

# **2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW 2017)**

**Orlando, Florida, USA  
29 May - 2 June 2017**

**Pages 1-804**



**IEEE Catalog Number: CFP1751J-POD  
ISBN: 978-1-5386-3409-7**

**Copyright © 2017 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:	CFP1751J-POD
ISBN (Print-On-Demand):	978-1-5386-3409-7
ISBN (Online):	978-1-5386-3408-0

**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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2017 IEEE International Parallel and Distributed Processing Symposium Workshops

## IPDPSW 2017

### Table of Contents

Message from the General Chair .....	xxiii
Message from the Workshops Chair and Vice-Chair.....	xxv

---

#### **HCW: Heterogeneity in Computing Workshop**

Introduction to HCW Workshop .....	1
<i>Erik Saule and Emmanuel Jeannot</i>	
Message from the HCW Steering Committee Chair .....	2
<i>Behrooz Shirazi</i>	
Message from the HCW General Chair .....	3
<i>Erik Saule</i>	
Message from the HCW Program Committee Chair .....	4
<i>Emmanuel Jeannot</i>	
HCW Keynote Talk .....	5
<i>Ricky Kwok</i>	

#### **Session 1: Managing the Different Components of Heterogeneous Systems**

Portable Implementation of Advanced Driver-Assistance Algorithms on Heterogeneous Architectures .....	6
<i>Oliver Jakob Arndt, Fabian David Träger, Tobias Moß, and Holger Blume</i>	
Improving CPU Performance Through Dynamic GPU Access Throttling in CPU-GPU Heterogeneous Processors .....	18
<i>Siddharth Rai and Mainak Chaudhuri</i>	
Transparent Heterogeneous Backing Store for File Systems .....	30
<i>Benjamin Marks and Tia Newhall</i>	

## Session 2: Scheduling and Resource Allocation

Alternative Processor Within Threshold: Flexible Scheduling on Heterogeneous Systems .....	42
<i>Sonia Lopez and Stavan Satish Karia</i>	
Preemptive Resource Management for Dynamically Arriving Tasks in an Oversubscribed Heterogeneous Computing System .....	54
<i>Dylan Machovec, Sudeep Pasricha, Anthony A. Maciejewski, Howard Jay Siegel, Gregory A. Koenig, Michael Wright, Marcia Hilton, Rajendra Rambharos, Thomas Naughton, and Neena Imam</i>	
Modeling of Applications and Hardware to Explore Task Mapping and Scheduling Strategies on a Heterogeneous Micro-Server System .....	65
<i>Lilia Zaourar, Massinissa Ait Aba, David Briand, and Jean-Marc Philippe</i>	
Consumer-and-Provider-Oriented Efficient IaaS Resource Allocation .....	77
<i>Thibaud Ecarot, Djamel Zeglache, and Cedric Brandily</i>	

## RAW: Reconfigurable Architectures Workshop

Introduction to RAW Workshop .....	86
<i>Marco D. Santambrogio and Ramachandran Vaidyanathan</i>	
RAW Keynote Speakers .....	88
<i>Ronald F. DeMara and Georgi Gaydadjiev</i>	

## Session 1: Architectures for Convolutional Neural Networks and Sliding Window

A Pipelined and Scalable Dataflow Implementation of Convolutional Neural Networks on FPGA .....	90
<i>Marco Bacis, Giuseppe Natale, Emanuele Del Sozzo, and Marco Domenico Santambrogio</i>	
On-Chip Memory Based Binarized Convolutional Deep Neural Network Applying Batch Normalization Free Technique on an FPGA .....	98
<i>Haruyoshi Yonekawa and Hiroki Nakahara</i>	
A Modified Sliding Window Architecture for Efficient BRAM Resource Utilization .....	106
<i>Murad Qasaimah, Joseph Zambreno, and Phillip H. Jones</i>	

## Session 2: Design and Programming Methods

Automatic Flow Selection and Quality-of-Result Estimation for FPGA Placement .....	115
<i>G. Grewal, S. Areibi, M. Westrik, Z. Abuowaimer, and B. Zhao</i>	
Exploiting Decoupled OpenCL Work-Items with Data Dependencies on FPGAs: A Case Study .....	124
<i>Javier Alejandro Varela, Norbert Wehn, Qian Liang, and Songyin Tang</i>	
Exploiting FPGAs from Higher Level Languages A Signal Analysis Case Study .....	132
<i>L. Stornaiuolo, A. Parravicini, G. Durelli, and M. D. Santambrogio</i>	

ReEP: A Toolset for Generation and Programming of Reconfigurable Datapaths for Event Processing .....	141
<i>Philip Gottschling and Christian Hochberger</i>	

### **Session 3: Acceleration of Curran's Approximation and Elliptic Curve Crypto**

A Scalable Dataflow Implementation of Curran's Approximation Algorithm .....	150
<i>Anna Maria Nestorov, Enrico Reggiani, Hristina Palikareva, Pavel Burovskiy, Tobias Becker, and Marco D. Santambrogio</i>	
A Generic Approach to the Development of Coprocessors for Elliptic Curve Cryptosystems .....	158
<i>Rabia Shahid, Ted Winograd, and Kris Gaj</i>	

### **Session 4: Acceleration of Biological Signal Processing**

A Hardware Acceleration for Surface EMG Non-Negative Matrix Factorization .....	168
<i>Luca Cerina, Pierandrea Cancian, Giuseppe Franco, and Marco Domenico Santambrogio</i>	
On-FPGA Real-Time Processing of Biological Signals From High-Density MEAs: a Design Space Exploration .....	175
<i>Giovanni Pietro Seu, Gian Nicola Angotzi, Giuseppe Tuveri, Luigi Raffo, Luca Berdondini, Alessandro Maccione, and Paolo Meloni</i>	

