computing and Optimization

Scenario Decomposition for 0-1 Stochastic Programs: Improvements and Asynchronous Implementation	722
<i>Kevin Ryan, Deepak Rajan, and Shabbir Ahmed</i>	
PIPS-SBB: A Parallel Distributed-Memory Branch-and-Bound Algorithm for Stochastic Mixed-Integer Programs	730
<i>Lluís-Miquel Munguía, Geoffrey Oxberry, and Deepak Rajan</i>	
Counting Triangles in Large Graphs on GPU	740
<i>Adam Polak</i>	

Session II: Parallel Algorithms for Scheduling problems

GPU-Based Two Level Parallel B&B for the Blocking Job Shop Scheduling Problem	747
<i>Adel Dabah, Ahcène Bendjoudi, Didier El-Baz, and Abdelhakim Aitzai</i>	
Parallel Ant Colony Optimization for Flow Shop Scheduling Subject to Limited Machine Availability	756
<i>Yumei Huo and Jun Xiong Huang</i>	
GPGPU-Based Parallel Algorithms for Scheduling Against Due Date	766
<i>Abhishek Awasthi, Jörg Lässig, Jens Leuschner, and Thomas Weise</i>	
Performance Analysis of Bio-Inspired Scheduling Algorithms for Cloud Environments	776
<i>Ali Al Buhussain, Robson E. De Grande, and Azzedine Boukerche</i>	

Session III: Parallel Heuristics and Metaheuristics

Optimizing Metaheuristics and Hyperheuristics through Multi-level Parallelism on a Many-Core System	786
<i>José Matías Cutillas Lozano, Domingo Giménez, and Luis Pedro García</i>	
A Parallel Ant Colony Optimization for the Maximum-Weight Clique Problem	796
<i>Didier El Baz, Mhand Hifi, Lei Wu, and Xiaochuan Shi</i>	
Evaluating the Performance of A4SDN on Various Network Topologies	801
<i>Giovanni Cammarata, Antonella Di Stefano, Giovanni Morana, and Daniele Zito</i>	
Hybrid Heuristics for Mapping Task Problem on Large Scale Heterogeneous Platforms	809
<i>Ania Kaci, Huy-Nam Nguyen, Amir Nakib, and Patrick Siarry</i>	
A Semi-Greedy Heuristic for the Mapping of Large Task Graphs	817
<i>Karl-Eduard Berger, François Galea, Bertrand Le Cun, and Renaud Sirdey</i>	

Session IV: Combinatorial Scientific Computing

A High Performance Implementation of Spectral Clustering on CPU-GPU Platforms	825
<i>Yu Jin and Joseph F. Jaja</i>	
Testing Fine-Grained Parallelism for the ADMM on a Factor-Graph	835
<i>Ning Hao, Amirreza Oghbaee, Mohammad Rostami, Nate Derbinsky, and José Bento</i>	
High Performance Parallel Graph Coloring on GPGPUs	845
<i>Pingfan Li, Xuhao Chen, Zhe Quan, Jianbin Fang, Huayou Su, Tao Tang, and Canqun Yang</i>	

Workshop 8-GABB - Graph Algorithms Building Blocks

GABB Introduction and Committees	855
<i>Tim Mattson</i>	
GABB 2016 Keynote	856
<i>David Bader</i>	
Array Types for a Graph Processing Language	857
<i>Mark Tullsen and Matthew Sottile</i>	
The Right Way to Search Evolving Graphs	867
<i>Jiahao Chen and Weijian Zhang</i>	
Updating PageRank for Streaming Graphs	877
<i>Jason Riedy</i>	
Application of Graph Sparsification in Developing Parallel Algorithms for Updating Connected Components	885
<i>Sriram Srinivasan, Sanjukta Bhowmick, and Sajal Das</i>	
Towards a Distributed Large-Scale Dynamic Graph Data Store	892
<i>Keita Iwabuchi, Scott Sallinen, Roger Pearce, Brian Van Essen, Maya Gokhale, and Satoshi Matsuoka</i>	
Enforced Sparse Non-negative Matrix Factorization	902
<i>Brendan Gavin, Vijay Gadepally, and Jeremy Kepner</i>	
GBTL-CUDA: Graph Algorithms and Primitives for GPUs	912
<i>Peter Zhang, Marcin Zalewski, Andrew Lumsdaine, Samantha Misurda, and Scott McMillan</i>	
Jaccard Coefficients as a Potential Graph Benchmark	921
<i>Peter M. Kogge</i>	

