

2014 IEEE International Parallel & Distributed Processing Symposium Workshops (IPDPSW 2014)

**Phoenix, Arizona, USA
19 - 23 May 2014**

Pages 1-874



**IEEE Catalog Number: CFP1451J-POD
ISBN: 978-1-4799-4115-5**

2014 IEEE 28th International Parallel & Distributed Processing Symposium Workshops

IPDPSW 2014

Table of Contents

Message from the General Chair.....	xxiii
Message from the Workshops Chair.....	xxv

Workshop 1: HCW — Heterogeneity in Computing Workshop

HCW Introduction	1
<i>Behrooz Shirazi and Uwe Schwiegelshohn</i>	
Message from the HCW Steering Committee Chair	3
<i>Behrooz Shirazi</i>	
Message from the HCW General Chair	4
<i>Uwe Schwiegelshohn</i>	
Message from the HCW Program Chair	5
<i>Shoukat Ali</i>	
HCW 2014 Keynote Talk	6
<i>David Abramson</i>	

HCW Session 1: Heterogeneous Environments for Basic Linear Algebra

Hybrid Multi-elimination ILU Preconditioners on GPUs	7
<i>Dimitar Lukarski, Hartwig Anzt, Stanimire Tomov, and Jack Dongarra</i>	
Searching for the Optimal Data Partitioning Shape for Parallel Matrix Matrix Multiplication on 3 Heterogeneous Processors	17
<i>Ashley DeFlumere and Alexey Lastovetsky</i>	
Taking Advantage of Hybrid Systems for Sparse Direct Solvers via Task-Based Runtimes	29
<i>Xavier Lacoste, Mathieu Faverge, George Bosilca, Pierre Ramet, and Samuel Thibault</i>	
Topology-Aware Optimization of Communications for Parallel Matrix Multiplication on Hierarchical Heterogeneous HPC Platform	39
<i>Tania Malik, Vladimir Rychkov, Alexey Lastovetsky, and Jean-Noël Quintin</i>	

HCW Session 2: Scheduling and Resource Allocation

Scheduling Methods for Accelerating Applications on Architectures with Heterogeneous Cores	48
<i>Linchuan Chen, Xin Huo, and Gagan Agrawal</i>	
Utility Driven Dynamic Resource Management in an Oversubscribed Energy-Constrained Heterogeneous System	58
<i>Bhavesh Khemka, Ryan Friese, Sudeep Pasricha, Anthony A. Maciejewski, Howard Jay Siegel, Gregory A. Koenig, Sarah Powers, Marcia Hilton, Rajendra Rambharos, and Steve Poole</i>	
An Efficient Algorithm for Scheduling Jobs in Volunteer Computing Platforms	68
<i>Adel Essafi, Denis Trystram, and Zied Zaidi</i>	

HCW Session 3: Resource-Related Performance Optimization

Resource Centered Computing Delivering High Parallel Performance	77
<i>Jens Gustedt, Stephane Vialle, and Patrick Mercier</i>	
Point-to-Point and Congestion Bandwidth Estimation: Experimental Evaluation on PlanetLab Data	89
<i>Lionel Eyraud-Dubois and Przemysław Uznański</i>	
Runtime Behavior Comparison of Modern Accelerators and Coprocessors	97
<i>Ayman Tarakji and Niels Ole Salscheider</i>	

Workshop 2: RAW — Reconfigurable Architectures Workshop

RAW Introduction and Committees	109
<i>Jürgen Becker, Ramachandran Vaidyanathan, Marco Santambrogio, Jim Tørresen, Ron Sass, and Philip Leong</i>	
RAW 2014 Keynotes	111
<i>Joshua Walstrom and Maya Gokhale</i>	

RAW Session 1: Compilers and Binary Translation for Reconfigurable Architectures

Twill: A Hybrid Microcontroller-FPGA Framework for Parallelizing Single-Threaded C Programs	112
<i>Doug Gallatin, Aaron Keen, Chris Lupo, and John Oliver</i>	
A New Dataflow Compiler IR for Accelerating Control-Intensive Code in Spatial Hardware	122
<i>Ali Mustafa Zaidi and David Greaves</i>	
Efficient Software-Based Runtime Binary Translation for Coarse-Grained Reconfigurable Architectures	132
<i>Toan X. Mai and Jongeun Lee</i>	

RAW Session 2: New Reconfigurable Architectures

A Dependable Coarse-Grain Reconfigurable Multicore Array	141
<i>Georgios Smaragdos, Danish Anis Khan, Ioannis Sourdis, Christos Strydis, Alirad Malek, and Stavros Tzilis</i>	
Automated Hybrid Interconnect Design for FPGA Accelerators Using Data Communication Profiling	151
<i>Cuong Pham-Quoc, Zaid Al-Ars, and Koen Bertels</i>	
SmartBricks: A Visual Environment to Design and Explore Novel Custom Domain-Specific Architectures	161
<i>AnilKumar Sistla, Xiaozhong Luo, Mukund Malladi, Marc Reisner, Rajasekhar Ganduri, and Gayatri Mehta</i>	

RAW Session 3: ViPES Papers

A Framework for Mapping Dynamic Virtual Kernels onto Heterogeneous Reconfigurable Platforms	170
<i>Harry Sidiropoulos, Kostas Siozios, and Dimitrios Soudris</i>	
A Hybrid ILP-CP Model for Mapping Directed Acyclic Task Graphs to Multicore Architectures	176
<i>Andreas Emeretlis, George Theodoridis, Panayiotis Alefragis, and Nikolaos Voros</i>	
A Framework for Customizing Virtual 3-D Reconfigurable Platforms at Run-Time	183
<i>Kostas Siozios, Dimitrios Soudris, and Michael Hübner</i>	

