

2013 IEEE International Symposium on Parallel & Distributed Processing, Workshops and Phd Forum

(IPDPSW 2013)

**Cambridge, Massachusetts, USA
20 – 24 May 2013**

Pages 1-770



**IEEE Catalog Number: CFP1351J-POD
ISBN: 978-1-4799-1372-5**

2013 IEEE 27th International Symposium on Parallel & Distributed Processing Workshops and PhD Forum

IPDPSW 2013

Table of Contents

Message from the IPDPS General Co-Chairs to the Workshops.....	xxix
Message from the Workshops Chair.....	xxx

Heterogeneity in Computing Workshop—HCW

HCW Introduction	1
<i>Alexey Lastovetsky and Uwe Schwiegelshohn</i>	
Message from the HCW Steering Committee Chair	3
<i>Behrooz Shirazi</i>	
Message from the HCW General Chair	4
<i>Alexey Lastovetsky</i>	
Message from the HCW Program Chair	5
<i>Uwe Schwiegelshohn</i>	
HCW 2013 Keynote Talk	6
<i>Jack Dongarra</i>	

HCW Session 1: Scheduling and Resource Allocation

Network Delay-Aware Load Balancing in Selfish and Cooperative Distributed Systems	7
<i>Piotr Skowron and Krzysztof Rzadca</i>	
An Analysis Framework for Investigating the Trade-Offs between System Performance and Energy Consumption in a Heterogeneous Computing Environment	19
<i>Ryan Friese, Bhavesh Khemka, Anthony A. Maciejewski, Howard Jay Siegel, Gregory A. Koenig, Sarah Powers, Marcia Hilton, Jendra Rambharos, Gene Okonski, and Stephen W. Poole</i>	
Scheduling Tightly-Coupled Applications on Heterogeneous Desktop Grids	31
<i>Henri Casanova, Fanny Dufossé, Yves Robert, and Frédéric Vivien</i>	
SDBATS: A Novel Algorithm for Task Scheduling in Heterogeneous Computing Systems	43
<i>Ehsan Ullah Munir, Sajjad Mohsin, Altaf Hussain, Muhammad Wasif Nisar, and Shoukat Ali</i>	

HCW Session 2: Heterogeneous Processors

An On-chip Heterogeneous Implementation of a General Sparse Linear Solver	54
<i>Arash Sadrieh, Stefano Charissis, and Adam P. Hill</i>	
Parallel Macro Pipelining on the Intel SCC Many-Core Computer	64
<i>Tim Süß, Andrew Schoenrock, Sebastian Meisner, and Christian Plessl</i>	
Brawny vs. Wimpy: Evaluation and Analysis of Modern Workloads on Heterogeneous Processors	74
<i>Vishal Gupta and Karsten Schwan</i>	

HCW Session 3: Communication in Heterogeneous Systems

Seeds for a Heterogeneous Interconnect	84
<i>Adam Hackett, Deepak Ajwani, Shoukat Ali, Steve Kirkland, and John P. Morrison</i>	
Issues in Communication Heterogeneity for Message-Passing Concurrent Computing	93
<i>Jaroslav Slawinski, Umberto Villa, Tiziano Passerini, Alessandro Veneziani, and Vaidy Sunderam</i>	

Reconfigurable Architectures Workshop—RAW

RAW Introduction	103
<i>Jürgen Becker, Ramachandran Vaidyanathan, Peter Athanas, Marco D. Santambrogio, René Cumplido, and Oliver Sander</i>	

RAW Session 1: Systems and Applications

RTL Simulation of High Performance Dynamic Reconfiguration: A Video Processing Case Study	106
<i>Lingkan Gong, Oliver Diessel, Johnny Paul, and Walter Stechele</i>	
High-Performance Pipelined Architecture for Tree-Based IP Lookup Engine on FPGA	114
<i>Yun Qu and Viktor K. Prasanna</i>	
A Comparison of Ruleset Feature Independent Packet Classification Engines on FPGA	124
<i>Andrea Sanny, Thilan Ganegedara, and Viktor K. Prasanna</i>	

RAW Session 2: Architectures and Algorithms I

Architecture Exploration of High-Performance Floating-Point Fused Multiply-Add Units and their Automatic Use in High-Level Synthesis	134
<i>Björn Liebig, Jens Huthmann, and Andreas Koch</i>	
A Flexible Memory Controller Supporting Deep Belief Networks with Fixed-Point Arithmetic	144
<i>Jingfei Jiang, Rongdong Hu, and Mikel Luján</i>	
Hardware Supported Adaptive Data Collection for Networks on Chip	153
<i>Jan Heißwolf, Andreas Weichslgartner, Aurang Zaib, Ralf König, Thomas Wild, Andreas Herkersdorf, Jürgen Teich, and Jürgen Becker</i>	

RAW Session 3: Software and Tools

An FPGA Router for Alternative Reconfiguration Flows	163
<i>Wenwei Zha and Peter Athanas</i>	
Automated Partitioning for Partial Reconfiguration Design of Adaptive Systems	172
<i>Kizheppatt Vipin and Suhaib A. Fahmy</i>	
A Novel FPGA-based Evolvable Hardware System Based on Multiple Processing Arrays	182
<i>Ángel Gallego, Javier Mora, Andrés Otero, Rubén Salvador, Eduardo de la Torre, and Teresa Riesgo</i>	

RAW Session 4: Architectures and Algorithms II

A Flexible Interconnection Structure for Reconfigurable FPGA Dataflow Applications	192
<i>Gianluca Durelli, Alessandro A. Nacci, Riccardo Cattaneo, Christian Pilato, Donatella Sciuto, and Marco D. Santambrogio</i>	
Design and Implementation of High Performance Architectures with Partially Reconfigurable CGRAs	202
<i>Mansureh Shahraki Moghaddam, Kolin Paul, and M. Balakrishnan</i>	
Portable Memory Consistency for Software Managed Distributed Memory in Many-Core SoC	212
<i>Jochem H. Rutgers, Marco J.G. Bekooij, and Gerard J.M. Smit</i>	

RAW Session 5: Software and Tools II

Cross-Architectural Study of Custom Reconfigurable Devices Using Crowdsourcing	222
<i>AnilKumar Sistla, Natalie Parde, Krunalkumar Patel, and Gayatri Mehta</i>	
HERA Project's Holistic Evolutionary Framework	231
<i>Davide B. Bartolini, Matteo Carminati, Fabio Cancare, Marco D. Santambrogio, and Donatella Sciuto</i>	
A Hybrid FPGA Model to Estimate On-Chip Crossbar Logic Utilization in SoC Platforms	239
<i>Yoon Kah Leow and Ali Akoglu</i>	

RAW Poster Session 1

Reconfigurable Architecture of a RRC Fir Interpolator for Multi-standard Digital Up Converter	247
<i>Indranil Hatai, Indrajit Chakrabarti, and Swapna Banerjee</i>	
Virtual UARTs for Reconfigurable Multi-processor Architectures	252
<i>Pierre Bomel, Kevin Martin, and Jean-Philippe Diguët</i>	
Heterogeneous CPU/FPGA Reconfigurable Computing System for Avionic Test Application	260
<i>George Afonso, Zeineb Baklouti, David Duvivier, Rabie Ben Atitallah, Eli Billauer, and Stephan Stilkerich</i>	
FPGA Implementation of Subcarrier Index Modulation OFDM Transceiver	268
<i>Micahel Mefenza and Christophe Bobda</i>	
Hardware MPI-2 Functions for Multi-Processing Reconfigurable System on Chip	273
<i>Roland Christian Gamom Ngounou Ewo, Emmanuel Kiegaing, Martin Mbouenda, Hilaire Bertrand Fotsin, and Bertrand Granado</i>	