### **Session 5: Design Methods**

Combining Boolean Gates and Branching Programs in One Model can Lead to Faster Circuits .....	184
<i>Yosi Ben-Asher, Esti Stein, and Ramachandran Vaidyanathan</i>	
Efficient Totally-Ordered Subset Generation, with Application in Partial Reconfiguration .....	192
<i>Utsav Agarwal and Ramachandran Vaidyanathan</i>	

### **Short Papers**

FAReP: Fragmentation-Aware Replacement Policy for Task Reuse on Reconfigurable FPGAs .....	202
<i>Godwin Enemali, Adewale Adetomi, and Tughrul Arslan</i>	
Power Analysis of HLS-Designed Customized Instruction Set Architectures .....	207
<i>Tejaswini Ananthanarayana, Sonia Lopez, and Marcin Lukowiak</i>	
A Near Optimal Integrated Solution for Resource Constrained Scheduling, Binding and Routing on CGRAs .....	213
<i>Tajas Ruschke, Lukas Johannes Jung, and Christian Hochberger</i>	
Clock Buffers, Nets, and Trees for On-Chip Communication: A Novel Network Access Technique in FPGAs .....	219
<i>Adewale Adetomi, Godwin Enemali, and Tughrul Arslan</i>	
Pearson Correlation Coefficient Acceleration for Modeling and Mapping of Neural Interconnections .....	223
<i>Enrico Reggiani, Eleonora D'Arnese, Andrea Purgato, and Marco D. Santambrogio</i>	

Out-of-Order Execution of Buffered Function Units in Exposed Data Path Architectures .....	229
<i>Tripti Jain, Klaus Schneider, and Frederik Walk</i>	
Dynamic Dual Fixed-Point CORDIC Implementation .....	235
<i>Andres Jacoby and Daniel Llamocca</i>	
A Highly Scalable and Efficient Parallel Design of N-Body Simulation on FPGA .....	241
<i>Emanuele Del Sozzo, Lorenzo Di Tucci, and Marco D. Santambrogio</i>	
Feasibility Study of Real-Time Spiking Neural Network Simulations on a Swarm Intelligence Based Digital Architecture .....	247
<i>Francesca Palumbo, Carlo Sau, Danilo Pani, Paolo Meloni, and Luigi Raffo</i>	

## **HiCOMB: 16th IEEE International Workshop on High Performance Computational Biology**

Introduction to HiCOMB Workshop .....	251
<i>Alex Pothén and Ananth Grama</i>	
HiCOMB Keynote .....	252
<i>Radu Marculescu</i>	

### **Session 1**

Scalable FRaC Variants: Anomaly Detection for Precision Medicine .....	253
<i>Cyrus Cousins, Chirstopher M. Pietras, and Donna K. Slonim</i>	
Exploratory Modeling and Simulation of the Evolutionary Dynamics of Single-Stranded RNA Virus Populations .....	263
<i>Jae-Seung Yeom, Tanya Kostova-Vassilevska, Peter D. Barnes Jr., David R. Jefferson, and Tomas Opielstrup</i>	

### **Session 2**

Parallel NGS Assembly Using Distributed Assembly Graphs Enriched with Biological Knowledge .....	273
<i>Julia D. Warnke-Sommer and Hesham H. Ali</i>	
Parallel and Memory-Efficient Preprocessing for Metagenome Assembly .....	283
<i>Vasudevan Rengasamy, Paul Medvedev, and Kamesh Madduri</i>	

### **Session 3**

Scalable Parallelization of a Markov Coalescent Genealogy Sampler .....	293
<i>Philip E. Davis, Adam M. Terwilliger, David Zeitler, and Greg Wolfe</i>	
Par-eXpress: A Tool for Analysis of Sequencing Experiments With Ambiguous Assignment of Fragments in Parallel .....	303
<i>Mucahid Kutlu, Gagan Agrawal, and James S. Blachly</i>	

## **EduPar: NSF/TCPP Workshop on Parallel and Distributed Computing Education**

Introduction to EduPar Workshop .....	311
<i>Sheikh Ghafoor, Sushil K Prasad, and Satish Puri</i>	
EduPar Keynote .....	314
<i>Jack Dongarra</i>	

### **Session 1: Tools and Programming Environment**

RAI: A Scalable Project Submission System for Parallel Programming Courses .....	315
<i>Abdul Dakkak, Carl Pearson, Cheng Li, and Wen-mei Hwu</i>	
Introducing Parallel and Distributed Computing to K12 .....	323
<i>Brian Broll, Ákos Lédeczi, Péter Völgyesi, János Sallai, Miklós Maróti, and Chris Vanags</i>	
Log Visualization Tool for Message-Passing Programming in Pilot .....	331
<i>Tianyi Bao and William B. Gardner</i>	
I Can Has Supercomputer? A Novel Approach to Teaching Parallel and Distributed Computing Concepts Using a Meme-Based Programming Language .....	339
<i>David A Richie and James A Ross</i>	