RAW Session 4: Circuit-Level Applications

Over-clocking of Linear Projection Designs through Device Specific Optimisations	189
<i>Rui Policarpo Duarte and Christos-Savvas Bouganis</i>	
Influence of Magnetic Fields and X-Radiation on Ring Oscillators in FPGAs	199
<i>Michael Raitza, Markus Vogt, Christian Hochberger, and Thilo Pionteck</i>	
Radiation Tolerance of Color Configuration on an Optically Reconfigurable Gate Array	205
<i>Takumi Fujimori and Minoru Watanabe</i>	

RAW Session 5: Numerical Reconfigurable Computing Applications

Adaptive Booth Algorithm for Three-Integers Multiplication for Reconfigurable Mesh	211
<i>Esti Stein and Yosi Ben Asher</i>	
An FPGA Implementation of the Hestenes-Jacobi Algorithm for Singular Value Decomposition	220
<i>Xinying Wang and Joseph Zambreno</i>	

RAW Session 6: Applications of Reconfigurable Computing

CyGraph: A Reconfigurable Architecture for Parallel Breadth-First Search	228
<i>Osama G. Attia, Tyler Johnson, Kevin Townsend, Philip Jones, and Joseph Zambreno</i>	
Adaptive Raytracing Implementation Using Partial Dynamic Reconfiguration	236
<i>Gianluca Durelli, Fabrizio Spada, Riccardo Cattaneo, Christian Pilato, Danilo Pau, and Marco D. Santambrogio</i>	

PaRA-Sched: A Reconfiguration-Aware Scheduler for Reconfigurable Architectures	243
<i>Riccardo Cattaneo, Riccardo Bellini, Gianluca Durelli, Christian Pilato, Marco D. Santambrogio, and Donatella Sciuto</i>	

RAW Poster Session 1

An ILP-Based Optimal Circuit Mapping Method for PLDs	251
<i>Hiroki Nishiyama, Masato Inagi, Shin'ichi Wakabayashi, Shinobu Nagayama, Keisuke Inoue, and Mineo Kaneko</i>	
High-Level Synthesis from C vs. a DSL-Based Approach	257
<i>Cristiano Bacelar de Oliveira, João M. P. Cardoso, and Eduardo Marques</i>	
An Evaluation of User Satisfaction Driven Scheduling in a Polymorphic Embedded System	263
<i>Zhang Zhang, Swamy D. Ponpandi, and Akhilesh Tyagi</i>	
A Low-Latency Algorithm and FPGA Design for the Min-Search of LDPC Decoders	269
<i>Georgios Tzimpragos, Christoforos Kachris, Dimitrios Soudris, and Ioannis Tomkos</i>	
FPGA Redundancy Configurations: An Automated Design Space Exploration	275
<i>Jahanzeb Anwer, Marco Platzner, and Sebastian Meisner</i>	

RAW Poster Session 2

Hierarchical Pipeline Optimization of Coarse Grained Reconfigurable Processor for Multimedia Applications	281
<i>Chen Mei, Peng Cao, Yang Zhang, Bo Liu, and Leibo Liu</i>	
Module Placement Using Constraint Programming in Run-Time Reconfigurable Systems	287
<i>Alexander Wold, Andreas Agne, and Jim Torresen</i>	
An Efficient Heterogeneous Register File Implementation for FPGAs	293
<i>Hasan Erdem Yantir and Arda Yurdakul</i>	
Minimizing Scrubbing Effort through Automatic Netlist Partitioning and Floorplanning	299
<i>Bernhard Schmidt, Daniel Ziener, and Jürgen Teich</i>	
Virtualization Support for FPGA-Based Coprocessors Connected via PCI Express to an Intel Multicore Platform	305
<i>Duy Viet Vu, Timo Sandmann, Steffen Baehr, Oliver Sander, and Juergen Becker</i>	

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

HIPS Introduction and Committees	311
<i>John Cavazos</i>	

HIPS Session 1: System Support

Bohrium: A Virtual Machine Approach to Portable Parallelism	312
<i>Mads R.B. Kristensen, Simon A.F. Lund, Troels Blum, Kenneth Skovhede, and Brian Vinter</i>	
HATI: Hardware Assisted Thread Isolation for Concurrent C/C++ Programs	322
<i>Juan Carlos Martínez Santos and Yunsi Fei</i>	

A General Model Checking Framework for Various Memory Consistency Models	332
<i>Tatsuya Abe and Toshiyuki Maeda</i>	
HIPS Session 2: Optimization	
Autotuning Tensor Transposition	342
<i>Lai Wei and John Mellor-Crummey</i>	
Automatic MPI-IO Tuning with the Periscope Tuning Framework	352
<i>Weifeng Liu, Isaiás A. Comprés Ureña, Michael Gerndt, and Bin Gong</i>	
Optimizing Collective Communication in UPC	361
<i>Jithin Jose, Khaled Hamidouche, Jie Zhang, Akshay Venkatesh, and Dhabaleswar K. (DK) Panda</i>	
HIPS Session 3: Effective Communication	
SWIFT: A Transparent and Flexible Communication Layer for PCIe-Coupled Accelerators and (Co-)Processors	371
<i>Simon Pickartz, Pablo Reble, Carsten Clauss, and Stefan Lankes</i>	
Deterministic Synchronization of Multi-threaded Programs with Operational Transformation	381
<i>Christopher Boelmann, Lorenz Schwittmann, and Torben Weis</i>	
ABC2: Adaptively Balancing Computation and Communication in a DSM Cluster of Multicores for Irregular Applications	391
<i>Sai Charan Koduru, Keval Vora, and Rajiv Gupta</i>	
Workshop 4: NIDISC — Workshop on Nature Inspired Distributed Computing	
NIDISC Introduction and Committees	401
<i>Pascal Bouvry, Franciszek Seredynski, and El-Ghazali Talbi</i>	
NIDISC Session 1: Applications of Bio-Inspired Algorithms	
Using Physical Stigmergy in Decentralized Optimization under Multiple Non-separable Constraints: Formal Methods and an Intelligent Lighting Example	402
<i>Theodore P. Pavlic</i>	
Hybrid Metaheuristic for Annual Hydropower Generation Optimization	412
<i>A. Nakib, El-Ghazali Talbi, and A. Fuser</i>	
Machine-Learning-Based Identification of Defect Patterns in Semiconductor Wafer Maps: An Overview and Proposal	420
<i>Fatima Adly, Paul D. Yoo, Sami Muhaidat, and Yousof Al-Hammadi</i>	
Data Quality, Consistency, and Interpretation Management for Wind Farms by Using Neural Networks	430
<i>Alain Fuser, Florent Fontaine, and Jack Copper</i>	