PageRank Pipeline Benchmark: Proposal for a Holistic System Benchmark for Big-Data Platforms	929
---	-----

*Patrick Dreher, Chansup Byun, Chris Hill, Vijay Gadepally, Bradley Kuszmaul,
and Jeremy Kepner*

Workshop 9-EduPar - NSF/TCPP Workshop on Parallel and Distributed Computing Education

EduPar Introduction and Committees	938
<i>Ramachandran Vaidyanathan, Sushil K Prasad, and Satish Puri</i>	
EduPar 2016 Keynote	941
<i>Randal E. Bryant</i>	

Session 1: Programming Framework and Tools

WebGPU: A Scalable Online Development Platform for GPU Programming Courses	942
<i>Abdul Dakkak, Carl Pearson, and Wen-Mei Hwu</i>	
Parallel Programming with Pictures in a Snap!	950
<i>Annette Feng and Wu-Chun Feng</i>	
Modules to Teach Parallel and Distributed Computing Using MPI for Python and Disco	958
<i>José Ortiz-Ubarri, Rafael Arce-Nazario, and Edusmildo Orozco</i>	
VIPLE: Visual IoT/Robotics Programming Language Environment for Computer Science Education	963
<i>Yinong Chen and Gennaro De Luca</i>	

Session 2: Instruction Techniques and Experiences

Seeing Multithreaded Behavior Using TSGL	972
<i>Joel C. Adams, Patrick A. Crain, and Christopher P. Dilley</i>	
The Suzaku Pattern Programming Framework	978
<i>Barry Wilkinson and Clayton Ferner</i>	
A Flipped Classroom Approach to Teaching Concurrency and Parallelism	987
<i>Shirley V. Moore and Steven R. Dunlop</i>	
A Parallel Programming Course Based on an Execution Time-Energy Consumption Optimization Problem	996
<i>Javier Cuenca and Domingo Giménez</i>	

Workshop 10-HPDAV - High Performance Data Analysis and Visualization

HPDAV Introduction and Committees	1004
<i>Wes Bethel</i>	
HPDAV 2016 Keynote	1006
<i>Jim Jeffers</i>	

Full Papers Session I

Visualization and Analysis for Near-Real-Time Decision Making in Distributed Workflows	1007
<i>David Pugmire, James Kress, Jong Choi, Scott Klasky, Tahsin Kurc, Randy Michael Churchill, Matthew Wolf, Greg Eisenhower, Hank Childs, Kesheng Wu, Alexander Sim, Junmin Gu, and Jonathan Low</i>	
High Performance Molecular Visualization: In-Situ and Parallel Rendering with EGL	1014
<i>John E. Stone, Peter Messmer, Robert Sisneros, and Klaus Schulten</i>	
Introducing Acacia-RDF: An X10-Based Scalable Distributed RDF Graph Database Engine	1024
<i>Miyuru Dayarathna, Isuru Herath, Yasima Dewmini, Gayan Mettananda, Sameera Nandasiri, Sanath Jayasena, and Toyotaro Suzumura</i>	

