

**2012 IEEE 26th International
Parallel and Distributed
Processing Symposium
Workshops & PhD Forum**

(IPDPSW 2012)

**Shanghai, China
21 – 25 May 2012**

Pages 1-854



**IEEE Catalog Number: CFP1251J-PRT
ISBN: 978-1-4673-0974-5**

2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum

IPDPSW 2012

Table of Contents

Message from the IPDPS General Co-Chairs to the Workshops.....	xxxiii
Message from the Workshops Chair.....	xxxv
Message from the PhD Forum Co-Chairs.....	xxxvii

Heterogeneity in Computing Workshop - HCW

HCW Introduction	1
<i>Behrooz Shirazi and Alexey Lastovetsky</i>	
Message from the HCW Steering Committee Chair	3
<i>Howard Jay Siegel</i>	
Message from the HCW General Chair	4
<i>Behrooz Shirazi</i>	
Message from the HCW Program Chair	5
<i>Alexey Lastovetsky</i>	
HCW 2012 Keynote Talk	6
<i>David A. Bader</i>	

Session 1: Optimization and Evaluation of Heterogeneous Platforms

Experiences with the Sparse Matrix-Vector Multiplication on a Many-core Processor	7
<i>Juan C. Pichel and Francisco F. Rivera</i>	
Performance Benefits of Heterogeneous Computing in HPC Workloads	16
<i>Victor W. Lee, Ed Grochowski, and Robert Geva</i>	
Thermal-Aware Performance Optimization in Power Constrained Heterogenous Data Centers	27
<i>Abdulla M. Al-Qawasmeh, Sudeep Pasricha, Anthony M. Maciejewski, and Howard Jay Siegel</i>	
Experiences with Target-Platform Heterogeneity in Clouds, Grids, and On-Premises Resources	41
<i>Jaroslav Slawinski, Tiziano Passerini, Umberto Villa, Alessandro Veneziani, and Vaidy Sunderam</i>	

BLOR: Bandwidth and Latency Sensitive Overlay Routing for Flash Data Dissemination	53
<i>Xiaoyong Li, Yijie Wang, Yongquan Fu, Xiaoling Li, and Weidong Sun</i>	
Session 2: Scheduling and Mapping in Heterogeneous Environments I	
Scheduling Batch and Heterogeneous Jobs with Runtime Elasticity in a Parallel Processing Environment	65
<i>Dinesh Kumar, Zon-yin Shae, and Hani Jamjoom</i>	
Task Scheduling in Large-scale Distributed Systems Utilizing Partial Reconfigurable Processing Elements	79
<i>M. Faisal Nadeem, Imran Ashraf, S. Arash Ostadzadeh, Stephan Wong, and Koen Bertels</i>	
Mixed Data-Parallel Scheduling for Distributed Continuous Integration	91
<i>Olivier Beaumont, Nicolas Bonichon, Ludovic Courtès, Eelco Dolstra, and Xavier Hanin</i>	
A Monte-Carlo Approach for Full-Ahead Stochastic DAG Scheduling	99
<i>Wei Zheng and Rizos Sakellariou</i>	
Session 3: Algorithms and Methods for Scientific Computing on Heterogeneous Platforms	
A Block-Asynchronous Relaxation Method for Graphics Processing Units	113
<i>Hartwig Anzt, Stanimire Tomov, Jack Dongarra, and Vincent Heuveline</i>	
Partitioning for Parallel Matrix-Matrix Multiplication with Heterogeneous Processors: The Optimal Solution	125
<i>Ashley DeFlumere, Alexey Lastovetsky, and Brett A. Becker</i>	
A Fast Parallel Implementation of Molecular Dynamics with the Morse Potential on a Heterogeneous Petascale Supercomputer	140
<i>Qiang Wu, Canqun Yang, Feng Wang, and Jingling Xue</i>	
Session 4: Heterogeneous Parallel and Distributed Programming	
High-Performance Distributed Multi-Model / Multi-Kernel Simulations: A Case-Study in Jungle Computing	150
<i>Niels Drost, Jason Maassen, Maarten A.J. van Meersbergen, Henri E. Bal, F. Inti Pelupessy, Simon Portegies Zwart, Michael Kliphuis, Henk A. Dijkstra, and Frank J. Seinstra</i>	
A Portable High-Productivity Approach to Program Heterogeneous Systems	163
<i>Zeki Bozkus and Basilio B. Fraguela</i>	
dOpenCL: Towards a Uniform Programming Approach for Distributed Heterogeneous Multi-/Many-Core Systems	174
<i>Philipp Kegel, Michel Steuwer, and Sergei Gorbachev</i>	

Session 5: Scheduling and Mapping in Heterogeneous Environments II

Scalable Communication-aware Task Mapping Algorithms for Interconnected Multicore Systems	187
<i>I-Hsin Chung, Che-Rung Lee, Jiazheng Zhou, Chung-Yi Chou, and Yeh-Ching Chung</i>	
A Combined Dual-stage Framework for Robust Scheduling of Scientific Applications in Heterogeneous Environments with Uncertain Availability	193
<i>Florina M. Ciorba, Timothy Hansen, Srishti Srivastava, Ioana Banicescu, Anthony A. Maciejewski, and Howard Jay Siegel</i>	

Reconfigurable Architectures Workshop - RAW

RAW Introduction	208
<i>Juergen Becker, Jinian Bian, Christophe Bobda, Rene Cumplido, and Michael Huebner</i>	

Session 1 - Physical Design of Partially Reconfigurable Architectures

Designing Nonvolatile Reconfigurable Switch-based FPGA through Overall Circuit Performance Evaluation	213
<i>Kazutaka Ikegami, Keiko Abe, Kumiko Nomura, Shinichi Yasuda, Masato Oda, and Shinobu Fujita</i>	
A Power and Cluster-Aware Technology Mapping and Clustering Scheme for Dual-VT FPGAs	221
<i>Wei Ting Loke, Yajun Ha, and Wenfeng Zhao</i>	
A Comparison of DAG and Mesh Topologies for Coarse-Grain Reconfigurable Array	227
<i>Jonathan Antusiak, Antoine Trouvé, and Kazuaki Murakami</i>	

Session 2 - Network on Chip for Reconfigurable Hardware

Hardware-assisted Decentralized Resource Management for Networks on Chip with QoS	234
<i>Jan Heißwolf, Aurang Zaib, Andreas Weichslgartner, Ralf König, Thomas Wild, Jürgen Teich, Andreas Herkersdorf, and Jürgen Becker</i>	
Self-Correction Trace Model: A Full-System Simulator for Optical Network-on-Chip	242
<i>Mingzhe Zhang, Liqiang He, and Dongrui Fan</i>	
Real-Time Monitoring of Multicore SoCs through Specialized Hardware Agents on NoC Network Interfaces	248
<i>George Kornaros and Dionisios Pnevmatikatos</i>	