RAW Poster Session 2

Thermal Aware Module Placement for Heterogeneous 3D-IC Based FPGAs	281
<i>Alexander Wold, Dirk Koch, and Jim Torresen</i>	
A Hierarchical Architectural Framework for Reconfigurable Logic Computing	287
<i>Peng Li, Angshuman Parashar, Michael Pellauer, Tao Wang, and Joel Emer</i>	
A Case Study into Predictable and Composable MPSoC Reconfiguration	293
<i>Pranav Tendulkar and Sander Stuijk</i>	
Real-Time Supply Voltage Sensor for Detecting/Debugging Electrical Timing Failures	301
<i>Miho Ueno, Masanori Hashimoto, and Takao Onoye</i>	
A Hardware Approach for Solving the Robot Localization Problem Using a Sequential EKF	306
<i>Sérgio Cruz, Daniel M. Muñoz, Milton E. Conde, Carlos H. Llanos, and Geovany A. Borges</i>	

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

HIPS Introduction	314
<i>Kathryn Mohror and Stephen L. Olivier</i>	

HIPS Session 1: Support for Accelerator and Manycore Architectures

GPU Programming in Rust: Implementing High-Level Abstractions in a Systems-Level Language	315
<i>Eric Holk, Milinda Pathirage, Arun Chauhan, Andrew Lumsdaine, and Nicholas D. Matsakis</i>	
TSHMEM: Shared-Memory Parallel Computing on Tilera Many-Core Processors	325
<i>Bryant C. Lam, Alan D. George, and Herman Lam</i>	
Introducing a Data Sliding Mechanism for Cooperative Caching in Manycore Architectures	335
<i>Safae Dahmani, Loïc Cudennec, and Guy Gogniat</i>	

HIPS Session 2: Scalable Tools for Correctness and Performance Analysis

Interactive Debugging of Dynamic Dataflow Embedded Applications	345
<i>Kevin Pouget, Patricia López Cueva, Miguel Santana, and Jean-François Méhaut</i>	
How to Scale Dynamic Tuning to Large Parallel Applications	355
<i>Andrea Martínez, Anna Sikora, Eduardo César, and Joan Sorribes</i>	
Model Checking Stencil Computations Written in a Partitioned Global Address Space Language	365
<i>Tatsuya Abe, Toshiyuki Maeda, and Mitsuhsa Sato</i>	

HIPS Session 3: Programming Models and Abstractions

Loop Chaining: A Programming Abstraction for Balancing Locality and Parallelism	375
<i>Christopher D. Krieger, Michelle Mills Strout, Catherine Olschanowsky, Andrew Stone, Stephen Guzik, Xinfeng Gao, Carlo Bertolli, Paul H.J. Kelly, Gihan Mudalige, Brian Van Straalen, and Sam Williams</i>	

Toward Abstracting the Communication Intent in Applications to Improve Portability and Productivity	385
<i>Tiffany M. Mintz, Oscar Hernandez, Christos Kartsaklis, David E. Bernholdt, Markus Eisenbach, and Swaroop Pophale</i>	

Programming Support for Speculative Execution with Software Transactional Memory	394
<i>Min Feng, Rajiv Gupta, and Iulian Neamtiu</i>	

Workshop on Nature Inspired Distributed Computing—NIDISC

NIDISC Introduction	404
<i>Pascal Bouvry, Franciszek Seredynski, and El-Ghazali Talbi</i>	

NIDISC Session 1: Multi-objective Optimization

A Possibilistic Framework for Solving Multi-objective Problems under Uncertainty: Definition of New Pareto Optimality	405
<i>Bahri Oumayma, Ben Amor Nahla, and El-Ghazali Talbi</i>	

A Parallel Multi-objective Local Search for AEDB Protocol Tuning	415
<i>Santiago Iturriaga, Patricia Ruiz, Sergio Nesmachnow, Bernabé Dorronsoro, and Pascal Bouvry</i>	

Scalability Analysis of the Asynchronous, Master-Slave Borg Multiobjective Evolutionary Algorithm	425
<i>David Hadka, Kamesh Madduri, and Patrick Reed</i>	

ShadObf: A C-Source Obfuscator Based on Multi-objective Optimisation Algorithms	435
<i>Benoît Bertholon, Sébastien Varrette, and Sébastien Martinez</i>	

NIDISC Session 2: Application of Bio-Inspired Algorithms

Application of Evolutionary Algorithms to Maximum Lifetime Coverage Problem in Wireless Sensor Networks	445
<i>Antonina Tretyakova and Franciszek Seredynski</i>	

Tree-Based Consensus Model for Proline Cis-Trans Isomerization Prediction	454
<i>Paul D. Yoo, Albert Y. Zomaya, Khalfan Alromaihi, and Sara Alshamsi</i>	

Multi-segment Green Light Optimal Speed Advisory	459
<i>Marcin Seredynski, Wojciech Mazurczyk, and Djamel Khadraoui</i>	

Novel Levy Based Particle Swarm Optimization Algorithm for Electrical Power Grid	466
<i>Boussaad Ismail, Amir Nakib, Frederic Heliodore, Serge Poullain, and Patrick Siarry</i>	

NIDISC Session 3: Scheduling and Resource Allocation

Optimal Peer Selection Strategy in P2P-VoD Systems Using Dynamic Evolution Strategy	474
<i>Thibaud Rohmer, Amir Nakib, and Abdelhamid Nafaa</i>	

Optimization of Shared-Memory Multicore Systems Using Game Theory and Genetic Algorithms on Cellular Automata Lattices	482
<i>Michail-Antisthenis I. Tsompanas, Christoforos Kachris, and Georgios Ch. Sirakoulis</i>	

Cost Minimization of Service Deployment in a Public Cloud Environment	491
<i>François Legillon, Nouredine Melab, Didier Renard, and El-Ghazali Talbi</i>	

Workshop on High Performance Computational Biology—HiCOMB

HiCOMB Introduction	499
<i>Jaroslav Zola, David A. Bader, and Srinivas Aluru</i>	

HiCOMB Session I: Sequence and Structure

Biological Sequence Comparison on Hybrid Platforms with Dynamic Workload Adjustment	501
<i>Fernando Machado Mendonca and Alba Cristina Magalhaes Alves de Melo</i>	
Accelerating All-to-All Protein Structures Comparison with TAlign Using a NoC Many-Cores Processor Architecture	510
<i>Anuj Sharma, Antonis Papanikolaou, and Elias S. Manolakos</i>	
Secondary Structure Predictions for Long RNA Sequences Based on Inversion Excursions and MapReduce	520
<i>Daniel T. Yehdego, Boyu Zhang, Vikram K.R. Kodimala, Kyle L. Johnson, Michela Taufer, and Ming-Ying Leung</i>	