### **Session 2: Pedagogy and Experience**

Teaching Future Big Data Analysts: Curriculum and Experience Report .....	346
<i>Joshua Eckroth</i>	
Hacking at the Divide Between Polar Science and HPC: Using Hackathons as Training Tools .....	352
<i>Jane Wyngaard, Heather Lynch, Jaroslaw Nabrzyski, Allen Pope, and Shantenu Jha</i>	
Preparing an Online Java Parallel Computing Course .....	360
<i>Vivek Sarkar, Max Grossman, Zoran Budimlić, and Shams Imam</i>	
A Laboratory Based Course on GPU Programming: Methods, Practices, and Lessons .....	367
<i>Jawwad Ahmed Shamsi</i>	

## **ParLearning: The 6th International Workshop on Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analytics**

Introduction to ParLearning Workshop .....	375
<i>Anand Panangadan</i>	
ParLearning Keynotes .....	377
<i>John Feo and Wei Tan</i>	

### **Session 1**

ExtDict: Extensible Dictionaries for Data- and Platform-Aware Large-Scale Learning .....	379
<i>Azalia Mirhoseini, Bitá Darvish Rouhani, Ebrahim Songhori, and Farinaz Koushanfar</i>	
Coded TeraSort .....	389
<i>Songze Li, Sucha Supittayapornpong, Mohammad Ali Maddah-Ali, and Salman Avestimehr</i>	

Scaling Deep Learning Workloads: NVIDIA DGX-1/Pascal and Intel Knights Landing .....	399
<i>Nitin A. Gawande, Joshua B. Landwehr, Jeff A. Daily, Nathan R. Tallent, Abhinav Vishnu, and Darren J. Kerbyson</i>	

Efficient and Portable ALS Matrix Factorization for Recommender Systems .....	409
<i>Jing Chen, Jianbin Fang, Weifeng Liu, Tao Tang, Xuhao Chen, and Canqun Yang</i>	

## Session 2

Large-Scale Stochastic Learning Using GPUs .....	419
<i>Thomas Parnell, Celestine Duenner, Kubilay Atasu, Manolis Sifalakis, and Haris Pozidis</i>	

Distributed and in-Situ Machine Learning for Smart-Homes and Buildings: Application to Alarm Sounds Detection .....	429
<i>Amaury Durand, Yanik Ngoko, and Christophe Cérin</i>	

The New Large-Scale RNNLM System Based on Distributed Neuron .....	433
<i>Dejjiao Niu, Rui Xue, Tao Cai, Hai Li, Kingsley Effah, and Hang Zhang</i>	

Cache Friendly Parallelization of Neural Encoder-Decoder Models Without Padding on Multi-core Architecture .....	437
<i>Yuchen Qiao, Kazuma Hashimoto, Akiko Eriguchi, Haixia Wang, Dongsheng Wang, Yoshimasa Tsuruoka, and Kenjiro Taura</i>	

## PDCO: 7th IEEE Workshop Parallel / Distributed Computing and Optimization

Introduction to PDCO Workshop .....	441
<i>Grégoire Danoy and Didier El Baz</i>	

### Session 1: Scheduling I

A Parallel Approximation Algorithm for Scheduling Parallel Identical Machines .....	442
<i>Laleh Ghalami and Daniel Grosu</i>	

Communication Aware task Placement for Workflow Scheduling on DaaS-Based Cloud .....	452
<i>Hadrien Croubois and Eddy Caron</i>	

Dynamic Mapping of Application Workflows in Heterogeneous Computing Environments .....	462
<i>Muhammad Qasim, Touseef Iqbal, Ehsan Ullah Munir, Nikos Tziritas, Samee U. Khan, and Laurence T. Yang</i>	

### Session 2: Scheduling II

Load-Aware Strategies for Cloud-Based VoIP Optimization with VM Startup Prediction .....	472
<i>Jorge M. Cortés-Mendoza, Andrei Tchernykh, Igor Bychkov, Alexander Feoktistov, Pascal Bouvry, and Loic Didelot</i>	

Multiobjective Vehicle-type Scheduling in Urban Public Transport .....	482
<i>David Peña, Andrei Tchernykh, Sergio Nesmachnow, Renzo Massobrio, Alexander Feoktistov, and Igor Bychkov</i>	



### **Session 3: Parallel Metaheuristics and Machine Learning**

A new Co-evolutionary Algorithm Based on Constraint Decomposition .....	492
<i>Emmanuel Kieffer, Grégoire Danoy, Pascal Bouvry, and Anass Nagih</i>	
Training Many Neural Networks in Parallel via Back-Propagation .....	501
<i>Javier A. Cruz-López, Vincent Boyer, and Didier El-Baz</i>	
Design of Metaheuristic Based on Machine Learning: A Unified Approach .....	510
<i>Amir Nakib, Mohamed Hilia, Frederic Heliodore, and El-Ghazali Talbi</i>	

### **Session 4: Graphs, Networks and Algorithms**

Shared Memory Parallel Subgraph Enumeration .....	519
<i>Raphael Kimmig, Henning Meyerhenke, and Darren Strash</i>	
Exploration of de Bruijn Graph Filtering for <i>de novo</i> Assembly Using GraphLab .....	530
<i>Julien Collet, Tanguy Sassolas, Yves Lhuillier, Renaud Sirdey, and Jacques Carlier</i>	
An Efficient CPP Solution for Resilience-Oriented SDN Controller Deployment .....	540
<i>He Li, Robson Eduardo De Grande, and Azzedine Boukerche</i>	