Session 3 - General Session 1

High Speed – Low Power Optical Configuration on an ORGA with a Phase-modulation Type Holographic Memory	256
<i>Takahiro Watanabe and Minoru Watanabe</i>	
Efficient Reconfiguration Algorithm for Three-dimensional VLSI Arrays	261
<i>Guiyuan Jiang, Wu Jigang, and Jizhou Sun</i>	

Algorithm for Communication Synchronization on Reconfigurable Processor Arrays with Faults	266
<i>Wu Jigang, Guiyuan Jiang, Yuanrui Zhang, and Yuanbo Zhu</i>	
A Heterogeneous Cache Distribution with Reconfigurable Interconnect	271
<i>Aishwariya Pattabiraman, Annie Avakian, and Ranga Vemuri</i>	
Study of an Automated Precise SEU Fault Injection Technique	277
<i>Zhou Jing, Liu Zengrong, Chen Lei, Wang Shuo, Wen Zhiping, Wu Lishuai, and Chen Xun</i>	
Detecting Data Hazards in Multi-Processor System-on-Chips on FPGA	282
<i>Chao Wang, Xi Li, Peng Chen, Xiaojing Feng, Junneng Zhang, and Xuehai Zhou</i>	
Mapping Algorithm for Coarse-Grained Reconfigurable Multimedia Architectures	288
<i>Naijin Chen and Jianhui Jiang</i>	
Reconfigurable Designs for Networking Silicon	294
<i>Tao Li, Zhentao Liu, Huimin Du, Lei Zhang, Jungang Han, Lin Jiang, and Qingang Dong</i>	
Fair Access to External Memory for Chip-multiprocessor	300
<i>Shufan Yang, Qiang Wu, Xiongren Xiao, Renfa Li, and Dominic Hillenbrand</i>	
Self-Adaptive Heterogeneous Cluster with Wireless Network	306
<i>Xinyu Niu, Kuen Hung Tsoi, and Wayne Luk</i>	
 Session 4 - Improving Security and Computing Efficiency of Reconfigurable Systems	
Managing Dynamic Reconfiguration for Fault-tolerance on a Manycore Architecture	312
<i>Zain ul-Abdin, Essayas Gebrewahid, and Bertil Svensson</i>	
Using Run-Time Reconfiguration to Implement Fault-Tolerant Coarse Grained Reconfigurable Architectures	320
<i>Thomas Schweizer, Anja Küster, Sven Eisenhardt, Tommy Kuhn, and Wolfgang Rosenstiel</i>	
On Supporting Efficient Partial Reconfiguration with Just-In-Time Compilation	328
<i>Harry Sidiropoulos, Kostas Siozios, Peter Figuli, Dimitrios Soudris, and Michael Hubner</i>	
An Enhanced Relocation Manager to Speedup Core Allocation in FPGA-based Reconfigurable Systems	336
<i>M.D. Santambrogio, F. Cancare, R. Cattaneo, S. Bhandariy, and D. Sciuto</i>	
Classification of Massively Parallel Computer Architectures	344
<i>Muhammad Ali Shami and Ahmed Hemani</i>	
 Session 5 - Applications and Special Purpose Architectures with Reconfigurable Hardware 1	
An Optimized Reconfigurable System for Computing the Phylogenetic Likelihood Function on DNA Data	352
<i>Simon A. Berger, Nikolaos Alachiotis, and Alexandros Stamatakis</i>	
FPGA-based Router Virtualization: A Power Perspective	360
<i>Thilan Ganegedara and Viktor K. Prasanna</i>	

Session 6 - Applications and Special Purpose Architectures with Reconfigurable Hardware 2

A Reconfigurable High Performance ASIP Engine for Image Signal Processing	368
<i>Hsuanchun Liao, Mochamad Asri, Tsuyoshi Isshiki, Dongju Li, and Hiroaki Kunieda</i>	
Area-Efficient FPGA Implementation of Quadruple Precision Floating Point Multiplier	376
<i>Manish Kumar Jaiswal and Ray C.C. Cheung</i>	
FPGA Implementation of SRAM-based Ternary Content Addressable Memory	383
<i>Zahid Ullah, Manish Kumar Jaiswal, Y.C. Chan, and Ray C.C. Cheung</i>	
Improved Bioinformatics Processing Unit for Multiple Applications	390
<i>Pei Liu, Ahmed Hemani, and Kolin Paul</i>	

Session 7 - General Session 2

RIVER: Reconfigurable Pre-Synthesized-Streaming Architecture for Signal Processing on FPGAs	397
<i>Dominic Hillenbrand, Christian Brugger, Jie Tao, Shufan Yang, and Matthias Balzer</i>	
Efficient On-line Hardware/Software Task Scheduling for Dynamic Run-time Reconfigurable Systems	401
<i>Ahmed Al-Wattar, Shawki Areibi, and Faycal Saffih</i>	
Model-Driven Approach for Automatic Dynamic Partially Reconfigurable IP Customization	407
<i>Gilberto Ochoa-Ruiz, Ouassila Labbani-Narsis, El-Bay Bourennane, and Phillipe Soulard</i>	
Embodied Computing: Self-adaptation in Bio-inspired Reconfigurable Architectures	413
<i>Laurent Rodriguez, Benoît Miramond, Imen Kalbousi, and Bertrand Granado</i>	
On Dynamic Run-time Processor Pipeline Reconfiguration	419
<i>Carsten Tradowsky, Florian Thoma, Michael Hübner, and Jürgen Becker</i>	
Pareto Optimal Temporal Partition Methodology for Reconfigurable Architectures Based on Multi-objective Genetic Algorithm	425
<i>Weiguang Sheng, Weifeng He, Jianfei Jiang, and Zhigang Mao</i>	
Hardware Index to Permutation Converter	431
<i>Jon T. Butler and Tsutomu Sasao</i>	
Mini-Robot's Performance Optimization via Online Reconfiguration and HW/SW Task Scheduling	437
<i>Gianluca Durelli, Federica Cresci, Donatella Sciuto, Mario Porrman, and Marco D. Santambrogio</i>	
SMPP: Generic SAT Solver over Reconfigurable Hardware Accelerator	443
<i>Zhongda Yuan, Yuchun Ma, and Jinian Bian</i>	
A High-Performance FPGA-Based Implementation of the LZSS Compression Algorithm	449
<i>Ivan Shcherbakov, Christian Weis, and Norbert Wehn</i>	