Short Papers Session

Towards Asynchronous Many-Task in Situ Data Analysis Using Legion	1033
<i>Philippe Pébaÿ, Janine C. Bennett, David Hollman, Sean Treichler, Patrick S. McCormick, Christine M. Sweeney, Hemanth Kolla, and Alex Aiken</i>	
Coupling LAMMPS and the vl3 Framework for Co-Visualization of Atomistic Simulations	1038
<i>Silvio Rizzi, Mark Hereld, Joseph Insley, Preeti Malakar, Michael E. Papka, Thomas Uram, and Venkatram Vishwanath</i>	
Developing a Scalable SNMP Monitor	1043
<i>Krishna Bharadwaj, Samuel Flores, Joshua Rodriguez, Lance Long, and G. Elisabeta Marai</i>	

Full Papers Session II

Immersive Molecular Visualization with Omnidirectional Stereoscopic Ray Tracing and Remote Rendering	1048
<i>John E. Stone, William R. Sherman, and Klaus Schulten</i>	
Tuned to Terrible: A Study of Parallel Particle Advection State of the Practice	1058
<i>Robert Sisneros and David Pugmire</i>	

Workshop 11-VarSys - Variability in Parallel and Distributed Systems

VarSys Introduction	1068
<i>Kirk Cameron, Todd Gamblin, and Dimitrios S. Nikolopoulos</i>	
Variability: A Tuning Headache	1069
<i>Allan Porterfield, Sridutt Bhalachandra, Wei Wang, and Rob Fowler</i>	
Mitigating Processor Variation through Dynamic Load Balancing	1073
<i>Bilge Acun and Laxmikant V. Kale</i>	
Characterizing and Reducing Cross-Platform Performance Variability Using OS-Level Virtualization	1077
<i>Ivo Jimenez, Carlos Maltzahn, Jay Lofstead, Adam Moody, Kathryn Mohror, Remzi Arpacı-Dusseau, and Andrea Arpacı-Dusseau</i>	
Towards Managing Variability in the Cloud	1081
<i>Ali Anwar, Yue Cheng, and Ali R. Butt</i>	
Near Real-Time Tracking of IoT Device Users	1085
<i>Jinseong Kim, Jae J. Jang, and Im Y. Jung</i>	

Workshop 12-HPPAC - High-Performance, Power-Aware Computing

HPPAC Introduction and Committees	1089
<i>Barry Rountree and Shuaiwen Leon Song</i>	
Lightning Talks A	
The Right Metric for Efficient Supercomputing: A Ten-Year Retrospective	1090
<i>Chung-Hsing Hsu and Wu-Chun Feng</i>	
Overcoming Challenges in Scalable Power Monitoring with the Power API	1094
<i>Ryan E. Grant, Michael Levenhagen, Stephen L. Olivier, David DeBonis, Kevin Pedretti, and James H. Laros III</i>	
Achieving Safety for Power Shifting in Overprovisioned High Performance Computing Systems	1098
<i>Shirley Moore</i>	
POSITION PAPER: Countering the Noise-Induced Critical Path Problem	1102
<i>Rogelio Long and Shirley Moore</i>	

Lightning Talks B

Re-Examining HPC Energy Efficiency Dashboard Elements	1106
<i>Natalie Bates, Chung-Hsing Hsu, Neena Imam, Torsten Wilde, and Dale Sartor</i>	
A Power-Aware Cost Model for HPC Procurement	1110
<i>Neha Ghokar, Frank Mueller, and Barry Rountree</i>	

Energy Claims at Scale: Decreasing the Energy Demand of HPC Workloads at OS Level	1114
--	------

Christopher Eibel, Timo Höning, and Wolfgang Schröder-Preikschat

Systemwide Power Management with Argo	1118
---	------

*Daniel Ellsworth, Tapasya Patki, Swann Perarnau, Sangmin Seo,
Abdelhalim Amer, Judicael Zounmevo, Rinku Gupta, Kazutomo Yoshii,
Henry Hoffman, Allen Malony, Martin Schulz, and Pete Beckman*

Regular Papers A

Best Practices for Scalable Power Measurement and Control	1122
---	------

Scott Walker and Marty McFadden

LibPowerMon: A Lightweight Profiling Framework to Profile Program Context and System-Level Metrics	1132
---	------