### **Session 5: Parallel Algorithms**

Optimal Bandwidth Selection for Kernel Regression Using a Fast Grid Search and a GPU .....	550
<i>Chris Rohlfis and Mohamed Zahran</i>	
Space-Efficient Pointwise Computation of the Distance Transform on GPUs .....	557
<i>Numair Khan and Mohamed Zahran</i>	
Optimizing One-Sided Communication of Parallel Applications Using Critical Path Methods .....	567
<i>Christian Herold, Olaf Krzikalla, and Andreas Knüpfer</i>	

### **GABB: Graph Algorithms Building Blocks**

Introduction to GABB Workshop .....	577
<i>Aydın Buluç and Tim Mattson</i>	
GABB Keynote .....	578
<i>Ümit V. Çatalyürek</i>	

### **Session 1**

Breadth-First Search with A Multi-Core Computer .....	579
<i>Maryia Belova and Ming Ouyang</i>	
Order or Shuffle: Empirically Evaluating Vertex Order Impact on Parallel Graph Computations .....	588
<i>George M Slotá, Sivasankaran Rajamanickam, and Kamesh Madduri</i>	
A Study of Graph Decomposition Algorithms for Parallel Symmetry Breaking .....	598
<i>Sayyad Nayyaroddeen, Mahak Gambhir, and Kishore Kothapalli</i>	

## Session 2

Constructing Adjacency Arrays from Incidence Arrays .....	608
<i>Hayden Jananathan, Karia Dibert, and Jeremy Kepner</i>	
Mini-Gunrock: A Lightweight Graph Analytics Framework on the GPU .....	616
<i>Yangzihao Wang, Sean Baxter, and John D. Owens</i>	
Algebraic Multigrid for Least Squares Problems on Graphs with Applications to HodgeRank .....	627
<i>Charles Colley, Junyuan Lin, Xiaozhe Hu, and Shuchin Aeron</i>	

## Session 3

Deriving Streaming Graph Algorithms from Static Definitions .....	637
<i>David Ediger and James P. Fairbanks</i>	

## Session 4

Design of the GraphBLAS API for C .....	643
<i>Aydin Buluç, Tim Mattson, Scott McMillan, José Moreira, and Carl Yang</i>	
A Linear Algebra-Based Programming Interface for Graph Computations in Scala and Spark .....	653
<i>William Horn, Gabriel Tanase, Hao Yu, and Pratap Pattnaik</i>	

## AsHES: The Seventh International Workshop on Accelerators and Hybrid Exascale Systems

Introduction to AsHES Workshop .....	660
<i>Sunita Chandrasekaran</i>	
AsHES Keynote .....	661
<i>Tim Mattson</i>	

## Session 1: Programming Models and Runtime Systems

Implementing the OpenACC Data Model .....	662
<i>Michael Wolfe, Seyong Lee, Jungwon Kim, Xiaonan Tian, Rengan Xu, Sunita Chandrasekaran, and Barbara Chapman</i>	
Exploring Translation of OpenMP to OpenACC 2.5: Lessons Learned .....	673
<i>Sergio Pino, Lori Pollock, and Sunita Chandrasekaran</i>	
Exploring the Performance Benefit of Hybrid Memory System on HPC Environments .....	683
<i>Ivy Bo Peng, Roberto Gioiosa, Gokcen Kestor, Pietro Cicotti, Erwin Laure, and Stefano Markidis</i>	

## Session 2: Algorithms

Performance-Portable Sparse Matrix-Matrix Multiplication for Many-Core Architectures .....	693
<i>Mehmet Deveci, Christian Trott, and Sivasankaran Rajamanickam</i>	
Time and Energy to Solution Evaluation for the Three-Point Angular Correlation Function .....	703
<i>Antonio Gomez-Iglesias and Miguel Cardenas-Montes</i>	
Auto-Tuning Strategies for Parallelizing Sparse Matrix-Vector (SpMV) Multiplication on Multi- and Many-Core Processors .....	713
<i>Kaixi Hou, Wu-chun Feng, and Shuai Che</i>	

## Session 3: Scheduling and Architectures

A Pluggable Framework for Composable HPC Scheduling Libraries .....	723
<i>Max Grossman, Vivek Kumar, Nick Vrvilo, Zoran Budimlić, and Vivek Sarkar</i>	
Static Versus Dynamic Task Scheduling of the Lu Factorization on ARM big. LITTLE Architectures .....	733
<i>Sandra Catalán, Rafael Rodríguez-Sánchez, Enrique S. Quintana-Ortí, and José R. Herrero</i>	
Benchmarking SW26010 Many-Core Processor .....	743
<i>Zhigeng Xu, James Lin, and Satoshi Matsuoka</i>	

## HIPS: 22nd International Workshop on High Level Programming Models and Supportive Environments

Introduction to HIPS Workshop .....	753
<i>Bo Wu and Andreas Knüpfer</i>	
HIPS Keynote .....	755
<i>Zizhong Chen</i>	

### Session 1

Performance Study of Multithreaded MPI and OpenMP Tasking in a Large Scientific Code .....	756
<i>Dana Akhmetova, Roman Iakymchuk, Örjan Ekeberg, and Erwin Laure</i>	
Comparison of Threading Programming Models .....	766
<i>Solmaz Salehian, Jiawen Liu, and Yonghong Yan</i>	
Annotation-Based Parallelization of Java Code .....	775
<i>Mostafa Mehrabi, Nasser Giacaman, and Oliver Sinnen</i>	

### Session 2

Using LLVM for Optimized Lightweight Binary Re-Writing at Runtime .....	785
<i>Alexis Engelke and Josef Weidendorfer</i>	
Snowflake: A Lightweight Portable Stencil DSL .....	795
<i>Nathan Zhang, Michael Driscoll, Charles Markley, Samuel Williams, Protonu Basu, and Armando Fox</i>	