Session 8 - Tools for Partially Reconfigurable FPGAs

DGECS: Description Generator for Evolved Circuits Synthesis	454
<i>Fabio Cancare, Davide B. Bartolini, Matteo Carminati, Donatella Sciuto, and Marco D. Santambrogio</i>	
A Compiler Back-End for Reconfigurable, Mixed-ISA Processors with Clustered Register Files	462
<i>Timo Stripf, Ralf Koenig, Patrick Rieder, and Juergen Becker</i>	
Modeling for Synthesis with System#	470
<i>C. Köllner, F. Mendoza, and K.D. Müller-Glaser</i>	
FPM: A Flexible Programming Model for MPSoC on FPGA	477
<i>Chao Wang, Xi Li, Junneng Zhang, Peng Chen, Xiaojing Feng, and Xuehai Zhou</i>	

High-Level Parallel Programming Models & Supportive Environments - HIPS

HIPS Introduction	485
<i>Matthias S. Müller</i>	

Case Studies

Improving High-Performance Sparse Libraries Using Compiler-Assisted Specialization: A PETSc Case Study	487
<i>Shreyas Ramalingam, Mary Hall, and Chun Chen</i>	
An Empirical Performance Study of Chapel Programming Language	497
<i>Nan Dun and Kenjiro Taura</i>	
Simulating the Spread of Infectious Disease over Large Realistic Social Networks Using Charm++	507
<i>Keith R. Bisset, Ashwin M. Aji, Eric Boehm, Laxmikant V. Kale, Tariq Kamal, Madhav V. Marathe, and Jae-Seung Yeom</i>	

Methods

A New Method of MHP Analysis for Languages with Dynamic Barriers	519
<i>Saurabh Joshi, R.K. Shyamasundar, and Sanjeev K. Aggarwal</i>	
Awareness of MPI Virtual Process Topologies on the Single-Chip Cloud Computer	529
<i>Steffen Christgau and Bettina Schnor</i>	
Speedup for Multi-Level Parallel Computing	537
<i>Shanjiang Tang, Bu-Sung Lee, and Bingsheng He</i>	

Tools, Compilers and Libraries

Conflict Avoidance Scheduling Using Grouping List for Transactional Memory	547
<i>Dongmin Choi, Seung Hun Kim, and Won W. Ro</i>	
Compile-Time Detection of False Sharing via Loop Cost Modeling	557
<i>Munara Tolubaeva, Yonghong Yan, and Barbara Chapman</i>	

Communication Library to Overlap Computation and Communication for OpenCL Application	567
<i>Toshiya Komoda, Shinobu Miwa, and Hiroshi Nakamura</i>	
HERCULES: A Pattern Driven Code Transformation System	574
<i>Christos Kartsaklis, Oscar Hernandez, Chung-Hsing Hsu, Thomas Ilsche, Wayne Joubert, and Richard L. Graham</i>	
Nature Inspired Distributed Computing - NIDISC	
NIDISC Introduction	584
<i>Pascal Bouvry, Franciszek Seredynski, and El-Ghazali Talbi</i>	
Session 1: Large-scale Distributed Systems	
When Distributed Hash Tables Meet Chemical Programming for Autonomic Computing	585
<i>Marko Obrovac and Cédric Tedeschi</i>	
An Efficient Stochastic Local Search for Heterogeneous Computing Scheduling	593
<i>Sergio Nesmachnow, Francisco Luna, and Enrique Alba</i>	
Using Simulated Annealing to Find Lower Bounds of Localization with Noisy Measurements	601
<i>Farhan Ahammed, Javid Taheri, and Albert Y. Zomaya</i>	
The Necessity for Strong Reciprocators in Mobile Ad Hoc Networks	609
<i>Marcin Seredynski, Grégoire Danoy, and Pascal Bouvry</i>	
Session 2: Hybrid Metaheuristics	
A Novel Approach for Regularized Signal Deconvolution Based on Hybrid Swarm Intelligence: Application to Neutron Radiography	617
<i>Slami Saadi, Abderrezak Guessoum, and Maamar Bettayeb</i>	
Parallel Hybrid Metaheuristic for Multi-objective Biclustering in Microarray Data	625
<i>Khedidja Seridi, Laetitia Jourdan, and El-Ghazali Talbi</i>	
Hybrid Differential Evolution Using Low-Discrepancy Sequences for Image Segmentation	634
<i>A. Nakib, B. Daachi, and P. Siarry</i>	
Session 3: Metaheuristics Parallelization and Application	
Parallel Simulated Annealing for Fragment Based Sequence Alignment	641
<i>Jan Mendonca Correa, Alba Cristina Magalhaes Alves de Melo, Ricardo P. Jacobi, and Azzedine Boukerche</i>	
Discovering Cellular Automata Rules for Binary Classification Problem with Use of Genetic Algorithm	649
<i>Anna Piwonska, Franciszek Seredynski, and Mirosław Szaban</i>	

Enhanced Parallel Cooperative Model for Trajectory Based Metaheuristics: A Scalability Analysis	656
<i>Gabriel Luque, Francisco Luna, and Enrique Alba</i>	
High Performance Computational Biology - HiCOMB	
HiCOMB Introduction	663
<i>Mark J. Clement, Quinn Snell, Srinivas Aluru, and David A. Bader</i>	
Next-Generation Data Analysis	
Investigating Memory Optimization of Hash-index for Next Generation Sequencing on Multi-core Architecture	665
<i>Wendi Wang, Wen Tang, Linchuan Li, Guangming Tan, Peiheng Zhang, and Ninghui Sun</i>	
Parallel Pair-HMM SNP Detection	675
<i>Nathan L. Clement, Brent A. Shepherd, Paul Bodily, Sukhbat Tumor-Ochir, Younghoon Gim, Quinn Snell, Mark J. Clement, and W. Evan Johnson</i>	
Evaluation of GPU-based Seed Generation for Computational Genomics Using Burrows-Wheeler Transform	684
<i>Yongchao Liu and Bertil Schmidt</i>	
Analysis of Evolution	
The Multi-Processor Scheduling Problem in Phylogenetics	691
<i>Jiajie Zhang and Alexandros Stamatakis</i>	
Optimizing the Execution of Statistical Simulations for Human Evolution in Hyper-threaded Multicore Architectures	699
<i>Raquel Dias, César A.F. De Rose, Antônio Tadeu Azevedo Gomes, and Nelson J.R. Fagundes</i>	
SlimCodeML: An Optimized Version of CodeML for the Branch-Site Model	706
<i>Hannes Schabauer, Mario Valle, Christoph Pacher, Heinz Stockinger, Alexandros Stamatakis, Marc Robinson-Rechavi, Ziheng Yang, and Nicolas Salamin</i>	
Sequence and Structure	
Finding Common RNA Secondary Structures: A Case Study on the Dynamic Parallelization of a Data-driven Recurrence	715
<i>Steven T. Stewart, Eric Aubanel, and Patricia A. Evans</i>	
The Development of Parallel Adaptive Sampling Algorithms for Analyzing Biological Networks	725
<i>Kathryn Dempsey, Kanimathi Duraisamy, Sanjukta Bhowmick, and Hesham Ali</i>	