HiCOMB Session II: Phylogenetics and Metagenomics

A Generic Vectorization Scheme and a GPU Kernel for the Phylogenetic Likelihood Library	530
<i>Fernando Izquierdo-Carrasco, Nikolaos Alachiotis, Simon Berger, Tomas Flouri, Solon P. Pissis, and Alexandros Stamatakis</i>	
Boosting the Performance of Bayesian Divergence Time Estimation with the Phylogenetic Likelihood Library	539
<i>Diego Darriba, Andre Aberer, Tomas Flouri, Tracy A. Heath, Fernando Izquierdo-Carrasco, and Alexandros Stamatakis</i>	
A Map-Reduce Framework for Clustering Metagenomes	549
<i>Zeehasham Rasheed and Huzefa Rangwala</i>	

HiCOMB Session III: Molecular and Systems Biology

GPU-Accelerated Protein Family Identification for Metagenomics	559
<i>Changjun Wu and Ananth Kalyanaraman</i>	
Polarization Energy on a Cluster of Multicores	569
<i>Jesmin Jahan Tithi and Rezaul A. Chowdhury</i>	
GPU-Based Steady-State Solution of the Chemical Master Equation	579
<i>Marco Maggioni, Tanya Berger-Wolf, and Jie Liang</i>	

Advances in Parallel and Distributed Computing Models—APDCM

APDCM Introduction	589
<i>Oscar H. Ibarra</i>	

APDCM Session 1: Parallel Computing 1

The Hierarchical Memory Machine Model for GPUs	591
<i>Koji Nakano</i>	
Toward a Generic Hybrid CPU-GPU Parallelization of Divide-and-Conquer Algorithms	601
<i>Alejandro López-Ortiz, Alejandro Salinger, and Robert Suderman</i>	
Time-Power Tradeoffs for Sorting on a Mesh-Connected Computer with Optical Connections	611
<i>Patrick Poon and Quentin F. Stout</i>	

APDCM Session 2: Parallel Computing 2

Asynchronous P Systems for the Maximum Independent Set and Related Graph Problems	620
<i>Kohei Tanaka and Akihiro Fujiwara</i>	
EasyHPS: A Multilevel Hybrid Parallel System for Dynamic Programming	630
<i>Jun Du, Ce Yu, Jizhou Sun, Chao Sun, Shanjiang Tang, and Yanlong Yin</i>	
Hardware and Software Support for NUMA Computing on Configurable Emulated Shared Memory Architectures	640
<i>Martti Forsell, Erik Hansson, Christoph Kessler, Jari-Matti Mäkelä, and Ville Leppänen</i>	
Agent-Based Traffic Merging in Network-on-Chip	649
<i>Mengjia Yan, Weiwei Fu, Chao Wang, Tianzhou Chen, and Li Liu</i>	

APDCM Session 3: Distributed Computing 1

Self-Stabilizing Master-Slave Token Circulation Algorithm in an Undirected Ring of Arbitrary Size and Its Orientation	659
<i>Yihua Ding, James Wang, and Pradip K. Srimani</i>	
A Unified Approach for Different Tasks on Rings in Robot-Based Computing Systems	667
<i>Gianlorenzo D'Angelo, Gabriele Di Stefano, Alfredo Navarra, Nicolas Nisse, and Karol Suchan</i>	
Parallel File Download in Peer-to-Peer Networks with Random Service Capacities	677
<i>Keqin Li</i>	
Distributed Query Processing in an Ad-hoc Semantic Web Data Sharing System	687
<i>Jing Zhou, Gregor V. Bochmann, and Zhongzhi Shi</i>	

APDCM Session 4: Distributed Computing 2

Network Decontamination from a Black Virus	696
<i>Jie Cai, Paola Flocchini, and Nicola Santoro</i>	
Revisiting the Double Checkpointing Algorithm	706
<i>Jack Dongarra, Thomas Herault, and Yves Robert</i>	

APDCM Session 5: Short Presentations

Fast Leader (Full) Recovery Despite Dynamic Faults	716
<i>Ajoy K. Datta, Lawrence L. Larmore, Stéphane Devismes, and Sébastien Tixeuil</i>	

K-Selection Protocols from Energetic Complexity Perspective	726
<i>Marcin Kardas, Marek Klonowski, Dominik Pajak, and Kamil Wolny</i>	
Cooperative MIMO Paradigms for Cognitive Radio Networks	734
<i>Wei Chen and Liang Hong</i>	
Design Optimization of Energy- and Security-Critical Distributed Real-Time Embedded Systems	741
<i>Xia Zhang, Jinyu Zhan, Wei Jiang, Yue Ma, and Ke Jiang</i>	
On Analyzing Large Graphs Using GPUs	751
<i>Amlan Chatterjee, Sridhar Radhakrishnan, and John K. Antonio</i>	
Operation Synchronization Technique on Pipeline-Based Hardware Synthesis Applying Stream-Based Computing Framework	761
<i>Shinichi Yamagiwa, Ryoyu Watanabe, and Koichi Wada</i>	
Efficient Hough Transform on the FPGA using DSP Slices and Block RAMs	771
<i>Xin Zhou, Norihiro Tomagou, Yasuaki Ito, and Koji Nakano</i>	

Communication Architecture for Scalable Systems—CASS

CASS Introduction	779
<i>José Flich, Scott Pakin, and Craig Stunkel</i>	

CASS Session I

Network-on-Chip with Long-Range Wireless Links for High-Throughput Scientific Computation	781
<i>Turbo Majumder, Partha Pratim Pande, and Ananth Kalyanaraman</i>	
Fault Localization in NoCs Exploiting Periodic Heartbeat Messages in a Many-Core Environment	791
<i>Arne Garbade, Sebastian Weis, Sebastian Schlingmann, Bernhard Fechner, and Theo Ungerer</i>	
Head-of-Line Blocking Avoidance in Networks-on-Chip	796
<i>José V. Escamilla, José Flich, and Pedro J. Garcia</i>	

CASS Session II

GPU Peer-to-Peer Techniques Applied to a Cluster Interconnect	806
<i>Roberto Ammendola, Massimo Bernaschi, Andrea Biagioni, Mauro Bisson, Massimiliano Fatica, Ottorino Frezza, Francesca Lo Cicero, Alessandro Lonardo, Enrico Mastrostefano, Pier Stanislao Paolucci, Davide Rossetti, Francesco Simula, Laura Tosoratto, and Piero Vicini</i>	
Direct MPI Library for Intel Xeon Phi Co-Processors	816
<i>Min Si, Yutaka Ishikawa, and Masamichi Tatagi</i>	
Building Scalable PGAS Communication Subsystem on Blue Gene/Q	825
<i>Abhinav Vishnu, Darren J. Kerbyson, Kevin Barker, and Hubertus van Dam</i>	

CASS Session III

RRR: A Load Balanced Routing Scheme for Slimmed Fat-Trees	834
<i>Xin Yuan, Santosh Mahapatra, Michael Lang, and Scott Pakin</i>	
Predict-More Router: A Low Latency NoC Router with More Route Predictions	842
<i>Yuan He, Hiroshi Sasaki, Shinobu Miwa, and Hiroshi Nakamura</i>	

High-Performance, Power-Aware Computing—HPPAC

HPPAC Introduction	851
<i>Bronis R. de Supinski and Dong Li</i>	

HPPAC Session 1: Power Efficient Hardware

Measuring Power Consumption on IBM Blue Gene/Q	853
<i>Sean Wallace, Venkatram Vishwanath, Susan Coghlan, Zhiling Lan, and Michael E. Papka</i>	
PowerTune: Differentiated Power Allocation in Over-Provisioned Multicore Systems	860
<i>Vishal Gupta and Karsten Schwan</i>	
Decreasing Network Power with on-off Links Informed by Scientific Applications	868
<i>Gilbert Hendry</i>	