Enabling One-Sided Communication Semantics on ARM .....	805
<i>Pavel Shamis, M. Graham Lopez, and Gilad Shainer</i>	

### **Session 3**

Towards a Language Framework for Thick Control Flows .....	814
<i>Jari-Matti Mäkelä, Martti Forsell, and Ville Leppänen</i>	
Pure Concurrent Programming .....	824
<i>Benjamin J. L. Wang and Uwe R. Zimmer</i>	

### **APDCM: 19th Workshop on Advances in Parallel and Distributed Computational Models**

Introduction to APDCM Workshop .....	832
<i>Oscar H. Ibarra and Koji Nakano</i>	
APDCM Keynote .....	833
<i>Hong Shen</i>	

### **Session 1: Distributed Computing**

Complete Visibility for Mobile Agents with Lights Tolerating a Faulty Agent .....	834
<i>Aisha Aljohani and Gokarna Sharma</i>	
A Self-Stabilizing Algorithm for Constructing (1,1)-Maximal Directed Acyclic Graph .....	844
<i>Yonghwan Kim, Haruka Ohno, Yoshiaki Katayama, and Toshimitsu Masuzawa</i>	
Fault Tolerance for Cooperative Lifeline-Based Global Load Balancing in Java with APGAS and Hazelcast .....	854
<i>Jonas Posner and Claudia Fohry</i>	
Applications of Ear Decomposition to Efficient Heterogeneous Algorithms for Shortest Path/Cycle Problems .....	864
<i>Debarshi Dutta, Meher Chaitanya, Kishore Kothapalli, and Debajyoti Bera</i>	

### **Session 2: Scheduling and Hardware Models**

Co-Scheduling Algorithms for Cache-Partitioned Systems .....	874
<i>Guillaume Aupy, Anne Benoit, Loïc Pottier, Padma Raghavan, Yves Robert, and Manu Shantharam</i>	
Minimizing I/Os in Out-of-Core Task Tree Scheduling .....	884
<i>Loris Marchal, Samuel McCauley, Bertrand Simon, and Frédéric Vivien</i>	
Approximate Count and Queue Objects in Transactional Memory .....	894
<i>Basem Assiri and Costas Busch</i>	
Assessing NUMA Performance Based on Hardware Event Counters .....	904
<i>Max Plauth, Christoph Sterz, Felix Eberhardt, Frank Feinbube, and Andreas Polze</i>	

### **Session 3: Parallel Computing**

An Analysis of Resilience Techniques for Exascale Computing Platforms .....	914
<i>Daniel Dauwe, Sudeep Pasricha, Anthony A. Maciejewski, and Howard Jay Siegel</i>	
A Compression Method for Storage Formats of a Sparse Matrix in Solving the Large-Scale Linear Systems .....	924
<i>Tomoki Kawamura, Yoneda Kazunori, Takashi Yamazaki, Takashi Iwamura, Masahiro Watanabe, and Yasushi Inoguchi</i>	
Accelerating the Smith-Waterman Algorithm Using Bitwise Parallel Bulk Computation Technique on GPU .....	932
<i>Takahiro Nishimura, Jacir L. Bordim, Yasuaki Ito, and Koji Nakano</i>	
Photomosaic Generation by Rearranging Subimages, with GPU Acceleration .....	942
<i>Yi Yang, Yasuaki Ito, and Koji Nakano</i>	

### **HPPAC: 13th Workshop on High-Performance, Power-Aware Computing**

HPPAC Workshop Introduction .....	952
<i>Shuaiwen Leon Song and Richard Vuduc</i>	
HPPAC Keynote Talk .....	953
<i>Kirk W. Cameron</i>	

### **Session 1**

Using Machine Learning for Data Center Cooling Infrastructure Efficiency Prediction .....	954
<i>Hayk Shoukourian, Torsten Wilde, Detlef Labrenz, and Arndt Bode</i>	
Design of an Energy Aware Petaflops Class High Performance Cluster Based on Power Architecture .....	964
<i>Wissam Abu Ahmad, Andrea Bartolini, Francesco Beneventi, Luca Benini, Andrea Borghesi, Marco Cicala, Privato Forestieri, Cosimo Gianfreda, Daniele Gregori, Antonio Libri, Filippo Spiga, and Simone Tinti</i>	
Towards a Unified Monitoring Framework for Power, Performance and Thermal Metrics: A Case Study on the Evaluation of HPC Cooling Systems .....	974
<i>Aniruddha Marathe, Ghaleb Abdulla, Barry L. Rountree, and Kathleen Shoga</i>	

### **Session 2**

When Good Enough Is Better: Energy-Aware Scheduling for Multicore Servers .....	984
<i>Xinning Hui, Zhihui Du, Jason Liu, Hongyang Sun, Yuxiong He, and David A. Bader</i>	
A Runtime Workload Distribution with Resource Allocation for CPU-GPU Heterogeneous Systems .....	994
<i>Shouq Alsubaihi and Jean-Luc Gaudiot</i>	

### Session 3

Power Measurements of Hartree-Fock Algorithms Using Different Storage Devices .....	1004
<i>Vladimir Mironov, Alexander Moskovsky, and Yuri Alexeev</i>	
A Statistical Approach to Power Estimation for x86 Processors .....	1012
<i>Mohak Chadha, Thomas Ilsche, Mario Bielert, and Wolfgang E. Nagel</i>	