GPU Approaches

A Speculative HMMER Search Implementation on GPU	735
<i>Xiaoqiang Li, Wenting Han, Gu Liu, Hong An, Mu Xu, Wei Zhou, and Qi Li</i>	
Probabilistic Brain Fiber Tractography on GPUs	742
<i>Mo Xu, Xiaorui Zhang, Yu Wang, Ling Ren, Ziyu Wen, Yi Xu, Gaolang Gong, Ningyi Xu, and Huazhong Yang</i>	
Quantitative Trait Locus Analysis Using a Partitioned Linear Model on a GPU Cluster	752
<i>Peter E. Bailey, Tapasya Patki, Gregory M. Striemer, Ali Akoglu, David K. Lowenthal, Peter Bradbury, Matt Vaughn, Liya Wang, and Stephen Goff</i>	

Advances in Parallel and Distributed Computing Models - APDCM

APDCM Introduction	761
<i>Oscar H. Ibarra</i>	

Programming Models

Dataflow-like Synchronization in a PGAS Programming Model	762
<i>Jens Breitbart</i>	
On the Feasibility of a Distributed Runtime for the Chemical Programming Model	770
<i>Marko Obrovac and Cédric Tedeschi</i>	
Multi-core Portability Abstraction	778
<i>Martti Forsell and Mikko Hiivala</i>	
An Extended PRAM-NUMA Model of Computation for TCF Programming	786
<i>Martti Forsell and Ville Leppänen</i>	

Parallel Architectures and GPUs

Simple Memory Machine Models for GPUs	794
<i>Koji Nakano</i>	
Counting Problems on Graphs: GPU Storage and Parallel Computing Techniques	804
<i>Amlan Chatterjee, Sridhar Radhakrishnan, and John K. Antonio</i>	
On Running Windowed Image Computations on a Pipeline	813
<i>Ramachandran Vaidyanathan and Phaneendra Vinukonda</i>	
Very Large-Scale Integrated Processor	821
<i>Shigeyuki Takano</i>	

Wireless Networks

Optimal Number of Annuli for Maximizing the Lifetime of Sensor Networks	829
<i>Keqin Li</i>	
Distributed Algorithms for TDMA Link Scheduling in Sensor Networks	839
<i>Thamer Alsulaiman, Sushil K. Prasad, and Alexander Zelikovsky</i>	

An Energy-Optimum and Communication-Time Efficient Protocol for Allocation, Scheduling and Routing in Wireless Networks	848
<i>Thiago F. Neves, Marcos F. Caetano, and Jacir L. Bordim</i>	
An Opportunistic MAC Protocol Based on Statistical Spectrum Analysis	855
<i>Felipe M. Modesto, Marcos F. Caetano, and Jacir L. Bordim</i>	
Parallel and Distributed Algorithms	
A Self-stabilizing Algorithm for the Maximal 2-packing in a Cactus Graph	863
<i>Joel Antonio Trejo-Sánchez and José Alberto Fernández-Zepeda</i>	
An Extension of Matthews' Bound to Multiplex Random Walks	872
<i>Yusuke Hosaka, Yukiko Yamauchi, Shuji Kijima, Hirotaka Ono, and Masafumi Yamashita</i>	
Tight Approximation for Scheduling Parallel Jobs on Identical Clusters	878
<i>Marin Bougeret, Pierre-Francois Dutot, Klaus Jansen, Christina Robenek, and Denis Trystram</i>	
On-line Batch Scheduling in Distributed Optical Networks	886
<i>Yang Wang, Xiaojun Cao, Adrian Caciula, and Qian Hu</i>	
Communication Architecture for Scalable Systems - CASS	
CASS Introduction	894
<i>José Flich, Scott Pakin, and Craig Stunkel</i>	
Session I: Messaging Layers	
On the Portability and Performance of Message-Passing Programs on Embedded Multicore Platforms	896
<i>Shih-Hao Hung, Po-Hsun Chiu, Chia-Heng Tu, Wei-Ting Chou, and Wen-Long Yang</i>	
Optimized Reduce for Mesh-Based NoC Multiprocessors	904
<i>Adán Kohler and Martin Radetzki</i>	
Estimating Application Hierarchical Bandwidth Requirements Using BSP Family Models	914
<i>Adrian Soviani and Jaswinder Pal Singh</i>	
Session II: Routing and Data Transfer	
Design of Direct Communication Facility for Many-Core Based Accelerators	924
<i>Min Si and Yutaka Ishikawa</i>	
Achieving Global Fairness for On-Chip Network Using Group Allocation	930
<i>Shan-Jung Miao, Yin Men, and Yarsun Hsu</i>	
Limited Multi-path Routing on Extended Generalized Fat-trees	938
<i>Santosh Mahapatra, Xin Yuan, and Wickus Nienaber</i>	

High-Performance, Power-Aware Computing - HPPAC

HPPAC Introduction	946
<i>Bronis R. de Supinski and Roberto Gioiosa</i>	

Session 1: Capping and Scheduling

Beyond DVFS: A First Look at Performance under a Hardware-Enforced Power Bound	947
<i>Barry Rountree, Dong H. Ahn, Bronis R. de Supinski, David K. Lowenthal, and Martin Schulz</i>	
A Power Provision and Capping Architecture for Large Scale Systems	954
<i>Yongpeng Liu, Hong Zhu, Kai Lu, and Yongyan Liu</i>	
Dynamic Thread Scheduling in Asymmetric Multicores to Maximize Performance-per-Watt	964
<i>Arunachalam Annamalai, Rance Rodrigues, Israel Koren, and Sandip Kundu</i>	

Session 2: Power Efficient Hardware

Energy-Efficient and Fault-Tolerant Unified Buffer and Bufferless Crossbar Architecture for NoCs	972
<i>Yixuan Zhang, Randy Morris, Dominic DiTomaso, and Avinash Kodi</i>	
Optimizing Data Allocation and Memory Configuration for Non-Volatile Memory Based Hybrid SPM on Embedded CMPs	982
<i>Jingtong Hu, Qingfeng Zhuge, Chun Jason Xue, Wei-Che Tseng, and Edwin H.M. Sha</i>	

Session 3: Energy Workloads

Modeling Power and Energy Usage of HPC Kernels	990
<i>Ananta Tiwari, Michael A. Laurenzano, Laura Carrington, and Allan Snaveley</i>	
Power-Efficient Schemes via Workload Characterization on the Intel's Single-Chip Cloud Computer	999
<i>Gustavo A. Chaparro-Baquero, Qi Zhou, Chen Liu, Jie Tang, and Shaoshan Liu</i>	