HPPAC Session 2 : Energy/Power Measurement and Profiling

Design of a Concentrated Torus Topology with Channel Buffers and Efficient Crossbars in NoCs	876
<i>Dominic DiTomaso, Randy Morris, Evan Jolley, Ashwini Sarathy, Ahmed Louri, and Avinash Kodi</i>	
Power Measurement and Concurrency Throttling for Energy Reduction in OpenMP Programs	884
<i>Allan K. Porterfield, Stephen L. Olivier, Sridutt Bhalachandra, and Jan F. Prins</i>	
General Recommendations for High Performance Computing Data Center Energy Management Dashboard Display	892
<i>Dale Sartor, Rod Mahdavi, Ben D. Radhakrishnan, Natalie Bates, Anna Maria Bailey, and Ralph Wescott</i>	

HPPAC Session 3: Large Scale Power Management

Energy Consumption Models and Predictions for Large-Scale Systems	899
<i>Taghrid Samak, Christine Morin, and David Bailey</i>	
Analysis of a Self-Organizing Algorithm for Energy Saving in Data Centers	907
<i>Carlo Mastroianni, Michela Meo, and Giuseppe Papuzzo</i>	
Toward Runtime Power Management of Exascale Networks by on/off Control of Links	915
<i>Ehsan Tottoni, Nikhil Jain, and Laxmikant V. Kalé</i>	

HPPAC Session 4: Compiler and Runtime Techniques

A Compiler Analysis to Determine Useful Cache Size for Energy Efficiency	923
<i>Sanket Tavarageri and P. Sadayappan</i>	
Energy-Efficient Sparse Matrix Autotuning with CSX—A Trade-off Study	931
<i>Jan Christian Meyer, Juan Manuel Cebrian, Lasse Natvig, Vasileios Karakasis, Dimitris Siakavaras, and Konstantinos Nikas</i>	
Evaluation of Energy Characteristics of MPI Communication Primitives with RAPL	938
<i>Akshay Venkatesh, Krishna Kandalla, and Dhabaleswar K. Panda</i>	

High-Performance Grid and Cloud Computing Workshop—HPGC

HPGC Introduction	946
<i>Eric Aubanel and Michael Frumkin</i>	

HPGC Session 1: High Performance Cloud Computing

Dynamic Sharing of GPUs in Cloud Systems	947
<i>Khaled M. Diab, M. Mustafa Rafique, and Mohamed Hefeeda</i>	
Distributed Loop Scheduling Schemes for Cloud Systems	955
<i>Yiming Han and Anthony T. Chronopoulos</i>	
Distributing Storage in Cloud Environments	963
<i>Petra Berenbrink, Andre Brinkmann, Tom Friedetzky, Dirk Meister, and Lars Nagel</i>	
BOINC as a Service for the SlapOS Cloud: Tools and Methods	974
<i>Christophe Cérin and Alain Takoudjou</i>	

HPGC Session 2: Virtual Machines and Data Clouds

vGreenNet: Managing Server and Networking Resources of Co-Located Heterogeneous VMs	984
<i>Liyi Zhang, Gaurav Dhiman, and Tajana Simunic Rosing</i>	
Ninja Migration: An Interconnect-Transparent Migration for Heterogeneous Data Centers	992
<i>Ryousei Takano, Hidemoto Nakada, Takahiro Hirofuchi, Yoshio Tanaka, and Tomohiro Kudoh</i>	
Performance Implications from Sizing a VM on Multi-core Systems: A Data Analytic Application’s View	1001
<i>Seung-Hwan Lim, James Horey, Yanjun Yao, Edmon Begoli, and Qing Cao</i>	
MapReduce Algorithms for GIS Polygonal Overlay Processing	1009
<i>Satish Puri, Dinesh Agarwal, Xi He, and Sushil K. Prasad</i>	

Accelerators and Hybrid Exascale Systems—ASHES

ASHES Introduction	1017
<i>Jiayuan Meng</i>	

ASHES Session 1: Programing Model and Performance Optimizations

Synchronization and Ordering Semantics in Hybrid MPI+GPU Programming	1020
<i>Ashwin M. Aji, Pavan Balaji, James Dinan, Wu-chun Feng, and Rajeev Thakur</i>	
Tightly Coupled Accelerators Architecture for Minimizing Communication Latency among Accelerators	1030
<i>Toshihiro Hanawa, Yuetsu Kodama, Taisuke Boku, and Mitsuhsisa Sato</i>	
Analyzing Optimization Techniques for Power Efficiency on Heterogeneous Platforms	1040
<i>Yash Ukidave and David R. Kaeli</i>	
Composing Multiple StarPU Applications over Heterogeneous Machines: A Supervised Approach	1050
<i>Andra-Ecaterina Hugo, Abdou Guermouche, Pierre-Andre Wacrenier, and Raymond Namyst</i>	

ASHES Session 2: Accelerated Applications

Fast, Scalable Parallel Comparison Sort on Hybrid Multicore Architectures	1060
<i>Dip Sankar Banerjee, Parikshit Sakurikar, and Kishore Kothapalli</i>	
Tridiagonalization of a Symmetric Dense Matrix on a GPU Cluster	1070
<i>Ichitaro Yamazaki, Tingxing Dong, Stanimire Tomov, and Jack Dongarra</i>	
A Multi-Level Optimization Method for Stencil Computation on the Domain that is Bigger than Memory Capacity of GPU	1080
<i>Guanghao Jin, Toshio Endo, and Satoshi Matsuoka</i>	
Accelerating the 3D Elastic Wave Forward Modeling on GPU and MIC	1088
<i>Yang You, Haohuan Fu, Xiaomeng Huang, Guojie Song, Lin Gan, Wenjian Yu, and Guangwen Yang</i>	

ASHES Session 3: Emerging Hybrid Systems

Improving GPU Performance Prediction with Data Transfer Modeling	1097
<i>Michael Boyer, Jiayuan Meng, and Kalyan Kumaran</i>	
Use of SIMD Vector Operations to Accelerate Application Code Performance on Low-Powered ARM and Intel Platforms	1107
<i>Gaurav Mitra, Beau Johnston, Alistair P. Rendell, Eric McCreath, and Jun Zhou</i>	
Using MIC to Accelerate a Typical Data-Intensive Application: The Breadth-first Search	1117
<i>Gao Tao, Lu Yutong, and Suo Guang</i>	
Dynamic Load Balancing of the Adaptive Fast Multipole Method in Heterogeneous Systems	1126
<i>Robert E. Overman, Jan F. Prins, Laura A. Miller, and Michael L. Minion</i>	

Multicore and GPU Programming Models, Languages, and Compilers Workshop—PLC

PLC Introduction	1136
<i>Barbara Chapman and Weimin Zheng</i>	

PLC Session 1: Languages Extensions and Compilers for Multicore-SIMD Processors and GPUs

clMPI: An OpenCL Extension for Interoperation with the Message Passing Interface	1138
<i>Hiroyuki Takizawa, Makoto Sugawara, Shoichi Hirasawa, Isaac Gelado, Hiroaki Kobayashi, and Wen-Mei W. Hwu</i>	
Practical SIMD Vectorization Techniques for Intel® Xeon Phi™ Coprocessors	1149
<i>Xinmin Tian, Hideki Saito, Serguei V. Preis, Eric N. Garcia, Sergey S. Kozhukhov, Matt Masten, Aleksei G. Cherkasov, and Nikolay Panchenko</i>	
Mis-speculation-Driven Compiler Framework for Aggressive Loop Automatic Parallelization	1159
<i>Jin Lin, Xinmin Tian, and John Ng</i>	