NIDISC Session 2: Wireless Networks and Mobility Management

Graph-Based Cellular Automata Approach to Maximum Lifetime Coverage Problem in Wireless Sensor Networks	439
<i>Antonina Tretyakova, Franciszek Seredynski, and Pascal Bouvry</i>	
GPU Accelerated Nature Inspired Methods for Modelling Large Scale Bi-directional Pedestrian Movement	448
<i>Sankha Baran Dutta, Robert McLeod, and Marcia Friesen</i>	
Improving Bus Ride Comfort Using GLOSA-Based Dynamic Speed Optimisation	457
<i>Marcin Seredynski, Patricia Ruiz, Krzysztof Szczypiorski, and Djamel Khadraoui</i>	
A Genetic Algorithm-Based Sparse Coverage over Urban VANETs	464
<i>Huang Cheng, Xin Fei, Azzedine Boukerche, and Mohammed Almulla</i>	

NIDISC Session 3: Multi-objective Optimization

A Game-Theoretic Approach to Multiobjective Job Scheduling in Cloud Computing Systems	470
<i>Jakub Gasiór and Franciszek Seredynski</i>	
Multi-level and Multi-objective Survey on Cloud Scheduling	480
<i>Yacine Kessaci, Nouredine Melab, and El-Ghazali Talbi</i>	
Comparison of Multi-objective Optimization Algorithms for the JShadObf JavaScript Obfuscator	489
<i>Benoît Bertholon, Sébastien Varrette, and Pascal Bouvry</i>	

Workshop 5: HiCOMB — Workshop on High Performance Computational Biology

HiCOMB Introduction and Committees	497
<i>Alba Cristina Magalhaes Alves de Melo, Srinivas Aluru, and David A. Bader</i>	
HiCOMB Keynote and Invited Talks	499
<i>Stephen Larson, Ümit V. Çatalyürek, and Ananth Kalyanaraman</i>	

HiCOMB Session 1: Parallel Algorithms for Biological Sequence Analysis

Constructing Similarity Graphs from Large-Scale Biological Sequence Collections	500
<i>Jaroslav Zola</i>	
Removing Sequential Bottlenecks in Analysis of Next-Generation Sequencing Data	508
<i>Yi Wang, Gagan Agrawal, Gulcin Ozer, and Kun Huang</i>	

HiCOMB Session 2: Parallel/Distributed Architectures for Biological Applications

Efficient Computation of the Phylogenetic Likelihood Function on the Intel MIC Architecture	518
<i>Alexey M. Kozlov, Christian Goll, and Alexandros Stamatakis</i>	
Process Simulation of Complex Biochemical Pathways in Explicit 3D Space Enabled by Heterogeneous Computing Platform	528
<i>Jie Li, Amin Salighehdar, and Narayan Ganesan</i>	

Exploring Large Scale Receptor-Ligand Pairs in Molecular Docking Workflows in HPC Clouds	536
<i>Kary Ocaña, Silvia Benza, Daniel de Oliveira, Jonas Dias, and Marta Mattoso</i>	
A Comparison of a Campus Cluster and Open Science Grid Platforms for Protein-Guided Assembly Using Pegasus Workflow Management System	546
<i>Natasha Pavlovikj, Kevin Begcy, Sairam Behera, Malachy Campbell, Harkamal Walia, and Jitender S. Deogun</i>	
HiCOMB Session 3: Metagenomics and Assembly	
Design and Optimization of a Metagenomics Analysis Workflow for NVRAM	556
<i>Sasha Ames, Jonathan E. Allen, David A. Hysom, G. Scott Lloyd, and Maya B. Gokhale</i>	
Parallelization of the Trinity Pipeline for De Novo Transcriptome Assembly	566
<i>V. Sachdeva, C.S. Kim, K.E. Jordan, and M.D. Winn</i>	
HiPGA: A High Performance Genome Assembler for Short Read Sequence Data	576
<i>Xiaohui Duan, Kun Zhao, and Weiguo Liu</i>	
Workshop 6: APDCM — Advances in Parallel and Distributed Computing Models	
APDCM Introduction and Committees	585
<i>Oscar H. Ibarra</i>	
APDCM Session 1	
Bulk Execution of Oblivious Algorithms on the Unified Memory Machine, with GPU Implementation	586
<i>Kazuya Tani, Daisuke Takafuji, Koji Nakano, and Yasuaki Ito</i>	
A Linear Performance-Breakdown Model for GPU Programming Optimization Guidance	596
<i>Mario A. Chapa M. and Sato Hiroyuki</i>	
A Hybrid Parallel Tridiagonal Solver on Multi-core Architectures	604
<i>Guangping Tang, Kenli Li, Keqin Li, Hang Chen, and Jiayi Du</i>	
A Novel Computational Model for GPUs with Application to I/O Optimal Sorting Algorithms	614
<i>Atsushi Koike and Kunihiko Sadakane</i>	
Predicting Cache Contention for Multithread Applications at Compile Time	624
<i>Munara Tolubaeva, Yonghong Yan, and Barbara Chapman</i>	
APDCM Session 2	
Parallelism Extraction Algorithm from Stream-Based Processing Flow Applying Spanning Tree	632
<i>Guyue Wang, Shinichi Yamagiwa, and Koichi Wada</i>	
EEWA: Energy-Efficient Workload-Aware Task Scheduling in Multi-core Architectures	642
<i>Quan Chen, Long Zheng, Minyi Guo, and Zhiyi Huang</i>	
A Platform-Specific Code Smell Alert System for High Performance Computing Applications	652
<i>Chunyan Wang, Shoichi Hirasawa, Hiroyuki Takizawa, and Hiroaki Kobayashi</i>	
Optimizing Buffer Sizes for Pipeline Workflow Scheduling with Setup Times	662
<i>Anne Benoit, Jean-Marc Nicod, and Veronika Rehn-Sonigo</i>	