Session 4: Power and Energy Profiling and Metrics

The Green Index: A Metric for Evaluating System-Wide Energy Efficiency in HPC Systems	1007
<i>Balaji Subramaniam and Wu-chun Feng</i>	
Energy Efficiency Analysis of GPUs	1014
<i>Juan M. Cebrián, Ginés D. Guerrero, and José M. García</i>	
Tracing and Visualization of Energy-Related Metrics	1023
<i>Timo Minartz, Julian Kunkel, and Thomas Ludwig</i>	

High-Performance Grid and Cloud Computing - HPGC

HPGC Introduction	1031
<i>Eric Aubanel, Virendra Bhavsar, and Michael Frumkin</i>	

Session 1: Applications/Performance Analysis

Improving Parallelisation of a Monte Carlo Radiotherapy Simulation Using MPI	1033
<i>Gagarine Yaikhom, David W. Walker, and Coral Walker</i>	
A Simulation Study on Urban Water Threat Detection in Modern Cyberinfrastructures	1040
<i>Lizhe Wang, Dan Chen, Ze Deng, and Rajiv Ranjan</i>	
AzureBench: Benchmarking the Storage Services of the Azure Cloud Platform	1048
<i>Dinesh Agarwal and Sushil K. Prasad</i>	
Analysis and Optimization of Data Import with Hadoop	1058
<i>Weijia Xu, Wei Luo, and Nicholas Woodward</i>	

Session 2: Scheduling/Provisioning

Online Scheduling for Cloud Computing and Different Service Levels	1067
<i>Uwe Schwiegelshohn and Andrei Tchernykh</i>	
Cooperative Game Theoretical Techniques for Energy-Aware Task Scheduling in Cloud Computing	1075
<i>Nickolas Bielik and Ishfaq Ahmad</i>	
Provisioning Policies for Elastic Computing Environments	1085
<i>Paul Marshall, Henry Tufo, and Kate Keahey</i>	

Session 3: Grid/Cloud Infrastructure

Middleware Support for RDMA-based Data Transfer in Cloud Computing	1095
<i>Yufei Ren, Tan Li, Dantong Yu, Shudong Jin, and Thomas Robertazzi</i>	
Instant GridFTP	1104
<i>Rajkumar Kettimuthu, Lukasz Lacinski, Mike Link, Karl Pickett, Steve Tuecke, and Ian Foster</i>	
Using Static Code Analysis to Improve Performance of GridRPC Applications	1113
<i>Oleg Girko and Alexey Lastovetsky</i>	
Distributed Virtual Diskless Checkpointing: A Highly Fault Tolerant Scheme for Virtualized Clusters	1120
<i>Ben Eckart, Xubin He, Chentao Wu, Ferrol Aderholdt, Fang Han, and Stephen Scott</i>	
Different Approaches to Distributed Compilation	1128
<i>Josef Gattermayer and Pavel Tvrdik</i>	

System Management Techniques, Processes, and Services - SMTPS

SMTPS Introduction	1135
<i>Liana L. Fong, Renato J. Figueiredo, and Kyung Dong Ryu</i>	
Monitoring and Predicting Hardware Failures in HPC Clusters with FTB-IPMI	1136
<i>Raghunath Rajachandrasekar, Xavier Besseron, and Dhabaleswar K. Panda</i>	
VM Performance Isolation to Support QoS in Cloud	1144
<i>Marcio Silva, Kyung Dong Ryu, and Dilma Da Silva</i>	
Eucalyptus: Support for Effective Use of Persistent Memory	1152
<i>Mohammad Banikazemi and Bulent Abali</i>	
Designing Network Failover and Recovery in MPI for Multi-Rail InfiniBand Clusters	1160
<i>S. Pai Raikar, H. Subramoni, K. Kandalla, J. Vienne, and Dhabaleswar K. Panda</i>	
Scalla: Structured Cluster Architecture for Low Latency Access	1168
<i>Andrew Hanushevsky and Daniel L. Wang</i>	

Security and Trust of Distributed Networking Systems - STDN

STDN Introduction	1176
<i>Jinshu Su</i>	

Security

Online Anomaly Detection Based on Web Usage Mining	1177
<i>Yi Xie and Shensheng Tang</i>	
A Secure and Hierarchical Architecture for P2PSIP Session Initiation	1183
<i>Xianghan Zheng, Wenzhong Guo, Shangping Zhong, and Zhiyong Yu</i>	
Hardware/Software Mechanisms for Protecting an IDS against Algorithmic Complexity Attacks	1190
<i>Govind Sreekar Shenoy, Jordi Tubella, and Antonio González</i>	
LDC: Detecting BGP Prefix Hijacking by Load Distribution Change	1197
<i>Yujing Liu, Jinshu Su, and Rocky K.C. Chang</i>	
A Stream Reassembly Mechanism Based on DPI	1204
<i>Shuhui Chen and Yong Tang</i>	

Trust

A Practical Privacy-preserving Password Authentication Scheme for Cloud Computing	1210
<i>Ali A. Yassin, Hai Jin, Ayad Ibrahim, Weizhong Qiang, and Deqing Zou</i>	
Defeating against Sybil-attacks in Peer-to-peer Networks	1218
<i>Xu Xiang</i>	

An Efficient Property-Based Attestation Scheme with Flexible Revocation Mechanisms	1223
<i>Yue Xiao-han and Zhou Fucai</i>	
Relational Query Authentication Based on Hierarchical Hash Chain in Distributed Environments	1231
<i>Jian Xu, Fuxiang Li, Qiong Xiao, and Fucai Zhou</i>	
A 3N Approach to Network Control and Management	1237
<i>Feng Zhao, Dan Zhao, Xiaofeng Hu, Wei Peng, Baosheng Wang, and Zexin Lu</i>	
Management	
A New Task Allocation Algorithm Based on Dynamic Coalition in WSNs	1243
<i>Chengyu Chen, Wenzhong Guo, and Guolong Chen</i>	
Energy-balanced Sleep Scheduling Based on Particle Swarm Optimization in Wireless Sensor Network	1249
<i>Chaolong Yu, Wenzhong Guo, and Guolong Chen</i>	
Resiliency Controlling of Wireless Sensor Networks for the Protecting from Internal Attacks	1256
<i>Xu Huang and Dharmendra Sharma</i>	
Mining of Attack Models in IDS Alerts from Network Backbone by a Two-stage Clustering Method	1263
<i>Lin-Bo Qiao, Bo-Feng Zhang, Zhi-Quan Lai, and Jin-Shu Su</i>	
Characterizing Enough Vantage Points for Pinpointing Routing Instability	1270
<i>Hongjun Liu, Xiaofeng Hu, Dan Zhao, and Xicheng Lu</i>	
NSF/TCPP Workshop on Parallel and Distributed Computing Education - EduPar	
EduPar Introduction	1276
<i>Sushil K. Prasad</i>	
Curricular Change and Adoption	
Engineering a New Curriculum: Experiences at Ohio University in Incorporating the IEEE-TCPP Curriculum Initiative During a Transition to Semesters	1279
<i>David W. Juedes and Frank Drews</i>	
Experiences in Teaching a Specialty Multicore Computing Course	1283
<i>Peter E. Strazdins</i>	
An Experience of Early Initiation to Parallelism in the Computing Engineering Degree at the University of Murcia, Spain	1289
<i>Manuel E. Acacio, Javier Cuenca, Lorenzo Fernández, Ricardo Fernández-Pascual, Joaquín Cervera, Domingo Giménez, M. Carmen Garrido, Juan A. Sánchez Laguna, José Guillén, Juan Alejandro Palomino Benito, and María-Eugenia Requena</i>	