Aniruddha Marathe, Hormozd Gahvari, Jae-Seung Yeom, and Abhinav Bhatale

Power Balancing in an Emulated Exascale Environment	1142
---	------

Matthias Maiterth, Martin Schulz, Dieter Kranzmüller, and Barry Rountree

Combining Power and Performance Modeling for Application Analysis: A Case Study Using Aspen	1150
--	------

Sand L. Correa, Mariam Umar, and Kirk W. Cameron

Effective Utilization of CUDA Hyper-Q for Improved Power and Performance Efficiency	1160
--	------

Ryan S. Luley and Qinru Qiu

Regular Papers B

Identification of Critical Parameters for MapReduce Energy Efficiency Using Statistical Design of Experiments	1170
--	------

Nidhi Tiwari, Umesh Bellur, Santonu Sarkar, and Maria Indrawan

Utilizing Hardware Performance Counters to Model and Optimize the Energy and Performance of Large Scale Scientific Applications on Power-Aware Supercomputers	1180
---	------

Xingfu Wu and Valerie Taylor

Energy, Power, and Performance Characterization of GPGPU Benchmark Programs	1190
--	------

Jared Coplin and Martin Burtscher

Workshop 13-PDSEC - Parallel and Distributed Scientific and Engineering Computing

PDSEC Introduction and Committees	1200
<i>Peter Strazdins, Raphaël Couturier, Keita Teranishi, Alan Gray, Thomas Rauber, Gudula Rünger, and Laurence T. Yang</i>	

Session 1: Application and Task Parallelism

Electron Dynamics Simulation with Time-Dependent Density Functional Theory on Large Scale Symmetric Mode Xeon Phi Cluster	1202
<i>Yuta Hirokawa, Taisuke Boku, Shunsuke A. Sato, and Kazuhiro Yabana</i>	
Towards an Efficient Task-Based Parallelization over a Runtime System of an Explicit Finite-Volume CFD Code with Adaptive Time Stepping	1212
<i>Jean Marie Couteyen Carpaye, Jean Roman, and Pierre Brenner</i>	
Radiative Heat Transfer Calculation on 16384 GPUs Using a Reverse Monte Carlo Ray Tracing Approach with Adaptive Mesh Refinement	1222
<i>Alan Humphrey, Daniel Sunderland, Todd Harman, and Martin Berzins</i>	

Session 2: Resilience

Application Fault Tolerance for Shrinking Resources via the Sparse Grid Combination Technique	1232
<i>Peter E. Strazdins, Md. Mohsin Ali, and Bert Debuisschere</i>	
Two-Level Checkpointing and Verifications for Linear Task Graphs	1239
<i>Anne Benoit, Aurélien Cavelan, Yves Robert, and Hongyang Sun</i>	

Session 3: Performance

On the Development of Variable Size Batched Computation for Heterogeneous Parallel Architectures	1249
<i>Ahmad Abdelfattah, Azzam Haidar, Stanimire Tomov, and Jack Dongarra</i>	
Synapse: Synthetic Application Profiler and Emulator	1259
<i>Andre Merzky and Shantenu Jha</i>	

Workshop 14-DPDNS - Dependable Parallel, Distributed and Network-Centric Systems

DPDNS Introduction and Committees	1269
<i>Dimiter Avresky, Erik Maeble, and Roberto Palmieri</i>	
DPDNS 2016 Keynote	1270
<i>Shlomi Dolev</i>	

Session 1: Distributed Services

A Game Theoretic Approach for Diagnosing Faults in Asymmetric Comparison-Based Distributed and Parallel Systems	1271
<i>Mourad Elhadef</i>	
Distributed Decentralized Domain Name Service	1279
<i>Brendan Benshoof, Andrew Rosen, Anu G. Bourgeois, and Robert W. Harrison</i>	
Management Software for Protocol-Level Adaptations in Dependable Network Services	1288
<i>Kaliappa Ravindran</i>	