WECPAR: List Ranking Algorithm and Relative Computational Power	671
<i>Hatem M. El-Boghdadi</i>	

APDCM Session 3

Assessing the Impact of ABFT and Checkpoint Composite Strategies	679
<i>George Bosilca, Aurelien Bouteiller, Thomas Herault, Yves Robert, and Jack Dongarra</i>	
Memory-Aware List Scheduling for Hybrid Platforms	689
<i>Julien Herrmann, Loris Marchal, and Yves Robert</i>	
A Parallel Framework for Handling Non-determinism with Expressive Description Logics	699
<i>Jocelyne Faddoul and Wendy MacCaul</i>	
Prototyping the MBTAC Processor for the REPLICIA CMP	709
<i>Martti Forsell, Jussi Roivainen, and Ville Leppänen</i>	
Evaluation of the Global Address Space Programming Interface (GASPI)	717
<i>Jens Breitbart, Mareike Schmidtobreich, and Vincent Heuveline</i>	

APDCM Session 4

GPS: Towards Simplified Communication on SGL Model	727
<i>Chong Li and Gaétan Hains</i>	
Near-Optimal Location Tracking Using Sensor Networks	737
<i>Gokarna Sharma, Hari Krishnan, Costas Busch, and Steven R. Brandt</i>	
Self-Stabilizing Algorithm for Maximal 2-Packing with Safe Convergence in an Arbitrary Graph	747
<i>Yihua Ding, James Z. Wang, and Pradip K. Srimani</i>	
Minimum Set Cover of Sparsely Distributed Sensor Nodes by a Collection of Unit Disks	755
<i>Satoshi Fujita</i>	
An Efficient Implementation of the Gradient-Based Hough Transform Using DSP Slices and Block RAMs on the FPGA	762
<i>Xin Zhou, Yasuaki Ito, and Koji Nakano</i>	

Workshop 7: HPPAC — High-Performance, Power-Aware Computing

HPPAC Introduction and Committees	771
<i>Dong Li and Robert J. Fowler</i>	

HPPAC Session 1: Power and Energy Analysis and Profiling

Characterizing the Impact of Program Optimizations on Power and Energy for Explicit Hydrodynamics	773
<i>Edgar A. León and Ian Karlin</i>	
Application Power Signature Analysis	782
<i>Chung-Hsing Hsu, Jacob Combs, Jolie Nazor, Fabian Santiago, Rachelle Thysell, Suzanne Rivoire, and Stephen W. Poole</i>	

Metrics for Evaluating Energy Saving Techniques for Resilient HPC Systems	790
<i>Ryan E. Grant, Stephen L. Olivier, James H. Laros III, Ron Brightwell, and Allan K. Porterfield</i>	

HPPAC Session 2: Power-Efficient Hardware

Reducing Static and Dynamic Power of L1 Data Caches in GPGPUs	798
<i>Ehsan Atoofian</i>	
Exploiting DMA for Performance and Energy Optimized STREAM on a DSP	805
<i>Gilbert Netzer, Lennart Johnsson, Daniel Ahlin, Eric Stotzer, Pekka Varis, and Erwin Laure</i>	
A Study of Energy and Locality Effects Using Space-Filling Curves	815
<i>Nico Reissman, Jan Christian Meyer, and Magnus Jahre</i>	

HPPAC Session 3: Large Scale Power Management

Energy-Aware Load Balancing Policies for the Cloud Ecosystem	823
<i>Ashkan Paya and Dan C. Marinescu</i>	
Bag-of-Task Scheduling on Power-Aware Clusters Using a DVFS-Based Mechanism	833
<i>George Terzopoulos and Helen D. Karatza</i>	
A Criticality-Aware DVFS Runtime Utility for Optimizing Power Efficiency of Multithreaded Applications	841
<i>Haibo Zhang, Wenting Han, Feng Li, Songtao He, Yichao Cheng, Hong An, and Zhitao Chen</i>	

Workshop 8: HPGC — High-Performance Grid and Cloud Computing Workshop

HPGC Introduction and Committees	849
<i>Eric Aubanel, Virendrakumar C. Bhavsar, and Michael Frumkin</i>	
HPGC Keynotes	850
<i>Rajkumar Buyya and Derek Murray</i>	

HPGC Session 1

Evaluating GPU Passthrough in Xen for High Performance Cloud Computing	852
<i>Andrew J. Younge, John Paul Walters, Stephen Crago, and Geoffrey C. Fox</i>	
Scalable System Environment Caching and Sharing for Distributed Virtual Machines	860
<i>Teng Long, Ilchul Yoon, Alan Sussman, Adam Porter, and Atif Memon</i>	
Mega Data Center for Elastic Internet Applications	868
<i>Hangwei Qian and Michael Rabinovich</i>	