Novel Ways of Teaching

Teaching by Example: Using Analogies and Live Coding Demonstrations to Teach Parallel Computing Concepts to Undergraduate Students	1295
<i>Nasser Giacaman</i>	
Distributed Systems with Wireless Sensor Networks	1299
<i>Noemi Rodriguez and Silvana Rossetto</i>	
The Spanish Parallel Programming Contests and its Use as an Educational Resource	1303
<i>Francisco Almeida, Javier Cuenca, Ricardo Fernández-Pascual, Domingo Giménez, and Juan Alejandro Palomino Benito</i>	
Making Learning Parallel Processing Interesting	1307
<i>Jie Liu, Yanwei Wu, and John Marsaglia</i>	
Simplifying Hands-On Teaching of Distributed Algorithms with SPLAY	1311
<i>Etienne Rivière</i>	

Models for Incorporating

Lessons Learned after the Introduction of Parallel and Distributed Computing Concepts into ECE Undergraduate Curricula at UTN-Bahía Blanca Argentina	1317
<i>Javier Iparraguirre, Guillermo R. Friedrich, and Ricardo J. Coppo</i>	
Formal Methods for Surviving the Jungle of Heterogeneous Parallelism	1321
<i>Ganesh Gopalakrishnan</i>	
Incorporating the NSF/TCPP Curriculum Recommendations in a Liberal Arts Setting	1325
<i>Akshaye Dhawan</i>	
CSinParallel and Synergy for Rapid Incremental Addition of PDC Into CS Curricula	1329
<i>Richard Brown and Elizabeth Shoop</i>	
Courses in High-performance Computing for Scientists and Engineers	1335
<i>Richard W. Vuduc, Kenneth Czechowski, Aparna Chandramowlishwaran, and Jee Whan Choi</i>	

Parallel and Distributed Scientific and Engineering Computing - PDSEC

PDSEC Introduction	1341
<i>Thomas Rauber, Gudula Rünger, Peter E. Strazdins, Laurence T. Yang, Guangming Tan, and Yi Pan</i>	

Session 1 : Parallel Programming Techniques

Towards the Scalability of Dynamic Loop Scheduling Techniques via Discrete Event Simulation	1343
<i>Mahadevan Balasubramaniam, Nitin Sukhija, Florina M. Ciorba, Ioana Banicescu, and Srishti Srivastava</i>	
LDPLFS: Improving I/O Performance without Application Modification	1352
<i>S.A. Wright, S.D. Hammond, S.J. Pennycook, I. Miller, J.A. Herdman, and S.A. Jarvis</i>	

Deploying Scalable and Secure Secret Sharing with GPU Many-Core Architecture	1360
<i>Su Chen, Ling Bai, Yi Chen, Hai Jiang, and Kuan-Ching Li</i>	

Session 2: Numerical Algorithms

Parallelizing the Computation of Green Functions for Computational Electromagnetism Problems	1370
<i>Carlos Pérez-Alcaraz, Domingo Giménez, Alejandro Álvarez-Melcón, and Fernando D. Quesada</i>	
Implementation and Evaluation of Triple Precision BLAS Subroutines on GPUs	1378
<i>Daichi Mukunoki and Daisuke Takahashi</i>	
Scalable Parallel Algorithms for Boundary Control of Thermally Convective Flows	1387
<i>Haijian Yang and Xiao-Chuan Cai</i>	
Evaluating Polynomials in Several Variables and their Derivatives on a GPU Computing Processor	1397
<i>Jan Verschelde and Genady Yoffe</i>	
Deriving a Methodology for Code Deployment on Multi-Core Platforms via Iterative Manual Optimizations	1406
<i>Stuart McCool, Peter Milligan, and Paul Sage</i>	

Session 3: Network and Cloud Computing

Scientific Application Performance on HPC, Private and Public Cloud Resources: A Case Study Using Climate, Cardiac Model Codes and the NPB Benchmark Suite	1416
<i>Peter E. Strazdins, Jie Cai, Muhammad Atif, and Joseph Antony</i>	
An Effective Self-adaptive Load Balancing Algorithm for Peer-to-Peer Networks	1425
<i>Naixue Xiong, Kaihua Xu, Lilong Chen, Laurence T. Yang, and Yuhua Liu</i>	
A System for GIS Polygonal Overlay Computation on Linux Cluster - An Experience and Performance Report	1433
<i>Dinesh Agarwal, Satish Puri, Xi He, and Sushil K. Prasad</i>	
Scheduling of Tasks in the Parareal Algorithm for Heterogeneous Cloud Platforms	1440
<i>Hongtao Xiao and Eric Aubanel</i>	

Session 4: Science and Engineering Applications

Numerical Design of an Optimal Bypass for a Partially Blocked Artery	1449
<i>Rongliang Chen and Xiao-Chuan Cai</i>	
A Utility Based Power-Aware Autonomic Approach for Running Scientific Applications	1457
<i>Rajat Mehrotra, Ioana Banicescu, and Srishti Srivastava</i>	
Distributing Power Grid State Estimation on HPC Clusters - A System Architecture Prototype	1467
<i>Yan Liu, Wei Jiang, Shuangshuang Jin, Mark Rice, and Yousu Chen</i>	