Session 2: Cloud and Fault Tolerance

Mitigating Routing Inefficiencies to Cloud-Storage Providers: A Case Study	1298
<i>Soham Sinha, Di Niu, Zhi Wang, and Paul Lu</i>	
Leaderless Consensus: The State of the Art	1307
<i>Roberto Palmieri</i>	
Proactive Cloud Management for Highly Heterogeneous Multi-cloud Infrastructures	1311
<i>Alessandro Pellegrini, Pierangelo Di Sanzo, and Dimiter R. Avresky</i>	

Session 3: Multicore Computing

Towards Resiliency Evaluation of Vector Programs	1319
<i>Vishal Chandra Sharma, Ganesh Gopalakrishnan, and Sriram Krishnamoorthy</i>	
Analysis of Adaptive Mapping of Parallelized Application on Multicore System	1329
<i>Giiles Bizot, Dimiter Avresky, and Fabien Chaix</i>	

Workshop 15-LSPP - Large-Scale Parallel Processing

LSPP Introduction and Committees	1339
<i>Kevin J. Barker, Chris D. Carothers, and Eric van Hensbergen</i>	
LSPP 2016 Keynote	1340
<i>Michael Papka</i>	

Session 1: Making Efficient Use of Advanced Architectures

Evaluating the Performance Impact of Multiple Streams on the MIC-Based Heterogeneous Platform	1341
<i>Zhaokui Li, Jianbin Fang, Tao Tang, Xuhao Chen, Cheng Chen, and Canqun Yang</i>	

Parallel Implementation Strategies for Hierarchical Non-uniform Memory Access Systems by Example of the Scale-Invariant Feature Transform Algorithm	1351
---	------

Max Plauth, Wieland Hagen, Frank Feinbube, Felix Eberhardt, Lena Feinbube, and Andreas Polze

Efficient Genetic Algorithm Encoding for Large-Scale Multi-objective Resource Allocation	1360
--	------

Ryan D. Friese

Session 2: Workflow Modeling and Optimization and Modeling at Scale

Toward an End-to-End Framework for Modeling, Monitoring and Anomaly Detection for Scientific Workflows	1370
--	------

Anirban Mandal, Paul Ruth, Ilya Baldin, Dariusz Król, Gideon Juve, Rajiv Mayani, Rafael Ferreira Da Silva, Ewa Deelman, Jeremy Meredith, Jeffrey Vetter, Vickie Lynch, Ben Mayer, James Wynne III, Mark Blanco, Chris Carothers, Justin Lapre, and Brian Tierney

Modeling the Performance and Energy Impact of Dynamic Power Steering	1380
--	------

Kevin J. Barker and Darren J. Kerbyson

Workshop 16-ParLearning - Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analytics

ParLearning Introduction and Committees	1390
---	------

Charalampos Chelmis, Sutanay Choudhury, Arindam Pal, Anand Panangadan, Weiqin Tong, and Yinglong Xia

ParLearning 2016 Keynote	1392
--------------------------------	------

Peter M. Kogge

Session I

A Novel Scalable DBSCAN Algorithm with Spark	1393
--	------

Dianwei Han, Ankit Agrawal, Wei-Keng Liao, and Alok Choudhary

A Multi-Platform Evaluation of the Randomized CX Low-Rank Matrix Factorization in Spark	1403
---	------

Alex Gittens, Jey Kottalam, Jiyan Yang, Michael F. Ringenburg, Jatin Chhugani, Evan Racah, Mohitdeep Singh, Yushu Yao, Curt Fischer, Oliver Ruebel, Benjamin Bowen, Norman G. Lewis, Michael W. Mahoney, Venkat Krishnamurthy, and Prabhat