HPGC Session 2

Cloud-Based Simulation of a Smart Power Grid	875
<i>Ashkan Paya and Dan C. Marinescu</i>	
Analyzing Reliability of Virtual Machine Instances with Dynamic Pricing in the Public Cloud	885
<i>Seung-Hwan Lim, Gautam S. Thakur, and James L. Horey</i>	

Security of Applications Involving Multiple Organizations and Order Preserving Encryption in Hybrid Cloud Environments	894
<i>Mohammad Ahmadian, Ashkan Paya, and Dan C. Marinescu</i>	
Workshop 9: AsHES — Accelerators and Hybrid Exascale Systems	
AsHES Introduction and Committees	904
<i>Yunquan Zhang</i>	
AsHES Keynote	907
<i>Jeffrey Vetter</i>	
AsHES Session 1: Programming Model and Performance Optimizations	
Scalable Critical Path Analysis for Hybrid MPI-CUDA Applications	908
<i>Felix Schmitt, Robert Dietrich, and Guido Juckeland</i>	
Dymaxion++: A Directive-Based API to Optimize Data Layout and Memory Mapping for Heterogeneous Systems	916
<i>Shuai Che, Jiayuan Meng, and Kevin Skadron</i>	
Comparison of Parallel Programming Models on Intel MIC Computer Cluster	925
<i>Chenggang Lai, Zhijun Hao, Miaoqing Huang, Xuan Shi, and Haihang You</i>	
CoAdELL: Adaptivity and Compression for Improving Sparse Matrix-Vector Multiplication on GPUs	933
<i>Marco Maggioni and Tanya Berger-Wolf</i>	
AsHES Session 2: Accelerating Applications	
Optimizing Krylov Subspace Solvers on Graphics Processing Units	941
<i>Hartwig Anzt, William Sawyer, Stanimire Tomov, Piotr Luszczek, Ichitaro Yamazaki, and Jack Dongarra</i>	
XSW: Accelerating Biological Database Search on Xeon Phi	950
<i>Lipeng Wang, Yuandong Chan, Xiaohui Duan, Haidong Lan, Xiangxu Meng, and Weiguo Liu</i>	
Dynamically Balanced Synchronization-Avoiding LU Factorization with Multicore and GPUs	958
<i>Simlice Donfack, Stanimire Tomov, and Jack Dongarra</i>	
Scalable Fast Multipole Accelerated Vortex Methods	966
<i>Qi Hu, Nail A. Gumerov, Rio Yokota, Lorena Barba, and Ramani Duraiswami</i>	
AsHES Session 3: Emerging Hybrid Systems	
Infiniband-Verbs on GPU: A Case Study of Controlling an Infiniband Network Device from the GPU	976
<i>Lena Oden, Holger Fröning, and Franz-Joseph Pfreundt</i>	
Programming the Adapteva Epiphany 64-Core Network-on-Chip Coprocessor	984
<i>Anish Varghese, Bob Edwards, Gaurav Mitra, and Alistair P. Rendell</i>	
High-Performance Zonal Histogramming on Large-Scale Geospatial Rasters Using GPUs and GPU-Accelerated Clusters	993
<i>Jianting Zhang and Dali Wang</i>	

Workshop 10: PLC — Programming Models, Languages, and Compilers Workshop for Manycore and Heterogeneous Architectures

PLC Introduction and Committees	1001
<i>Barbara Chapman</i>	

PLC Session 1: Programming and Compilation Techniques for GPUs

Transparent GPU Execution of NumPy Applications	1002
<i>Troels Blum, Mads R.B. Kristensen, and Brian Vinter</i>	
KernelGen — The Design and Implementation of a Next Generation Compiler Platform for Accelerating Numerical Models on GPUs	1011
<i>Dmitry Mikushin, Nikolay Likhogrud, Eddy Z. Zhang, and Christopher Bergström</i>	
Using GPU Shared Memory with a Directive-Based Approach	1021
<i>Wei Ding, Ligang Lu, Mauricio Araya-Polo, Amik St-Cyr, Detlef Hohl, and Barbara M. Chapman</i>	

PLC Session 2: Libraries and Optimization Frameworks

CFD Builder: A Library Builder for Computational Fluid Dynamics	1029
<i>Jagan Jayaraj, Pei-Hung Lin, Paul R. Woodward, and Pen-Chung Yew</i>	
A Stream Processing Framework for On-Line Optimization of Performance and Energy Efficiency on Heterogeneous Systems	1039
<i>Benjamin Ranft, Oliver Denninger, and Philip Pfaffe</i>	

PLC Session 3: Tools and Performance Evaluation

OpenMP Task Scheduling Analysis via OpenMP Runtime API and Tool Visualization	1049
<i>Ahmad Qawasmeh, Abid M. Malik, and Barbara M. Chapman</i>	
A Case Study in Coordination Programming: Performance Evaluation of S-Net vs Intel’s Concurrent Collections	1059
<i>Pavel Zaichenkov, Bert Gijsbers, Clemens Grellck, Olga Tveretina, and Alex Shafarenko</i>	

Workshop 11: EduPar-NSF/TCPP Workshop on Parallel and Distributed Computing Education

EduPar Introduction and Committees	1068
<i>Sushil K Prasad</i>	
EduPar Keynote	1070
<i>Randy H. Katz</i>	

EduPar Session: Introductory Course and Across Curriculum

Limited Time and Experience: Parallelism in CS1	1071
<i>Steven A. Bogaerts</i>	
NSF/IEEE-TCPP Curriculum Implementation at the State University of Nizhni Novgorod	1079
<i>Viktor Gergel, Alexey Liniov, Iosif Meyerov, and Alexander Sysoyev</i>	
Parallel and Distributed Computing across the Computer Science Curriculum	1085
<i>David J. John and Stan J. Thomas</i>	