### HPBDC: 3rd IEEE International Workshop on High-Performance Big Data Computing

Introduction to HPBDC Workshop .....	1020
<i>Xiaoyi Lu, Jianfeng Zhan, and Dhabaleswar K. (DK) Panda</i>	

### Session 1: High-Performance Graph Processing

Performance Evaluation of Scale-Free Graph Algorithms in Low Latency Non-volatile Memory .....	1021
<i>Manu Shantharam, Keita Iwabuchi, Pietro Cicotti, Laura Carrington, Maya Gokhale, and Roger Pearce</i>	
High-Performance Data Analytics Beyond the Relational and Graph Data Models with GEMS .....	1029
<i>Vito Giovanni Castellana, Marco Minutoli, Shreyansh Bhatt, Khushbu Agarwal, Arthur Bleeker, John Feo, Daniel Chavarria-Miranda, and David Haglin</i>	
Graph Analytics: Complexity, Scalability, and Architectures .....	1039
<i>Peter M. Kogge</i>	

### Session 2: Benchmarking and Performance Analysis

Spark and HPC for High Energy Physics Data Analyses .....	1048
<i>Saba Sehrish, Jim Kowalkowski, and Marc Paterno</i>	
The Consistency Analysis of Secondary Index on Distributed Ordered Tables .....	1058
<i>Houliang Qi, Xu Chang, Xingwu Liu, and Li Zha</i>	
BigDataBench-S: An Open-Source Scientific Big Data Benchmark Suite .....	1068
<i>Xinhui Tian, Shaopeng Dai, Zhihui Du, Wanling Gao, Rui Ren, Yaodong Cheng, Zhifei Zhang, Zhen Jia, Peijian Wang, and Jianfeng Zhan</i>	
Scalable Architecture for Anomaly Detection and Visualization in Power Generating Assets .....	1078
<i>Paras Jain, Chirag Tailor, Sam Ford, Liexiao Ding, Michael Phillips, Fang Liu, Nagi Gebraeel, and Duen Horng Chau</i>	

### CHIUW: The Fourth Annual Chapel Implementers and Users Workshop

Introduction to CHIUW Workshop .....	1083
<i>Tom MacDonald and Michael Ferguson</i>	
CHIUW Keynote .....	1085
<i>Jonathan Dursi</i>	

Identifying Use-After-Free Variables in Fire-and-Forget Tasks .....	1086
<i>Jyothi Krishna V S and Vassily Litvinov</i>	
Towards a GraphBLAS Library in Chapel .....	1095
<i>Ariful Azad and Aydin Buluc</i>	
Comparative Performance and Optimization of Chapel in Modern Manycore Architectures .....	1105
<i>Engin Kayraklioglu, Wo Chang, and Tarek El-Ghazawi</i>	

## **PDSEC: 18th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing**

Introduction to PDSEC Workshop .....	1115
<i>Peter Strazdins, Keita Teranishi, Raphaël Couturier, Joseph Antony, Thomas Rauber, Gudula Rünger, and Laurence T. Yang</i>	
PDSEC Keynote .....	1117
<i>Pavan Balaji</i>	

### **Session 1: Best Paper**

Improving Performance of GMRES by Reducing Communication and Pipelining Global Collectives .....	1118
<i>Ichitaro Yamazaki, Mark Hoemmen, Piotr Luszczek, and Jack Dongarra</i>	

### **Session 2: Linear Algebra**

Simultaneously Solving Swarms of Small Sparse Systems on SIMD Silicon .....	1128
<i>Bryce Adelstein LeBach, Hans Johansen, and Samuel Williams</i>	
Sparse Supernodal Solver Using Block Low-Rank Compression .....	1138
<i>Gregoire Pichon, Eric Darve, Mathieu Faverge, Pierre Ramet, and Jean Roman</i>	
Task-Parallel LU Factorization of Hierarchical Matrices Using OmpSs .....	1148
<i>José I. Aliaga, Rocío Carratalá-Sáez, Ronald Kriemann, and Enrique S. Quintana-Ortí</i>	

### **Session 3: Applications**

Parallel Particle-in-Cell Performance Optimization: A Case Study of Electrospray Simulation .....	1158
<i>Ramachandran Kodanganallur Narayanan and Kamesh Madduri</i>	
Efficient Data Structures for a Hybrid Parallel and Vectorized Particle-in-Cell Code .....	1168
<i>Yann Barsamian, Sever A. Hirstoaga, and Éric Violard</i>	
A Locality-Based Threading Algorithm for the Configuration-Interaction Method .....	1178
<i>Hongzhang Shan, Samuel Williams, Calvin Johnson, and Kenneth McElvain</i>	
Architecting the Discontinuous Deformation Analysis Method Pipeline on the GPU .....	1188
<i>Yunfan Xiao, Min Huang, Qinghai Miao, Jun Xiao, and Ying Wang</i>	

## Session 4: Parallel Techniques

Redesigning OP2 Compiler to Use HPX Runtime Asynchronous Techniques .....	1198
<i>Zahra Khatami, Hartmut Kaiser, and J. Ramanujam</i>	
Automated Dynamic Data Redistribution .....	1208
<i>Thomas Marrinan, Joseph A. Insley, Silvio Rizzi, François Tessier, and Michael E. Papka</i>	
An Application-Aware Data Replacement Policy for Interactive Large-Scale Scientific Visualization .....	1216
<i>Lina Yu, Hongfeng Yu, Hong Jiang, and Jun Wang</i>	
Scalable Hierarchical Multipole Methods Using an Asynchronous Many-Tasking Runtime System .....	1226
<i>Jackson DeBuhr, Bo Zhang, and Luke D'Alessandro</i>	