A Parallel Resampling Algorithm for Particle Filtering on Shared-Memory Architectures	1477
<i>Peng Gong, Yuksel Ozan Basciftci, and Fusun Ozguner</i>	
A Unified Study of Epidemic Routing Protocols and their Enhancements	1484
<i>Zhenxin Feng and Kwan-Wu Chin</i>	
Dependable Parallel, Distributed and Network-Centric Systems - DPDNS	
DPDNS Introduction	1494
<i>Dimitar Avresky and Erik Maehle</i>	
Session 1: Dependable Wireless and Sensor Networks	
Using Localization for Fault-tolerant Radio Coverage in Wireless Mesh Networks	1496
<i>Svilen Ivanov and Edgar Nett</i>	
Joint and Simultaneous K-sensing Detection in Deterministic and Random Sensor Networks	1506
<i>Yun Wang and Andrew Kutta</i>	
A Fault-aware Sensor Architecture for Cooperative Mobile Applications	1512
<i>Jörg Kaiser and Sebastian Zug</i>	
Network Coding-Based On-Demand Multipath Routing in MANET	1520
<i>Baolin Sun, Xiaocheng Lu, Chao Gui, Ying Song, and Hua Chen</i>	
Session 2: Architectures and Cloud	
Automated Workload Characterization in Cloud-based Transactional Data Grids	1525
<i>Bruno Ciciani, Diego Didona, Pierangelo Di Sanzo, Roberto Palmieri, Sebastiano Peluso, Francesco Quaglia, and Paolo Romano</i>	
A Usage Control Based Architecture for Cloud Environments	1534
<i>Tina Tavizi, Mehdi Shajari, and Peyman Dodangeh</i>	
Near-Zero-Downtime Single-Machine Hypervisor Rejuvenation	1540
<i>Dmitry Zakharov and Felix Salfner</i>	
Session 3: Dependability Analysis and Verification	
Analysis of Data Reliability Tradeoffs in Hybrid Distributed Storage Systems	1546
<i>Bing Tang and Gilles Fedak</i>	
Component Substitutability Behavior Consistency Verification Based on Architecture-driven Development	1556
<i>Luxi Chen, Linpeng Huang, and Chen Li</i>	
Adaptive Mapping of Parallelized Application (Fork-join DAG) on Multicore System in the Presence of Multiple Failures	1563
<i>Gilles Bizot, Dimitar Avresky, Fabien Chaix, Nacer-Eddine Zergainoh, and Michael Nicolaidis</i>	

Multi-Threaded Architectures and Applications - MTAAP

MTAAP Introduction	1571
<i>Luiz DeRose</i>	

Runtime and Scheduling

Resilience to Various Failures for Read-mostly In-memory Data Structures	1572
<i>Larry Kaplan, Preston Briggs, Miles Ohlrich, and Will Leslie</i>	

Scheduling OR-parallelism in YapOr and ThOr on Multi-Core Machines	1581
<i>Inês Dutra, Ricardo Rocha, Vitor Santos Costa, Fernando Silva, and João Santos</i>	

A Discussion in Favor of Dynamic Scheduling for Regular Applications in Many-core Architectures	1591
<i>Elkin Garcia, Daniel Orozco, Robert Pavel, and Guang R. Gao</i>	

Algorithms and Applications

An Approach to Parallelize Kruskal's Algorithm Using Helper Threads	1601
<i>Anastasios Katsigiannis, Nikos Anastopoulos, Konstantinos Nikas, and Nectarios Koziris</i>	

Merge Path - Parallel Merging Made Simple	1611
<i>Saher Odeh, Oded Green, Zahi Mwassi, Oz Shmueli, and Yitzhak Birk</i>	

Scalable Multi-threaded Community Detection in Social Networks	1619
<i>Jason Riedy, David A. Bader, and Henning Meyerhenke</i>	

Architecture

An Early Evaluation of the Scalability of Graph Algorithms on the Intel MIC Architecture	1629
<i>Erik Saule and Ümit V. Çatalyürek</i>	

PMU-guided Priority Adjustment to Guarantee Thread Performance on IBM POWER SMT Processor	1640
<i>Zhengyu He and Bo Hong</i>	

Architecture Design Trade-offs among VLIW SIMD and Multi-core Schemes	1649
<i>Yaohua Wang, Shuming Chen, Kai Zhang, Hu Chen, and Xiaowen Chen</i>	

Large-Scale Parallel Processing - LSPP

LSPP Introduction	1659
<i>Darren J. Kerbyson, Ram Rajamony, and Charles Weems</i>	

Session 1: Algorithms

Dynamic Load Balancing for Unstructured Meshes on Space-Filling Curves	1661
<i>Daniel F. Harlacher, Harald Klimach, Sabine Roller, Christian Siebert, and Felix Wolf</i>	

Mesh Interface Resolution and Ghost Exchange in a Parallel Mesh Representation	1670
<i>Timothy J. Tautges, Jason A. Kraftcheck, Nathan Bertram, Vipin Sachdeva, and John Magerlein</i>	

Automatic Refinement of Parallel Applications Structure Detection	1680
<i>Juan Gonzalez, Kevin Huck, Judit Gimenez, and Jesus Labarta</i>	

Session 2: Specialized Systems

Scalable and Efficient Associative Processor Solution to Guarantee Real-Time Requirements for Air Traffic Control Systems	1688
<i>Mike Yuan, Johnnie W. Baker, Will Meilander, and Kevin Schaffer</i>	

Sparse Matrix-vector Multiplication on GPGPU Clusters: A New Storage Format and a Scalable Implementation	1696
<i>Moritz Kreutzer, Georg Hager, Gerhard Wellein, Holger Fehske, Achim Basermann, and Alan R. Bishop</i>	

An On-Demand Fast Parallel Pseudo Random Number Generator with Applications	1703
<i>Dip Sankar Banerjee, Aman Kumar Bahl, and Kishore Kothapalli</i>	

Session 3: Performance

High Volume Throughput Computing: Identifying and Characterizing Throughput Oriented Workloads in Data Centers	1712
<i>Jianfeng Zhan, Lixin Zhang, Ninghui Sun, Lei Wang, Zhen Jia, and Chunjie Luo</i>	

SWAPP: A Framework for Performance Projections of HPC Applications Using Benchmarks	1722
<i>Sameh Sharkawi, Don DeSota, Raj Panda, Stephen Stevens, Valerie Taylor, and Xingfu Wu</i>	

Reducing Migration-induced Cache Misses	1732
<i>Sajjid Reza and Gregory T. Byrd</i>	

Parallel Computing and Optimization - PCO

PCO Introduction	1742
<i>Didier El Baz</i>	

Session I: Combinatorial Scientific Computing

Scalable Hybrid Implementation of Graph Coloring Using MPI and OpenMP	1744
<i>Ahmet Erdem Saryüce, Erik Saule, and Ümit V. Çatalyürek</i>	

Two Edge Coloring Algorithms Using a Simple Matching Discovery Automata	1754
<i>J. Paul Daigle and Sushil K. Prasad</i>	

Session II: Parallel Optimization Algorithms

A Parallel BP Algorithm for the Discretizable Distance Geometry Problem	1762
<i>W. Gramacho, A. Mucherino, C. Lavor, and N. Maculan</i>	

GPU Implementation of the Branch and Bound Method for Knapsack Problems	1769
<i>Mohamed Esseghir Lalami and Didier El-Baz</i>	

Session III: Parallel Metaheuristics

Towards the Design of Systolic Genetic Search	1778
<i>Martín Pedemonte, Enrique Alba, and Francisco Luna</i>	
A Parallel Simulated Annealing Approach for the Mapping of Large Process Networks	1787
<i>François Galea and Renaud Sirdey</i>	
Parallel Local Search for the Costas Array Problem	1793
<i>Daniel Diaz, Florian Richoux, Yves Caniou, Philippe Codognet, and Salvador Abreu</i>	