EduPar Session: Software Engineering Courses

Service-Oriented Computing and Software Integration in Computing Curriculum	1091
<i>Yinong Chen and Zhizheng Zhou</i>	
EA: Research-Infused Teaching of Parallel Programming Concepts for Undergraduate Software Engineering Students	1099
<i>Nasser Giacaman and Oliver Sinnen</i>	
Using Patterns to Teach Parallel Computing	1106
<i>Clayton Ferner, Barry Wilkinson, and Barbara Heath</i>	

EduPar Session: Miscellaneous

Teaching HDFS/MapReduce Systems Concepts to Undergraduates	1114
<i>Linh Bao Ngo, Edward B. Duffy, and Amy W. Apon</i>	
Interactively Exploring the Connection between Nested Dissection Orderings for Parallel Cholesky Factorization and Vertex Separators	1122
<i>H. Martin Buecker and M. Ali Rostami</i>	
A Portable Cluster for Each Student	1130
<i>David Toth</i>	

Workshop 12: GABB — Graph Algorithms Building Blocks

GABB Introduction	1135
<i>Tim Mattson, David A. Bader, Aydın Buluç, John Gilbert, Joseph Gonzalez, and Jeremy Kepner</i>	

Workshop 13: PDSEC — Workshop on Parallel and Distributed Scientific and Engineering Computing

PDSEC Introduction and Committees	1138
<i>Peter Strazdins, Raphaël Couturier, Michelle Mills Strout, Keita Teranishi, Thomas Rauber, Gudula Rünger, and Laurence T. Yang</i>	

PDSEC Session 1: Best Papers

llamaOS: A Solution for Virtualized High-Performance Computing Clusters	1140
<i>William A. Magato and Philip A. Wilsey</i>	
New Algorithm for Computing Eigenvectors of the Symmetric Eigenvalue Problem	1150
<i>Azzam Haidar, Piotr Luszczek, and Jack Dongarra</i>	

PDSEC Session 2: Algorithms (I)

Exhaustive Key Search on Clusters of GPUs	1160
<i>Davide Barbieri, Valeria Cardellini, and Salvatore Filippone</i>	
Application Level Fault Recovery: Using Fault-Tolerant Open MPI in a PDE Solver	1169
<i>Md Mohsin Ali, James Southern, Peter Strazdins, and Brendan Harding</i>	
Nanoscale Cluster Detection in Massive Atom Probe Tomography Data	1179
<i>Sudip K. Seal, Srikanth B. Yeginath, and Michael K. Miller</i>	

Construction of Porous Networks Subjected to Geometric Restrictions by Using OpenMP	1189
<i>Angel González Méndez, Graciela Román Alonso, Fernando Rojas González,</i>	
<i>Miguel Alfonso Castro García, Miguel Aguilar Cornejo, and Salomón Cordero Sánchez</i>	

PDSEC Session 3: Systems and Performance Analysis

Integration and Evaluation of Decentralized Fairshare Prioritization (Aequus)	1198
<i>Daniel Espling, Per-Olov Östberg, and Erik Elmroth</i>	
Coordination Languages and MPI Perturbation Theory: The FOX Tuple Space Framework for Resilience	1208
<i>Jeremiah J. Wilke</i>	
DisSLib: CC: A Library for Distributed Search with a Central Common Search State	1218
<i>Tyson Kendon and Jörg Denzinger</i>	
Improving I/O Performance with Adaptive Data Compression for Big Data Applications	1228
<i>Hongbo Zou, Yongen Yu, Wei Tang, and Hsuanwei Michelle Chen</i>	
Analysis of MPI Shared-Memory Communication Performance from a Cache Coherence Perspective	1238
<i>Bertrand Putigny, Benoit Ruelle, and Brice Goglin</i>	

PDSEC Session 4: Algorithms (II)

Acceleration of GPU-Based Ultrasound Simulation via Data Compression	1248
<i>Andrew A. Haigh and Eric C. McCreath</i>	
Kd-Tree Based N-Body Simulations with Volume-Mass Heuristic on the GPU	1256
<i>Klaus Kofler, Dominik Steinhauser, Biagio Cosenza, Ivan Grasso, Sabine Schindler,</i> <i>and Thomas Fahringer</i>	
Nuclear Fusion Simulation Code Optimization and Performance Evaluation on GPU Cluster	1266
<i>Norihisa Fujita, Hideo Nuga, Taisuke Boku, and Yasuhiro Idomura</i>	
Acceleration of a Python-Based Tsunami Modelling Application via CUDA and OpenHMPP	1275
<i>Zhe Weng and Peter E. Strazdins</i>	
GPU Enhanced Path Finding for an Unmanned Aerial Vehicle	1285
<i>Roksana Hossain, Sebastian Magierowski, and Geoffery G. Messier</i>	

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

DPDNS Introduction and Committees	1294
<i>Dimitar Avresky, Erik Maehle, and Salvatore Distefano</i>	
DPDNS Keynote	1296
<i>Edgar Nett</i>	

DPDNS Session: Applications

Maintaining Dependable Communication Service for Mobile Stations in Wireless Mesh Networks by Tracking Capacity Demands	1297
<i>Timo Lindhorst, Burkhard Weseloh, and Edgar Nett</i>	
A Load Balancing Behavior for Underwater Robot Swarms to Increase Mission Time and Fault Tolerance	1306
<i>Ammar Amory, Thomas Tosik, and Erik Maehle</i>	
ExCoverly — A Framework for Distributed System Experiments and a Case Study of Service Discovery	1314
<i>Andreas Dittrich, Stefan Wanja, and Miroslaw Malek</i>	
Managing Soft-Errors in Transactional Systems	1324
<i>Mohamed Mohamedin, Roberto Palmieri, and Binoy Ravindran</i>	