PLC Session 2: Programming and Applications for MultiCore Processors and GPUs

Exploring Programming Multi-GPUs Using OpenMP and OpenACC-Based Hybrid Model	1169
<i>Rengan Xu, Sunita Chandrasekaran, and Barbara Chapman</i>	
OpenCL Performance Evaluation on Modern Multi Core CPUs	1177
<i>Joo Hwan Lee, Kaushik Patel, Nimit Nigania, Hyojong Kim, and Hyesoon Kim</i>	
Towards Virtual Shared Memory for Non-cache-coherent Multicore Systems	1186
<i>Bharath Ramesh, Calvin J. Ribbens, and Srinidhi Varadarajan</i>	
Unstructured Control Flow in GPGPU	1194
<i>Rodrigo Domínguez and David R. Kaeli</i>	

PLC Session 3: Runtime Implementation and Optimizations

An Abstract Object Oriented Runtime System for Heterogeneous Parallel Architecture	1203
<i>Jean-Marc Gratien</i>	
Offload Compiler Runtime for the Intel® Xeon Phi™ Coprocessor	1213
<i>Chris J. Newburn, Serguei Dmitriev, Ravi Narayanaswamy, John Wiegert, Ravi Murty, Francisco Chinchilla, Rajiv Deodhar, and Russ McGuire</i>	
Reducing the Cost of Measuring Memory Hierarchy Communication Parameters	1226
<i>Feng Jiang and Alan Sussman</i>	

NSF/TCPP Workshop on Parallel and Distributed Computing Education—EduPar

EduPar Introduction	1234
<i>Sushil K. Prasad</i>	

EduPar Session: Liberal Arts / Small Colleges

Integrating Parallel and Distributed Computing Topics into an Undergraduate CS Curriculum	1237
<i>Andrew Danner and Tia Newhall</i>	
Patterns and Exemplars: Compelling Strategies for Teaching Parallel and Distributed Computing to CS Undergraduates	1244
<i>Joel Adams, Richard Brown, and Elizabeth Shoop</i>	

EduPar Session: Introductory Courses (CS0/CS1/CS2)

Visualizing Parallelism in CS 2	1252
<i>Sean Massung and Cinda Heeren</i>	
Using Actors and the SALSA Programming Language to Introduce Concurrency in Computer Science II	1257
<i>Travis Desell</i>	
Hands-On Exploration of Parallelism for Absolute Beginners with Scratch	1263
<i>Steven Bogaerts</i>	

EduPar Session: Systems/Architecture Courses

Introducing Parallel Programming in Undergraduate Curriculum	1269
<i>Cordelia M. Brown, Yung-Hsiang Lu, and Samuel Midkiff</i>	
Adding GPU Computing to Computer Organization Courses	1275
<i>David Bunde, Karen L. Karavanic, Jens Mache, and Christopher T. Mitchell</i>	
A New Methodology for Studying Realistic Processors in Computer Science Degrees	1283
<i>Crispin Gómez Requena, Maria E. Gómez, and Julio Sahuquillo</i>	

EduPar Session: Multiple Course Adoptions and Techniques

Toward Using Higher-Level Abstractions to Teach Parallel Computing	1291
<i>Clayton Ferner, Barry Wilkinson, and Barbara Heath</i>	
Teaching Parallel and Distributed Computing to Undergraduate Computer Science Students	1297
<i>Marcelo Arroyo</i>	
Programming with Concurrency: Threads, Actors, and Coroutines	1304
<i>Zhen Li and Eileen Kraemer</i>	
Teaching Parallel and Distributed Computing Using a Cluster Computing Portal	1312
<i>Hong Lin</i>	
Teaching Concurrent and Distributed Computing—Initiatives in Rio de Janeiro	1318
<i>Adriano Branco, Ana Lúcia de Moura, Noemi Rodriguez, and Silvana Rossetto</i>	

Workshop on Parallel and Distributed Scientific and Engineering Computing—PDSEC

PDSEC Introduction	1324
<i>Peter Strazdins, Neal N. Xiong, Thomas Rauber, Yinglong Xia, Laurence T. Yang, and Gudula Rünger</i>	

PDSEC Session 1: Performance Evaluation on Large-Scale Systems

Performance Characteristics of Hardware Transactional Memory for Molecular Dynamics Application on BlueGene/Q: Toward Efficient Multithreading Strategies for Large-Scale Scientific Applications	1326
<i>Manaschai Kunaseth, Rajiv K. Kalia, Aiichiro Nakano, Priya Vashishta, David F. Richards, and James N. Glosli</i>	
Increasing the Scalability of PISM for High Resolution Ice Sheet Models	1336
<i>Phillip Dickens and Timothy Morey</i>	
Understanding the Performance of Two Production Supercomputers	1345
<i>Scott Pakin and Michael Lang</i>	

PDSEC Session 2: Applications via Heterogeneous Computing

Performance Dissection of a Molecular Dynamics Code across CUDA and GPU Generations	1355
<i>Matthew Wezowicz, Trilce Estrada, Sandeep Patel, and Michela Taufer</i>	
Leveraging GPUs in Ab Initio Nuclear Physics Calculations	1365
<i>Dossay Oryspayev, Hugh Potter, Pieter Maris, Masha Sosonkina, James P. Vary, Sven Binder, Angelo Calci, Joachim Langhammer, and Robert Roth</i>	
Orthogonalization on a General Purpose Graphics Processing Unit with Double Double and Quad Double Arithmetic	1373
<i>Jan Verschelde and Genady Yoffe</i>	

PDSEC Session 3: Linear Algebra and Infrastructure

A NUMA-Aware Fine Grain Parallelization Framework for Multi-core Architecture	1381
<i>Corentin Rossignon, Pascal Hénon, Olivier Aumage, and Samuel Thibault</i>	
On Partitioning and Reordering Problems in a Hierarchically Parallel Hybrid Linear Solver	1391
<i>Ichitaro Yamazaki, Xiaoye S. Li, Francois-Henry Rouet, and Bora Uçar</i>	
Scheduling a Parallel Sparse Direct Solver to Multiple GPUs	1401
<i>Kyungjoo Kim and Victor Eijkhout</i>	
A Scalable Implicit Solver for Phase Field Crystal Simulations	1409
<i>Chao Yang and Xiao-Chuan Cai</i>	

PDSEC Session 4: Cloud, Grid, and Distributed Systems

An Efficient Scheduling Model for Broadcasting in Wireless Sensor Networks	1417
<i>Hongju Cheng, Naixue Xiong, Xingbo Huang, and Laurence T. Yang</i>	
Evaluating the Flexibility of Dynamic Loop Scheduling on Heterogeneous Systems in the Presence of Fluctuating Load Using SimGrid	1429
<i>Nitin Sukhija, Ioana Banicescu, Srishti Srivastava, and Florina M. Ciorba</i>	
Efficient and Fault-Tolerant Static Scheduling for Grids	1439
<i>Patrick Cichowski and Jörg Keller</i>	
Briareus: Accelerating Python Applications with Cloud	1449
<i>Zhaomeng Zhu, Gongxuan Zhang, Yongping Zhang, Jian Guo, and Naixue Xiong</i>	