Cache-Aware Approximate Computing for Decision Tree Learning	1413
--	------

Orhan Kisyal, Mahmut T. Kandemir, and Jagadish Kotra

Accelerating Support Count for Association Rule Mining on GPUs	1423
--	------

Vasileios Zois, Anand Panangadan, and Viktor Prasanna

Session II

A Scheduling Algorithm for Hadoop MapReduce Workflows with Budget Constraints in the Heterogeneous Cloud	1433
<i>Andrew Wylie, Wei Shi, Jean-Pierre Corriveau, and Yang Wang</i>	
An Automatic Tuning System for Solving NP-Hard Problems in Clouds	1443
<i>Yanik Ngoko, Denis Trystram, Valentin Reis, and Christophe Cérin</i>	
GraQL: A Query Language for High-Performance Attributed Graph Databases	1453
<i>Daniel Chavarria-Miranda, Vito Giovanni Castellana, Alessandro Morari, David Haglin, and John Feo</i>	
Scalable Overlapping Community Detection	1463
<i>Ismail El-Helw, Rutger Hofman, Wenzhe Li, Sungjin Ahn, Max Welling, and Henri Bal</i>	

Session III

An Efficient Parallel Nonlinear Clustering Algorithm Using MapReduce	1473
<i>Xiang-You Peng, Yu-Bo Yang, Chang-Dong Wang, Dong Huang, and Jian-Huang Lai</i>	
A New Evaluation System for Scholars and Majors Based on Big-Data Techniques	1477
<i>Wenhua Yu, Lei Zhao, Xiangyu He, Jiacheng Zhou, Tong Cheng, Chengzhao Xue, and Fan Yang</i>	
Open Source Initiatives and Frameworks Addressing Distributed Real-Time Data Analytics	1481
<i>Sarwar Jahan Morshed, Juwel Rana, and Marcelo Milrad</i>	

Workshop 17-JSSPP - Job Scheduling Strategies for Parallel Processing

JSSPP Introduction and Committees	1485
<i>Walredo Cirne and Narayan Desai</i>	

Workshop 18-iWAPT - International Workshop on Automatic Performance Tuning

iWAPT Introduction and Committees	1486
<i>Weichung Wang</i>	

Session 1

Auto-Tuning of Hybrid MPI/OpenMP Execution with Code Selection by ppOpen-AT	1488
--	------

Takahiro Katagiri, Masaharu Matsumoto, and Satoshi Ohshima

Utilization and Expansion of ppOpen-AT for OpenACC	1496
--	------

Satoshi Ohshima, Takahiro Katagiri, and Masaharu Matsumoto

Session 2

Measurement Bias from Address Aliasing	1506
--	------

Lars Kirkholz Melhus and Rune Erlend Jensen

Blk-Tune: Blocking Parameter Auto-Tuning to Minimize Input-Output Traffic for Flash-Based Out-of-Core Stencil Computations	1516
---	------

Hiroko Midorikawa

A Time-Cost Based Automatic Scheduling Framework for Matrix Computation on Various Distributed Computing Platforms	1527
---	------

Rong Gu, Zhiqiang Liu, Chunfeng Yuan, and Yihua Huang

Session 3

Exploiting Performance Portability in Search Algorithms for Autotuning	1535
--	------

Amit Roy, Prasanna Balaprakash, Paul D. Hovland, and Stefan M. Wild

Search Space Generation and Pruning System for Autotuners	1545
---	------

Piotr Luszczek, Mark Gates, Jakub Kurzak, Anthony Danalis, and Jack Dongarra

Workshop 19-CHIUW - Chapel Implementers and Users Workshop

CHIUW Introduction and Committees	1555
---	------

Tom MacDonald and Greg Titus

CHIUW 2016 Keynote	1557
--------------------------	------

Nikhil Padmanabhan

Session 1: Benchmarking and Optimization

Optimizing Chapel for Single-Node Environments	1558
--	------

Richard B. Johnson and Jeff Hollingsworth

PGAS Access Overhead Characterization in Chapel	1568
---	------

*Engin Kayraklıoglu, Olivier Serres, Ahmad Anbar, Hashem Elezabi,
and Tarek El-Ghazawi*