DPDNS Session: Theoretical Aspects

Standby System Reliability through DRBD	1330
<i>Salvatore Distefano</i>	
Trust-Based Security for the Spanning Tree Protocol	1338
<i>Yingxu Lai, Qiuyue Pan, Zenghui Liu, Yinong Chen, and Zhizheng Zhou</i>	
Autonomy Requirements Engineering for Self-Adaptive Science Clouds	1344
<i>Emil Vassev and Mike Hinchey</i>	

Workshop 15: MTAAP — Workshop on Multi-threaded Architectures and Applications

MTAAP Introduction and Committees	1354
<i>Luiz DeRose</i>	

MTAAP Session: Algorithms and Position Papers

A New Parallel Algorithm for Two-Pass Connected Component Labeling	1355
<i>Siddharth Gupta, Diana Palsetia, Md. Mostofa Ali Patwary, Ankit Agrawal, and Alok Choudhary</i>	
Position Paper: Locality-Driven Scheduling of Tasks for Data-Dependent Multithreading	1363
<i>Jaime Arteaga, Stephane Zuckerman, Elkin Garcia, and Guang Gao</i>	
Position Paper: Leveraging Strength-Based Dynamic Slicing to Identify Control Reconvergence Instructions	1368
<i>Walid J. Ghandour and Nadine J. Ghandour</i>	

MTAAP Session: Graph Analytics

Parallel Heuristics for Scalable Community Detection	1374
<i>Hao Lu, Mahantesh Halappanavar, Ananth Kalyanaraman, and Sutanay Choudhury</i>	
Hardware/Software Vectorization for Closeness Centrality on Multi-/Many-Core Architectures	1386
<i>Ahmet Erdem Sariyuce, Erik Saule, Kamer Kaya, and Ümit V. Çatalyürek</i>	

Revisiting Edge and Node Parallelism for Dynamic GPU Graph Analytics	1396
<i>Adam McLaughlin and David A. Bader</i>	

MTAAP Session: Accelerators

A Validation Testsuite for OpenACC 1.0	1407
<i>Cheng Wang, Rengan Xu, Sunita Chandrasekaran, Barbara Chapman, and Oscar Hernandez</i>	
Extracting Maximal Exact Matches on GPU	1417
<i>Anas Abu-Doleh, Kamer Kaya, Mohamed Abouelhoda, and Ümit V. Çatalyürek</i>	
Predicting an Optimal Sparse Matrix Format for SpMV Computation on GPU	1427
<i>B. Neelima, G. Ram Mohana Reddy, and Prakash S. Raghavendra</i>	

Workshop 16: LSPP — Workshop on Large-Scale Parallel Processing

LSPP Introduction and Committees	1437
<i>Darren J. Kerbyson, Ram Rajamony, and Charles Weems</i>	

LSPP Session 1: Performance Analysis and Optimization

Higher Dimensional Gaussian Networks	1438
<i>Arash Shamaei, Bella Bose, and Mary Flahive</i>	

LSPP Session 2: Modeling Performance for Scaling

The Power-Performance Tradeoffs of the Intel Xeon Phi on HPC Applications	1448
<i>Bo Li, Hung-Ching Chang, Shuaiwen Song, Chun-Yi Su, Timmy Meyer, John Mooring, and Kirk W. Cameron</i>	
Performance Modeling for Hardware Thread-Level Speculation	1457
<i>Ying-Chieh Wang, Che-Rung Lee, Yeh-Ching Chung, I-Hsin Chung, and Michael Perrone</i>	
HMC-Sim: A Simulation Framework for Hybrid Memory Cube Devices	1465
<i>John D. Leidel and Yong Chen</i>	

LSPP Session 3: Large-Scale Systems

Online Monitoring System for Performance Fault Detection	1475
<i>Roberto Gioiosa, Gokcen Kestor, and Darren J. Kerbyson</i>	

LSPP Session 4: Scheduling

Towards Extreme-Scale Simulations with Next-Generation Trilinos: A Low Mach Fluid Application Case Study	1485
<i>Paul Lin, Matthew Bettencourt, Stefan Domino, Travis Fisher, Mark Hoemmen, Jonathan Hu, Eric Phipps, Andrey Prokopenko, Sivasankaran Rajamanickam, Christopher Siefert, Eric Cyr, and Stephen Kennon</i>	
Design and Implementation of a Large Scale Tree-Based QR Decomposition Using a 3D Virtual Systolic Array and a Lightweight Runtime	1495
<i>Ichitaro Yamazaki, Jakub Kurzak, Piotr Luszczek, and Jack Dongarra</i>	

SupMR: Circumventing Disk and Memory Bandwidth Bottlenecks for Scale-up MapReduce	1505
<i>Michael Sevilla, Ike Nassi, Kleoni Ioannidou, Scott Brandt, and Carlos Maltzahn</i>	

Workshop 17: PCO — Parallel Computing and Optimization

PCO Introduction and Committees	1515
<i>Didier El-Baz</i>	

PCO Session 1: Optimization Techniques for Parallel or Distributed Architectures

Towards Energy Efficient Allocation for Applications in Volunteer Cloud	1516
<i>Congfeng Jiang, Jian Wan, Christophe Cérin, Paolo Gianessi, and Yanik Ngoko</i>	
Fast Generation of Large Task Network Mappings	1526
<i>Karl-Eduard Berger, François Galea, Bertrand Le Cun, and Renaud Sirdey</i>	