Dependable Parallel, Distributed, and Network-Centric Systems—DPDNS

DPDNS Introduction	1457
<i>Dimiter Avresky, Erik Maehle, and Salvatore Distefano</i>	

DPDNS Session: Network Algorithms and Techniques

A Fault-Tolerant Strong Conjunctive Predicate Detection Algorithm for Large-Scale Networks	1460
<i>Min Shen and Ajay D. Kshemkalyani</i>	
Tolerating Packet Losses in Wireless Mesh Networks	1470
<i>Frank Engelhardt, Timo Lindhorst, and Edgar Nett</i>	
Reducing the Migration Times of Multiple VMs on WANs Using a Feedback Controller	1480
<i>Tae Seung Kang, Mauricio Tsugawa, José Fortes, and Takahiro Hirofuchi</i>	

DPDNS Session: Cloud Computing

DRACO PaaS: A Distributed Resilient Adaptable Cloud Oriented Platform	1490
<i>Antonio Celesti, Nicola Peditto, Fabio Verboso, Massimo Villari, and Antonio Puliafito</i>	
Sustained Resilience via Live Process Cloning	1498
<i>Arash Rezaei and Frank Mueller</i>	
A Model for Evaluation of User-Perceived Service Properties	1508
<i>Andreas Dittrich, Igor Kaitovic, Cristina Murillo, and Rafael Rezende</i>	

DPDNS Session: High Performance/Availability Computing

Symmetric Comparison-Based Fault Diagnosis of Multiprocessor and Distributed Systems Using Nonlinear Support Vector Machines	1518
<i>Mourad Elhadef</i>	
A Comprehensive Analysis of XOR-Based Erasure Codes Tolerating 3 or More Concurrent Failures	1528
<i>Pradeep Subedi and Xubin He</i>	
Monitoring and Controlling System for Microservers	1538
<i>Andreas C. Doering and Tibor Kiss</i>	

DPDNS Session: Fault Tolerance

Investigating Standby from a System Reliability Perspective	1542
<i>Salvatore Distefano</i>	
Towards Fault-Tolerant and Energy-Efficient Swarms of Underwater Robots	1550
<i>Ammar Amory, Benjamin Meyer, Christoph Osterloh, Thomas Tosik, and Erik Maehle</i>	

Workshop on Multi-Threaded Architectures and Applications—MTAAP

MTAAP Introduction	1554
<i>Luiz DeRose</i>	

MTAAP Session: Compiler and Tools

Capping Speculative Traces to Improve Performance in Simultaneous Multi-threading CPUs	1555
<i>Yilin Zhang and Wei-Ming Lin</i>	
CHiP: A Profiler to Measure the Effect of Cache Contention on Scalability	1565
<i>Bevin Brett, Pranith Kumar, Minjang Kim, and Hyesoon Kim</i>	
Compiler-Based Data Prefetching and Streaming Non-temporal Store Generation for the Intel(R) Xeon Phi(TM) Coprocessor	1575
<i>Rakesh Krishnaiyer, Emre Kultursay, Pankaj Chawla, Serguei Preis, Anatoly Zvezdin, and Hideki Saito</i>	

MTAAP Session: Scheduling and Runtime

The Pheet Task-Scheduling Framework on the Intel® Xeon Phi™ Coprocessor and other Multicore Architectures	1587
<i>Martin Wimmer, Manuel Pöter, and Jesper Larsson Träff</i>	
Toward a Scalable Heterogeneous Runtime System for the Convey MX Architecture	1597
<i>John D. Leidel, Joe Bolding, and Geoffrey Rogers</i>	
Towards Memory-Load Balanced Fast Fourier Transformations in Fine-Grain Execution Models	1607
<i>Chen Chen, Yao Wu, Stéphane Zuckerman, and Guang R. Gao</i>	

MTAAP Session: Graph Algorithms

Distributed Memory Breadth-First Search Revisited: Enabling Bottom-Up Search	1618
<i>Scott Beamer, Aydin Buluç, Krste Asanović, and David Patterson</i>	
Avoiding Locks and Atomic Instructions in Shared-Memory Parallel BFS Using Optimistic Parallelization	1628
<i>Jesmin Jahan Tithi, Dhruv Matani, Gaurav Menghani, and Rezaul A. Chowdhury</i>	
Investigating Graph Algorithms in the BSP Model on the Cray XMT	1638
<i>David Ediger and David A. Bader</i>	

MTAAP Session: Algorithms

Multithreaded Community Monitoring for Massive Streaming Graph Data	1646
<i>Jason Riedy and David A. Bader</i>	
Scalable, Multithreaded, Partially-in-Place Sorting	1656
<i>David J. Haglin, Robert D. Adolf, and Greg E. Mackey</i>	

Workshop on Large-Scale Parallel Processing—LSPP

LSPP Introduction	1665
<i>Darren J. Kerbyson, Ram Rajamony, and Charles Weems</i>	

LSPP Session 1: Performance Analysis and Optimization

Inferring Large-Scale Computation Behavior via Trace Extrapolation	1667
<i>Laura Carrington, Michael A. Laurenzano, and Ananta Tiwari</i>	
Systematic Reduction of Data Movement in Algebraic Multigrid Solvers	1675
<i>Hormozd Gahvari, William Gropp, Kirk E. Jordan, Martin Schulz, and Ulrike Meier Yang</i>	

LSPP Session 2: Modeling Performance for Scaling

Dataset Scaling and MapReduce Performance	1683
<i>Fan Zhang and Majd Sakr</i>	
Tabulated Equations of State with a Many-tasking Execution Model	1691
<i>Matthew Anderson, Maciej Brodowicz, Thomas Sterling, Hartmut Kaiser, and Bryce Adelstein-Lelbach</i>	
Toward Automatic Optimized Code Generation for Multiprecision Modular Exponentiation on a GPU	1700
<i>Niall Emmart and Charles Weems</i>	

LSPP Session 3: Large-Scale Systems

An Evaluation of Different I/O Techniques for Checkpoint/Restart	1708
<i>Faisal Shahzad, Markus Wittmann, Thomas Zeiser, Georg Hager, and Gerhard Wellein</i>	
Application Explorations for Future Interconnects	1717
<i>R.F. Barrett, C.T. Vaughan, S.D. Hammond, and D. Roweth</i>	

LSPP Session 4: Scheduling

High Performance Adaptive Distributed Scheduling Algorithm	1725
<i>Ankur Narang, Abhinav Srivastava, and R.K. Shyamasundar</i>	
Scalable Loop Self-Scheduling Schemes Implemented on Large-Scale Clusters	1735
<i>Yiming Han and Anthony T. Chronopoulos</i>	

Parallel Computing and Optimization—PCO

PCO Introduction	1743
<i>Didier El Baz</i>	

PCO Session 1: Algorithms

Semi-Matching Algorithms for Scheduling Parallel Tasks under Resource Constraints	1744
<i>Anne Benoit, Johannes Langguth, and Bora Uçar</i>	
Sequential and Parallel Restart Policies for Constraint-Based Local Search	1754
<i>Yves Caniou and Philippe Codognet</i>	
On a Hybrid MPI-Pthread Approach for Simplicial Branch-and-Bound	1764
<i>Juan F.R. Herrera, Leocadio G. Casado, Remigijus Paulavičius, Julius Žilinskas, and Eligius M.T. Hendrix</i>	
Anticipated Dynamic Load Balancing Strategy to Parallelize Constraint Programming Search	1771
<i>Tarek Menouer and Bertrand Le Cun</i>	