Session 2: Chapel Improvement

Chplvis: A Communication and Task Visualization Tool for Chapel	1578
<i>Philip A. Nelson and Greg Titus</i>	
Transparently Resilient Task Parallelism for Chapel	1586
<i>Konstantina Panagiotopoulou and Hans-Wolfgang Loidl</i>	

Workshop 20-HPBDC - High-Performance Big Data Computing

HPBDC Introduction and Committees	1596
<i>Dhabaleswar K. (DK) Panda, Jianfeng Zhan, and Xiaoyi Lu</i>	

Session I: High-Performance Big Data Applications and Systems

Evaluation of SMP Shared Memory Machines for Use with In-Memory and OpenMP Big Data Applications	1597
<i>Andrew J. Younge, Christopher Reidy, Robert Henschel, and Geoffrey C. Fox</i>	
Hadoop on HPC: Integrating Hadoop and Pilot-Based Dynamic Resource Management	1607
<i>Andre Luckow, Ioannis Paraskevakos, George Chantzialexiou, and Shantenu Jha</i>	
PACM: A Prediction-Based Auto-Adaptive Compression Model for HDFS	1617
<i>Ruijian Wang, Chao Wang, and Li Zha</i>	

Session II: High-Performance Streaming Systems

SamzaSQL: Scalable Fast Data Management with Streaming SQL	1627
<i>Milinda Pathirage, Julian Hyde, Yi Pan, and Beth Plale</i>	
Towards High Performance Processing of Streaming Data in Large Data Centers	1637
<i>Supun Kamburugamuve, Saliya Ekanayake, Milinda Pathirage, and Geoffrey Fox</i>	
Extracting Log Patterns from System Logs in LARGE	1645
<i>Yining Zhao and Haili Xiao</i>	

Session III (Short Papers): Performance Studies of Big Data Systems and Applications

Exploring the Performance of Spark for a Scientific Use Case	1653
<i>Saba Sehrish, Jim Kowalkowski, and Marc Paterno</i>	
Big Data for Medical Image Analysis: A Performance Study	1660
<i>Rui Zhang, Hongzhi Wang, Renu Tewari, Gero Schmidt, and Deepika Kakrania</i>	

Workshop 21-HPCMASPA - Monitoring and Analysis for High Performance Computing Systems Plus Applications

HPCMASPA Introduction and Committees	1665
<i>Benjamin Allan, Jim Brandt, Ann Gentile, Cory Lueninghoener, Nichamon Naksinehaboon, Boyana Norris, and Narate Taerat</i>	
HPCMASPA 2016 Keynote	1667
<i>William T.C. Kramer</i>	

Session 1: Instrumentation and Metrics

Calltree-Controlled Instrumentation for Low-Overhead Survey Measurements	1668
<i>Christian Iwinsky and Christian Bischof</i>	
Automatically Instrumenting Scientific Applications to Produce Heartbeat Events	1678
<i>Mohammed Tanash, Nasim Ghazanfari, Omar Aaziz, and Jonathan Cook</i>	
Defining Metrics to Distill Large-Scale HPC Platform and Application Performance Data into Actionable Quantities	1687
<i>Anthony Agelastos</i>	

Session 2: Monitoring Systems

Using Intrinsic Performance Counters to Assess Efficiency in Task-Based Parallel Applications	1692
<i>Patricia Grubel, Hartmut Kaiser, Kevin Huck, and Jeanine Cook</i>	
Understanding Application and System Performance Through System-Wide Monitoring	1702
<i>R. Todd Evans, James C. Browne, and William L. Barth</i>	
Large-Scale Persistent Numerical Data Source Monitoring System Experiences	1711
<i>J. Brandt, A. Gentile, M. Showerman, J. Enos, J. Fullop, and G. Bauer</i>	
Design and Implementation of a Scalable HPC Monitoring System	1721
<i>S. Sanchez, A. Bonnie, G. Van Heule, C. Robinson, A. DeConinck, K. Kelly, Q. Snead, and J. Brandt</i>	

Workshop 22-IPDRM - Emerging Parallel and Distributed Runtime Systems and Middleware