## JSSPP: 21st Workshop on Job Scheduling Strategies for Parallel Processing

Introduction to JSSPP Workshop .....	1235
<i>Walfredo Cirne, Narayan Desai, and Dalibor Klusáček</i>	

## DPDNS: 22nd IEEE Workshop on Dependable Parallel, Distributed and Network-Centric Systems

Introduction to DPDNS Workshop .....	1237
<i>Dimiter Avresky and Erik Maehle</i>	

## Session 1

Reliability Calculation of P2P Streaming Systems with Bottleneck Links .....	1238
<i>Satoshi Fujita</i>	
Lifetime and Full-View Coverage Guarantees Through Distributed Algorithms in Camera Sensor Networks .....	1245
<i>Chaoyang Li and Anu G. Bourgeois</i>	

## Session 2

A Small-Scale Testbed for Large-Scale Reliable Computing .....	1251
<i>Jason St. John and Thomas J. Hacker</i>	
LSTM-Based Memory Profiling for Predicting Data Attacks in Distributed Big Data Systems .....	1259
<i>Santosh Aditham, Nagarajan Ranganathan, and Srinivas Katkoori</i>	
An Outlook on Volunteer and Crowdsourcing Based Computing .....	1268
<i>Salvatore Distefano and Samuele Rodi</i>	
Exploring the Effect of Compiler Optimizations on the Reliability of HPC Applications .....	1274
<i>Rizwan A. Ashraf, Roberto Gioiosa, Gokcen Kestor, and Ronald F. DeMara</i>	



## **IPDRM: Second Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware**

IPDRM Workshop Introduction .....	1284
<i>Shuaiwen Leon Song and Torsten Hoefler</i>	

### **Session 1**

Characterizing and Improving the Performance of Many-Core Task-Based Parallel Programming Runtimes .....	1285
<i>Jaume Bosch, Xubin Tan, Carlos Álvarez, Daniel Jiménez-González, Xavier Martorell, and Eduard Ayguadé</i>	
A Memory Heterogeneity-Aware Runtime System for Bandwidth-Sensitive HPC Applications .....	1293
<i>Kavitha Chandrasekar, Xiang Ni, and Laxmikant V. Kale</i>	
SmartBlock: An Approach to Standardizing In Situ Workflow Components .....	1301
<i>Alexis Champsaur, Jay Lofstead, Jai Dayal, Matthew Wolf, Greg Eisenhauer, Patrick Widener, and Ada Gavrilovska</i>	

### **Session 2**

A Case Study in Computational Caching Microservices for HPC .....	1309
<i>John Jenkins, Galen Shipman, Jamaludin Mohd-Yusof, Kipton Barros, Philip Carns, and Robert Ross</i>	
A Load-Balanced Parallel and Distributed Sorting Algorithm Implemented with PGX.D .....	1317
<i>Zahra Khatami, Sungpack Hong, Jinsoo Lee, Siegfried Depner, Hassan Chafi, J. Ramanujam, and Hartmut Kaiser</i>	

### **Session 3**

Performance Prediction of HPC Applications on Intel Processors .....	1325
<i>Carlos Rosales, Antonio Gómez-Iglesias, Si Liu, Feng Chen, Lei Huang, Hang Liu, Antia Lamas-Linares, and John Cazes</i>	
vPHI: Enabling Xeon Phi Capabilities in Virtual Machines .....	1333
<i>Stefanos Gerangelos and Nectarios Koziris</i>	

## **iWAPT: 12th International Workshop on Automatic Performance Tuning**

Introduction to iWAPT Workshop .....	1341
<i>Osni Marques and Reiji Suda</i>	

### **Session 1: New Methodology of Auto-Tuning**

A Sampling Based Strategy to Automatic Performance Tuning of GPU Programs .....	1342
<i>Wilson Feng and Tarek S. Abdelrahman</i>	
Use of Synthetic Benchmarks for Machine-Learning-Based Performance Auto-Tuning .....	1350
<i>Tianyi David Han and Tarek S. Abdelrahman</i>	

## **Session 2: Auto-Tuning Software and Environment**

Automating Compiler-Directed Autotuning for Phased Performance Behavior .....	1362
<i>Tharindu Rusira, Mary Hall, and Protonu Basu</i>	
A Customizable Auto-Tuning Scenario with User-Defined Code Transformations .....	1372
<i>HiroYuki Takizawa, Daichi Sato, Shoichi Hirasawa, and Daisuke Takahashi</i>	
Online-Autotuning in the Presence of Algorithmic Choice .....	1379
<i>Philip Pfafe, Martin Tillmann, Sigmar Walter, and Walter F. Tichy</i>	

## **Session 3: Case-Study of Auto-Tuning and Optimization**

Performance Analysis and Optimization of Sparse Matrix-Vector Multiplication on Intel Xeon Phi .....	1389
<i>Athena Elafrou, Georgios Goumas, and Nectarios Koziris</i>	
Auto-Tuning on NUMA and Many-Core Environments with an FDM Code .....	1399
<i>Takahiro Katagiri, Satoshi Ohshima, and Masaharu Matsumoto</i>	
Autotuning Batch Cholesky Factorization in CUDA with Interleaved Layout of Matrices .....	1408
<i>Mark Gates, Jakub Kurzak, Piotr Luszczek, Yu Pei, and Jack Dongarra</i>	