PCO Session 2: GPU Computing and Optimization

Recent Advances on GPU Computing in Operations Research	1778
<i>Vincent Boyer and Didier El Baz</i>	
High Performance GPU Accelerated Local Optimization in TSP	1788
<i>Kamil Rocki and Reiji Suda</i>	
An Efficient Parallelization Strategy for Dynamic Programming on GPU	1797
<i>Karl-Eduard Berger and Francois Galea</i>	
High Throughput Parallel Implementation of Aho-Corasick Algorithm on a GPU	1807
<i>Nhat-Phuong Tran, Myungho Lee, Sugwon Hong, and Jaeyoung Choi</i>	

PCO Session 3

Task Scheduling Greedy Heuristics for GPU Heterogeneous Cluster Involving the Weights of the Processor	1817
<i>Keliang Zhang and Baifeng Wu</i>	
On the Optimality and Speed of the Deep Greedy Switching Algorithm for Linear Assignment Problems	1828
<i>Amgad Naiem and Mohammed El-Beltagy</i>	
Parallel Algorithms for Graph Optimization Using Tree Decompositions	1838
<i>Blair D. Sullivan, Dinesh Weerapurage, and Chris Groër</i>	
Subdomain Mapping Approach to Enhance the Coupling in Earth System Modeling	1848
<i>Yingsheng Ji, Guangwen Yang, Li Liu, and Shu Wang</i>	

Parallel and Distributed Computing for Machine Learning and Inference Problems—ParLearning

ParLearning Introduction	1856
<i>Sutanay Choudhury, George Chin, and Yinglong Xia</i>	
Combining Parallel Algorithms Solving the Same Application: What is the Best Approach?	1859
<i>Alfredo Goldman, Joachim Lepping, Yanik Ngoko, and Denis Trystram</i>	
Enhancing Accuracy and Performance of Collaborative Filtering Algorithm by Stochastic SVD and Its MapReduce Implementation	1869
<i>Che-Rung Lee and Ya-Fang Chang</i>	
Reducing False Transactional Conflicts with Speculative Sub-Blocking State—An Empirical Study for ASF Transactional Memory System	1879
<i>Lifeng Nai and Hsien-Hsin S. Lee</i>	
Revisiting a Pattern for Processing Combinatorial Objects in Parallel	1889
<i>Christian Trefftz and Jerry Scripps</i>	
EDA and ML—A Perfect Pair for Large-Scale Data Analysis	1894
<i>Ryan Hafen and Terence Critchlow</i>	
Combining Structure and Property Values is Essential for Graph-Based Learning	1899
<i>David J. Haglin and Lawrence B. Holder</i>	

High Performance Data Intensive Computing—HPDIC

HPDIC Introduction	1905
<i>Christophe Cérin, Cong-Feng Jiang, Yuqing Gao, and Jilin Zhang</i>	

HPDIC Session 1: MapReduce, Hadoop, and New Architecture Support for Data Intensive Computing

High-Performance RDMA-based Design of Hadoop MapReduce over InfiniBand	1908
<i>Md. Wasi-ur-Rahman, Nusrat Sharmin Islam, Xiaoyi Lu, Jithin Jose, Hari Subramoni, Hao Wang, and Dhabaleswar K. (DK) Panda</i>	
HadoopCL: MapReduce on Distributed Heterogeneous Platforms through Seamless Integration of Hadoop and OpenCL	1918
<i>Max Grossman, Mauricio Breternitz, and Vivek Sarkar</i>	
Minimizing Remote Accesses in MapReduce Clusters	1928
<i>Prateek Tandon, Michael J. Cafarella, and Thomas F. Wenisch</i>	
MapReducing GEPETO or Towards Conducting a Privacy Analysis on Millions of Mobility Traces	1937
<i>Sébastien Gambs, Marc-Olivier Killijian, Izabela Moise, and Miguel Núñez del Prado Cortez</i>	
Resource Provisioning for Staging Components	1947
<i>Tuan Anh Nguyen, Greg Eisenhauer, Karsten Schwan, Matthew Wolf, Hasan Abbasi, Scott Klasky, and Norbert Podhorszki</i>	
BPS: A Performance Metric of I/O System	1954
<i>Shuibing He, Xian-He Sun, and Yanlong Yin</i>	

HPDIC Session 2: File System and Storage for Data Intensive Computing

Transparent Optimization of Parallel File System I/O via Standard System Tool Enhancement	1963
<i>Paul Z. Kolano</i>	
Interference Resolver in Shared Storage Systems to Provide Fairness to I/O Intensive Applications	1971
<i>Ramya Prabhakar, Mahmut Kandemir, Padma Raghavan, and Myoungsoo Jung</i>	
InfoStor: Highly Available Distributed Block Store	1981
<i>YongJian Ren, YouQing Lin, JiLin Zhang, Jian Wan, and Cong-Feng Jiang</i>	
tpNFS: Efficient Support of Small Files Processing over pNFS	1989
<i>Bo Wang, Jinlei Jiang, and Guangwen Yang</i>	
HyCache: A User-Level Caching Middleware for Distributed File Systems	1997
<i>Dongfang Zhao and Ioan Raicu</i>	
Filesystem Aware Scalable I/O Framework for Data-Intensive Parallel Applications	2007
<i>Rengan Xu, Mauricio Araya-Polo, and Barbara Chapman</i>	

HPDIC Session 3: Data Analytics: Solutions and Case Studies

I/O Containers: Managing the Data Analytics and Visualization Pipelines of High End Codes	2015
<i>Jai Dayal, Jianting Cao, Greg Eisenhauer, Karsten Schwan, Matthew Wolf, Fang Zheng, Hasan Abbasi, Scott Klasky, Norbert Podhorszki, and Jay Lofstead</i>	
Proactive Support for Large-Scale Data Exploration	2025
<i>Mark Hereld, Tanu Malik, and Venkatram Vishwanath</i>	
Enhancement for Potential Target in Cryptography Algorithms by Applying Processor-in-Memory Architecture	2035
<i>Jed Kao-Tung Chang, Chen Liu, and Jean-Luc Gaudiot</i>	
Accelerating Dynamics Simulation of Solidification Processes of Liquid Metals Using GPU with CUDA	2045
<i>Jie Liang, KenLi Li, Lin Shi, and Yingqiang Liao</i>	
Data Evolution Analysis of Virtual DataSpace for Managing the Big Data Lifecycle	2054
<i>Xin Cheng, Chungjin Hu, Yang Li, Wei Lin, and Haolei Zuo</i>	
An Image Management System Implemented on Open-Source Cloud Platform	2064
<i>Jian Wan, Shuting Han, Jilin Zhang, Baojin Zhu, and Li Zhou</i>	
Acceleration of a High Order Finite-Difference WENO Scheme for Large-Scale Cosmological Simulations on GPU	2071
<i>Chen Meng, Long Wang, Zongyan Cao, Xianfeng Ye, and Long-Long Feng</i>	
A GPGPU Algorithm for c-Approximate r-Nearest Neighbor Search in High Dimensions	2079
<i>Lee A. Carraher, Philip A. Wilsey, and Fred S. Annexstein</i>	