IPDRM Introduction and Committees	1726
<i>Shuaiwen Leon Song and Todd Gamblin</i>	
IPDRM 2016 Keynote	1727
<i>Henry Hoffmann</i>	

Session 1

Non-intrusive Migration of MPI Processes in OS-Bypass Networks	1728
<i>Simon Pickartz, Carsten Clauss, Stefan Lankes, Stephan Krempel, Thomas Moschny, and Antonello Monti</i>	
Photon: Remote Memory Access Middleware for High-Performance Runtime Systems	1736
<i>Ezra Kissel and Martin Swany</i>	
Asynchronous Runtimes in Action: An Introspective Framework for a Next Gen Runtime	1744
<i>Joshua Suetterlein, Joshua Landwehr, Andrés Márquez, Joseph B. Manzano, and Guang R. Gao</i>	

Session 2

OWBP: Flash-Aware Offline Write Buffer Policy	1752
<i>Alireza Haghdoost and David H. C. Du</i>	
Topology-Aware Rank Reordering for MPI Collectives	1759
<i>Seyed H. Mirsadeghi and Ahmad Afsahi</i>	
GPUShare: Fair-Sharing Middleware for GPU Clouds	1769
<i>Anshuman Goswami, Jeffrey Young, Karsten Schwan, Naila Farooqui, Ada Gavrilovska, Matthew Wolf, and Greg Eisenhauer</i>	

Session 3

Performance Characterization of Hypervisor-and Container-Based Virtualization for HPC on SR-IOV Enabled InfiniBand Clusters	1777
<i>Jie Zhang, Xiaoyi Lu, and Dhabaleswar K. (DK) Panda</i>	
Macaca: A Scalable and Energy-Efficient Platform for Coupling Cloud Computing with Distributed Embedded Computing	1785
<i>Heng Zhang, Chunliang Hao, Yanjun Wu, and Mingshu Li</i>	
Benchmarking Streaming Computation Engines: Storm, Flink and Spark Streaming	1789
<i>Sanket Chintapalli, Derek Dagit, Bobby Evans, Reza Farivar, Thomas Graves, Mark Holderbaugh, Zhuo Liu, Kyle Nusbaum, Kishorkumar Patil, Boyang Jerry Peng, and Paul Poulosky</i>	

Workshop 23-ParSocial - Parallel and Distributed Processing for Computational Social Systems

ParSocial Introduction and Committees	1793
<i>Eunice E. Santos and John Korah</i>	
ParSocial 2016 Keynote	1795
<i>George Cybenko</i>	

Paper Session 1

Towards Reliable Social Sensing in Cyber-Physical-Social Systems	1796
<i>Chao Huang, Jermaine Marshall, Dong Wang, and Mianxiong Dong</i>	
Toward the New Version of D-MASON: Efficiency, Effectiveness and Correctness in Parallel and Distributed Agent-Based Simulations	1803
<i>Gennaro Cordasco, Carmine Spagnuolo, and Vittorio Scarano</i>	
Emergency-Driven Assured Information Sharing in Secure Online Social Networks: A Position Paper	1813
<i>Bhavani Thuraisingham, Murat Kantarcioglu, Latifur Khan, Barbara Carminati, Elena Ferrari, and Leila Bahri</i>	

Paper Session 2

Efficient Anytime Anywhere Algorithms for Closeness Centrality in Large and Dynamic Graphs	1821
<i>Eunice E. Santos, John Korah, Vairavan Murugappan, and Suresh Subramanian</i>	
Addressing Behavioral Uncertainty in Security Games: An Efficient Robust Strategic Solution for Defender Patrols	1831
<i>Thanh H. Nguyen, Arunesh Sinha, and Milind Tambe</i>	

Workshop 24-Roundtable I - PDC in Core Undergraduate Education

Workshop 24-Roundtable I Introduction	1839
<i>Dick Brown and Suzanne Matthews</i>	

Author Index