PCO Session 2: Parallel Optimization Algorithms

Adaptive N to P Portfolio for Solving Constraint Programming Problems on Top of the Parallel Bobpp Framework	1531
<i>Tarek Menouer and Bertrand Le Cun</i>	
Dependent Walks in Parallel Local Search	1541
<i>Yves Caniou and Philippe Codognet</i>	
A Parallel Large Neighborhood Search-Based Heuristic for the Disjunctively Constrained Knapsack Problem	1547
<i>Mhand Hifi, Stephane Negre, Toufik Saadi, Sagvan Saleh, and Lei Wu</i>	
Solving Hard MIPLIB2003 Problems with ParaSCIP on Supercomputers: An Update	1552
<i>Yuji Shinano, Tobias Achterberg, Timo Berthold, Stefan Heinz, Thorsten Koch, and Michael Winkler</i>	

PCO Session 3: Task Scheduling and Miscellaneous

A Task Scheduling Algorithm Based on Replication for Maximizing Reliability on Heterogeneous Computing Systems	1562
<i>Shuli Wang, Kenli Li, Jing Mei, Keqin Li, and Yan Wang</i>	
SkewControl: Gini Out of the Bottle	1572
<i>Si Zheng, Yunhuai Liu, Tian He, Li Shanshan, and Xiangke Liao</i>	
The Heuristic Static Load-Balancing Algorithm Applied to the Community Earth System Model	1581
<i>Yuri Alexeev, Sheri Mickelson, Sven Leyffer, Robert Jacob, and Anthony Craig</i>	
A Distributed Algorithm for a Reconfigurable Modular Surface	1591
<i>Didier El-Baz, Benoît Piranda, and Julien Bourgeois</i>	

Workshop 18: ParLearning — Workshop on Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analytics

ParLearning Introduction and Committees	1599
<i>Abhinav Vishnu and Yinglong Xia</i>	

ParLearning Keynote	1601
<i>Eric P. Xing</i>	

ParLearning Session 1

Wait-Free Primitives for Initializing Bayesian Network Structure Learning on Multicore Processors	1602
<i>Hsuan-Yi Chu, Yinglong Xia, Anand Panangadan, and Viktor K. Prasanna</i>	

gpuRF and gpuERT: Efficient and Scalable GPU Algorithms for Decision Tree Ensembles	1612
<i>Karl Jansson, Håkan Sundell, and Henrik Boström</i>	

Training Large Scale Deep Neural Networks on the Intel Xeon Phi Many-Core Coprocessor	1622
<i>Lei Jin, Zhaokang Wang, Rong Gu, Chunfeng Yuan, and Yihua Huang</i>	

Parallel Bayesian Network Modelling for Pervasive Health Monitoring System	1631
<i>Xiujuan Qian, Yongli Wang, and Xiaohui Jiang</i>	

ParLearning Session 2

Portfolio-Based Selection of Robust Dynamic Loop Scheduling Algorithms Using Machine Learning	1638
<i>Nitin Sukhija, Brandon Malone, Srishti Srivastava, Ioana Banicescu, and Florina M. Ciorba</i>	

A General P2P Scheme for Constructing Large-Scale Virtual Environments	1648
<i>Wei Wang, Guisong Yang, Naixue Xiong, Xingyu He, and Wenzhong Guo</i>	

ParLearning Session 3

Large Scale Discriminative Metric Learning	1656
<i>Peter D. Kirchner, Matthias Boehm, Berthold Reinwald, Daby Sow, Michael Schmidt, Deepak Turaga, and Alain Biem</i>	

YAFIM: A Parallel Frequent Itemset Mining Algorithm with Spark	1664
<i>Hongjian Qiu, Rong Gu, Chunfeng Yuan, and Yihua Huang</i>	

The Empirical Research of Virtual Enterprise Knowledge Transfer's Effectiveness Faced to the Independent Innovation Ability	1672
<i>Yang Bo, Naixue Xiong, and Wenzhong Guo</i>	

A Distributed Speech Algorithm for Large Scale Data Communication Systems	1680
<i>Naixue Xiong, Guoxiang Tong, Wenzhong Guo, Jian Tan, and Guanning Wu</i>	

Workshop 19: HPDIC - High Performance Data Intensive Computing

HPDIC Introduction and Committees	1688
<i>Christophe Cerin and Cong-Feng Jiang</i>	

HPDIC Session 1: Memory, I/O, and Performance Enhancement

Compactor: Optimization Framework at Staging I/O Nodes	1689
<i>Vishwanath Venkatesan, Mohamad Chaarawi, Quincey Koziol, and Edgar Gabriel</i>	
Hybrid BFS Approach Using Semi-external Memory	1698
<i>Keita Iwabuchi, Hitoshi Sato, Ryo Mizote, Yuichiro Yasui, Katsuki Fujisawa, and Satoshi Matsuoka</i>	
Model-Driven Data Layout Selection for Improving Read Performance	1708
<i>Jialin Liu, Surendra Byna, Bin Dong, Kesheng Wu, and Yong Chen</i>	

HPDIC Session 2: Clustering, Data Management, and Applications

Scalable and Reliable Data Broadcast with Cascade	1717
<i>Stéphane Martin, Tomasz Buchert, Pierric Willemet, Olivier Richard, Emmanuel Jeanvoine, and Lucas Nussbaum</i>	
SOM Clustering Using Spark-MapReduce	1727
<i>Tugdual Sarazin, Hanane Azzag, and Mustapha Lebbah</i>	
Optimizing the Join Operation on Hive to Accelerate Cross-Matching in Astronomy	1735
<i>Liang Li, Dixin Tang, Taoying Liu, Hong Liu, Wei Li, and Chenzhou Cui</i>	

Workshop 20: JSSPP — Workshop on Job Scheduling Strategies for Parallel Processing

JSSPP Introduction and Committees	1746
<i>Walfredo Cirne and Narayan Desai</i>	

Workshop 21: CHIUW — Chapel Implementers and Users Workshop

CHIUW Introduction and Committees	1747
<i>Brad Chamberlain</i>	

Author Index