## **Session 4: Scientific Applications by Auto-Tuning**

Quadruple-Precision BLAS Using Bailey's Arithmetic with FMA Instruction: Its Performance and Applications .....	1418
<i>Susumu Yamada, Toshiyuki Imamura, Takuya Ina, Narimasa Sasa, Yasuhiro Idomura, and Masahiko Machida</i>	
Fast Multidimensional Performance Parameter Estimation with Multiple One-Dimensional d-Spline Parameter Search .....	1426
<i>Masayoshi Mochizuki, Akihiro Fujii, and Teruo Tanaka</i>	
Algorithmic Performance-Accuracy Trade-off in 3D Vision Applications Using HyperMapper .....	1434
<i>Luigi Nardi, Bruno Bodin, Sajad Saeedi, Emanuele Vespa, Andrew J. Davison, and Paul H. J. Kelly</i>	

## **ParSocial: 2nd IEEE Workshop on Parallel and Distributed Processing for Computational Social System**

Introduction to ParSocial Workshop .....	1444
<i>Eunice E. Santos and John Korah</i>	
ParSocial Keynote .....	1446
<i>Boleslaw Szymanski</i>	

## Session 1

Predicting Viral News Events in Online Media .....	1447
<i>Xiaoyan Lu and Boleslaw Szymanski</i>	
Mobile Crowdsensing from a Selfish Routing Perspective .....	1457
<i>Julia Buwaya and José D. P. Rolim</i>	
Parallel Computing for Machine Learning in Social Network Analysis .....	1464
<i>George Cybenko</i>	

## Session 2

Work Partitioning on Parallel and Distributed Agent-Based Simulation .....	1472
<i>Gennaro Cordasco, Carmine Spagnuolo, and Vittorio Scarano</i>	
Parallel k-Core Decomposition on Multicore Platforms .....	1482
<i>Humayun Kabir and Kamesh Madduri</i>	
Endogenous Social Networks from Large-Scale Agent-Based Models .....	1492
<i>Eric Tatara, Nicholson Collier, Jonathan Ozik, and Charles Macal</i>	

## Session 3

Fast Parallel Graph Triad Census and Triangle Counting on Shared-Memory Platforms .....	1500
<i>Sindhuja Parimalarangan, George M. Slota, and Kamesh Madduri</i>	
Efficient Anytime Anywhere Algorithms for Vertex Additions in Large and Dynamic Graphs .....	1510
<i>Eunice E. Santos, John Korah, Vairavan Murugappan, and Suresh Subramanian</i>	
Accelerating Topic Exploration of Multi-Dimensional Documents .....	1520
<i>Hsu Wen-Jing, Lu You, and Lee Zhuo Qi</i>	

## BigDataEco: Big Data Regional Innovation Hubs and Spokes Workshop

Introduction to BigDataEco Workshop .....	1528
<i>Chaitan Baru, Fen Zhao, and Joanna Chan</i>	

## GraML: First Workshop on the Intersection of Graph Algorithms and Machine Learning

Introduction to GraML Workshop .....	1529
<i>Antonino Tumeo, Mahantesh Halappanavar, and John Feo</i>	
GraML Keynote .....	1531
<i>Sujith Ravi</i>	
Learning on Graphs for Predictions of Fracture Propagation, Flow and Transport .....	1532
<i>Hristo Djidjev, Daniel O'Malley, Hari Viswanathan, Jeffrey Hyman, Satish Karra, and Gowri Srinivasan</i>	
Analyzing Community Structure in Networks .....	N/A
<i>Hongyuan Zhan and Kamesh Madduri</i>	

Compound Analytics: Templates for Integrating Graph Algorithms and Machine Learning .....	1550
<i>Ronald D. Hagan, Charles A. Phillips, Bradley J. Rhodes, and Michael A. Langston</i>	

## **EMBRACE: Evolvable Methods for Benchmarking Realism and Community Engagement**

Introduction to EMBRACE Workshop .....	1557
<i>David Bader</i>	
EMBRACE Keynote .....	1558
<i>Torsten Hoefler</i>	

## **REPPAR: Workshop on Reproducibility in Parallel Computing**

Introduction to REPPAR Workshop .....	1559
<i>Sascha Hunold, Arnaud Legrand, and Lucas Nussbaum</i>	
REPPAR Keynote .....	1560
<i>Todd Gamblin</i>	

### **Session 1**

The Popper Convention: Making Reproducible Systems Evaluation Practical .....	1561
<i>Ivo Jimenez, Michael Sevilla, Noah Watkins, Carlos Maltzahn, Jay Lofstead, Kathryn Mohror, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau</i>	
Towards Trustworthy Testbeds Thanks to Throughout Testing .....	1571
<i>Lucas Nussbaum</i>	
Examining the Reproducibility of Using Dynamic Loop Scheduling Techniques in Scientific Applications .....	1579
<i>Franziska Hoffeins, Florina M. Ciorba, and Ioana Banicescu</i>	

### **Session 2**

Characterizing the Performance of Modern Architectures Through Opaque Benchmarks: Pitfalls Learned the Hard Way .....	1588
<i>Luka Stanisic, Lucas Mello Schnorr, Augustin Degomme, Franz C. Heinrich, Arnaud Legrand, and Brice Videau</i>	
Towards Reproducible Blocked LU Factorization .....	1598
<i>Roman Iakymchuk, Enrique S. Quintana-Ortí, Erwin Laure, and Stef Graillat</i>	

### **Author Index**