Workflow Models, Systems, Services, and Applications in the Cloud—CloudFlow

CloudFlow Introduction	2089
<i>Yong Zhao, Cui Lin, and Shiyong Lu</i>	

CloudFlow Session: Modeling

Dataflow Oriented Similarity Matching for Scientific Workflows	2091
<i>Philip Yeo and Syed Sibte Raza Abidi</i>	
Comparing Provisioning and Scheduling Strategies for Workflows on Clouds	2101
<i>Marc E. Frincu, Stéphane Genaud, and Julien Gossa</i>	
Modeling Many-Task Computing Workloads on a Petaflop IBM Blue Gene/P Supercomputer	2111
<i>Ke Wang, Zhangjie Ma, and Ioan Raicu</i>	
Investigating the Feasibility of Making Contexts Explicit in Designing Cloud Workflow	2121
<i>Xiaoliang Fan, Ruisheng Zhang, and Patrick Brézillon</i>	

CloudFlow Session: Application

Accelerating Distributed Workflows with Edge Resources	2129
<i>Siddharth Ramakrishnan, Robert Reutiman, Abhishek Chandra, and Jon Weissman</i>	
AzureBOT: A Framework for Bag-of-Tasks Applications on the Azure Cloud Platform	2139
<i>Dinesh Agarwal and Sushil K. Prasad</i>	
Adaptive Resource Management for Service Workflows in Cloud Environments	2147
<i>Yi Wei, M. Brian Blake, and Iman Saleh</i>	
A Data Intensive Statistical Aggregation Engine: A Case Study for Gridded Climate Records	2157
<i>David Chapman, Tyler A. Simon, Phuong Nguyen, and Milton Halem</i>	

Workshop on Job Scheduling Strategies for Parallel Processing—JSSPP

JSSPP Introduction	2165
<i>Walfredo Cirne and Narayan Desai</i>	

Virtual Prototyping of Parallel and Embedded Systems—ViPES

ViPES Introduction	2166
<i>Michael Hübner</i>	

ViPES Session 1: Novel Designs and Features for Full System Simulators

Experiences with Dynamic Binary Translation in a Full System Simulator	2168
<i>Aditya Kumar, Ahmed Gheith, and Michael Kistler</i>	
Co-simulation of Functional SystemC TLM Models with Power/Thermal Solvers	2176
<i>Tayeb Bouhadiba, Matthieu Moy, Florence Maraninchi, Jerome Cornet, Laurent Maillet-Contoz, and Ilija Materic</i>	

VirtualSoC: A Full-System Simulation Environment for Massively Parallel Heterogeneous System-on-Chip	2182
<i>Daniele Bortolotti, Christian Pinto, Andrea Marongiu, Martino Ruggiero, and Luca Benini</i>	
legaSCi: Legacy SystemC Model Integration into Parallel Systemc Simulators	2188
<i>Christoph Schumacher, Jan Henrik Weinstock, Rainer Leupers, Gerd Ascheid, Laura Tosoratto, Alessandro Lonardo, Dietmar Petras, and Thorsten Grötter</i>	

ViPES Session 2: Simulation Frameworks and Fault Tolerance for Virtual FPGAs

HVSoCs: A Framework for Rapid Prototyping of 3-D Hybrid Virtual System-on-Chips	2194
<i>D. Diamantopoulos, Kostas Siozios, E. Sotiriou-Xanthopoulos, G. Economakos, and Dimitrios Soudris</i>	
Simplify: A Framework for Enabling Fast Functional/Behavioral Validation of Multiprocessor Architectures in the Cloud	2200
<i>Gabriel Marchesan Almeida, Oliver Bellaver Longhi, Thomas Bruckschloegl, Michael Hübner, Fabiano Hessel, and Jürgen Becker</i>	
On Supporting Adaptive Fault Tolerant at Run-Time with Virtual FPGAs	2206
<i>Kostas Siozios, Dimitrios Soudris, and Michael Hübner</i>	

PhD Forum

Message from the PhD Forum Co-chairs	2212
<i>Luc Bougé and Bo Hong</i>	

PhD Forum: Algorithms

Algorithm/Architecture Codesign of Low Power and High Performance Linear Algebra Compute Fabrics	2214
<i>Ardavan Pedram</i>	
Energy Efficient Workflow Job Scheduling for Green Cloud	2218
<i>Fei Cao and Michelle M. Zhu</i>	
Toward Flexible and Fast Routing Strategies for Dynamic Network Provisioning	2222
<i>Liudong Zuo and Michelle M. Zhu</i>	
Discrete Min-Energy Scheduling on Restricted Parallel Processors	2226
<i>Xibo Jin, Fa Zhang, and Zhiyong Liu</i>	
LiPS: A Cost-Efficient Data and Task Co-Scheduler for MapReduce	2230
<i>Moussa Ehsan and Radu Sion</i>	
Identifying High betweenness Centrality Vertices in Large Noisy Networks	2234
<i>Vladimir Ufimtsev and Sanjukta Bhowmick</i>	

PhD Forum: Applications

Efficient Parallel and Distributed Algorithms for GIS Polygonal Overlay Processing	2238
<i>Satish Puri and Sushil K. Prasad</i>	
HPC System Software for Regular and Irregular Parallel Applications	2242
<i>Alessandro Morari and Mateo Valero</i>	
Wire Speed IPv6 Forwarding on Multi-core Platforms	2246
<i>Thilan Ganegedara and Viktor K. Prasanna</i>	
A Compression Framework for Multidimensional Scientific Datasets	2250
<i>Tekin Bicer and Gagan Agrawal</i>	

PhD Forum: Architecture

Performance and Power Simulation for Versatile GPGPU Global Memory	2254
<i>Bin Wang and Weikuan Yu</i>	
Exploiting Content Similarity to Improve Memory Performance in Large-Scale High-Performance Computing Systems	2258
<i>Scott Levy</i>	
Designing Hybrid Architectures for Massive-Scale Graph Analysis	2262
<i>David Ediger and David A. Bader</i>	
Reducing the Environmental Impact of Optical Networks	2266
<i>Thilo Schöndienst and Vinod M. Vokkarane</i>	
Fine-Grained Manipulation of FPGA Configuration for Incremental Design	2270
<i>Wenwei Zha and Peter Athanas</i>	
Applications Acceleration through Adaptive Hardware Components	2274
<i>Vito Giovanni Castellana and Fabrizio Ferrandi</i>	

PhD Forum: Software

SAGE: Geo-Distributed Streaming Data Analysis in Clouds	2278
<i>Radu Tudoran, Gabriel Antoniu, and Luc Bougé</i>	
Towards Dependability Testing of MapReduce Systems	2282
<i>João Eugenio Marynowski</i>	
Efficient I/O using Dedicated Cores in Large-Scale HPC Simulations	2286
<i>Matthieu Dorier</i>	
Self-Adaptive Cost-Efficient Consistency Management in the Cloud	2290
<i>Housseem-Eddine Chihoub</i>	
Towards Efficient Mapping, Scheduling, and Execution of HPC Applications on Platforms in Cloud	2294
<i>Abhishek Gupta and Laxmikant V. Kalé</i>	
Harnessing Adaptivity Analysis for the Automatic Design of Efficient Embedded and HPC Systems	2298
<i>Silvia Lovergine and Fabrizio Ferrandi</i>	

Adaptive Power and Resource Management Techniques for Multi-threaded Workloads2302
Can Hankendi and Ayse K. Coskun

Author Index