Session IV: Issues in Optimization of Parallel or Distributed Systems

Optimal Partitioning of a Multicore Server Processor	1803
<i>Keqin Li</i>	
Reducing Cache Pollution of Threaded Prefetching by Controlling Prefetch Distance	1812
<i>Yan Huang, Zhi-min Gu, Jie Tang, Min Cai, Jianxun Zhang, and Ninghan Zheng</i>	
A Class of an Almost-Optimal Size-Independent Parallel Prefix Circuits	1820
<i>Hatem M. El-Boghdadi</i>	

Accelerators and Hybrid Exascale Systems - ASHES

ASHES Introduction	1827
<i>Pavan Balaji</i>	

Session 1: Modeling and Optimization

Modeling and Predicting Performance of High Performance Computing Applications on Hardware Accelerators	1828
<i>Mitesh R. Meswani, Laura Carrington, Didem Unat, Allan Snively, Scott Baden, and Stephen Poole</i>	
Efficient Intranode Communication in GPU-Accelerated Systems	1838
<i>Feng Ji, Ashwin M. Aji, James Dinan, Darius Buntinas, Pavan Balaji, Wu-chun Feng, and Xiaosong Ma</i>	
Optimizing MPI Communication on Multi-GPU Systems Using CUDA Inter-Process Communication	1848
<i>S. Potluri, H. Wang, D. Bureddy, A.K. Singh, C. Rosales, and Dhabaleswar K. Panda</i>	

Session 2: Programming Models

Towards High-Level Programming of Multi-GPU Systems Using the SkelCL Library	1858
<i>Michel Steuwer, Philipp Kegel, and Sergei Gorlatch</i>	
Scaling Data-Intensive Applications on Heterogeneous Platforms with Accelerators	1866
<i>Ana Balevic and Bart Kienhuis</i>	

An Analysis of Multicore Specific Optimization in MPI Implementations	1874
<i>Pengqi Cheng and Yan Gu</i>	

Session 3: Accelerated Applications

Implementing High-performance Intensity Model with Blur Effect on GPUs for Large-scale Star Image Simulation	1879
<i>Chao Li, Yunquan Zhang, Changwen Zheng, and Xiaohui Hu</i>	

Parallelizing the Hamiltonian Computation in DQMC Simulations: Checkerboard Method for Sparse Matrix Exponentials on Multicore and GPU	1889
<i>Che-Rung Lee, Zhi-Hung Chen, and Quey-Liang Kao</i>	

Parallel Multi-Temporal Remote Sensing Image Change Detection on GPU	1898
<i>Huming Zhu, Yu Cao, Zhiqiang Zhou, and Maoguo Gong</i>	

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

ParLearning Introduction	1905
<i>Sutanay Choudhury, George Chin, and Yinglong Xia</i>	

Session 1

Accelerating the Training of HTK on GPU with CUDA	1907
<i>Zhihui Du, Xiangyu Li, and Ji Wu</i>	

Session 2

Dynamic Linear Solver Selection for Transient Simulations Using Machine Learning on Distributed Systems	1915
<i>Paul R. Eller, Jing-Ru C. Cheng, and Robert S. Maier</i>	

2D Partitioning Based Graph Search for the Graph500 Benchmark	1925
<i>Koji Ueno and Toyotaro Suzumura</i>	

OLAP Aggregation Based on Dimension-oriented Storage	1932
<i>Zhao Jing-hua, Song Ai-mei, and Song Ai-bo</i>	

Session 3

A GPU-accelerated Approximate Algorithm for Incremental Learning of Gaussian Mixture Model	1937
<i>Chunlei Chen, Dejun Mu, Huixiang Zhang, and Bo Hong</i>	

Session 4

Task Parallel Implementation of Belief Propagation in Factor Graphs	1944
<i>Nam Ma, Yinglong Xia, and Viktor K. Prasanna</i>	

PQH: A Multithreaded Parallel NN Search Index for Content-based Image Retrieval	1954
<i>Hui-zhong Chen, Ning Jing, Yong-guang Chen, and Luo Chen</i>	

High Performance Data Intensive Computing - HPDIC

HPDIC Introduction	1963
<i>Song Wu and Yong-Jian Ren</i>	

Session 1: GPU/GPGPU Applications for Data Intensive Computing

A Polyhedral Modeling Based Source-to-Source Code Optimization Framework for GPGPU	1964
<i>Chenxi Wang, Kang Kang, Maohua Zhu, and Yangdong Deng</i>	

A Massively Parallel Approach for Nonlinear Interdependency Analysis of Multivariate Signals with GPGPU	1971
<i>Dan Chen, Lizhe Wang, Dong Cui, Dongchuan Lu, Xiaoli Li, Samee U. Khan, and Joanna Kolodziej</i>	

Forecasting High Dimensional Volatility Using Conditional Restricted Boltzmann Machine on GPU	1979
<i>Xianggao Cai and Xiaola Lin</i>	

Task Scheduling for GPU Accelerated Hybrid OLAP Systems with Multi-core Support and Text-to-Integer Translation	1987
<i>Maria Malik, Lubomir Riha, Colin Shea, and Tarek El-Ghazawi</i>	

Session 2: MapReduce and Hadoop

A Large-Scale Graph Learning Framework of Technological Gatekeepers by MapReduce	1997
<i>Liu Tong and Guo Wensheng</i>	

MapReduce across Distributed Clusters for Data-intensive Applications	2004
<i>Lizhe Wang, Jie Tao, Holger Marten, Achim Streit, Samee U. Khan, Joanna Kolodziej, and Dan Chen</i>	

MTSD: A Task Scheduling Algorithm for MapReduce Base on Deadline Constraints	2012
<i>Zhuo Tang, Junqing Zhou, Kenli Li, and Ruixuan Li</i>	

A Multi-source Message Passing Model to Improve the Parallelism Efficiency of Graph Mining on MapReduce	2019
<i>Zeng Feng Zeng, Bin Wu, and Tian Tian Zhang</i>	

A Highly Efficient Consolidated Platform for Stream Computing and Hadoop	2026
<i>Hiroya Matsuura, Masaru Ganse, and Toyotaro Suzumura</i>	

MapReduce Based Skyline Services Selection for QoS-aware Composition	2035
<i>Liang Chen, Li Kuang, and Jian Wu</i>	

Statistics-based Workload Modeling for MapReduce	2043
<i>Hailong Yang, Zhongzhi Luan, Wenjun Li, Depei Qian, and Gang Guan</i>	

A MapReduce-based Algorithm for Motif Search	2052
<i>Hongwei Huo, Shuai Lin, Qiang Yu, Yipu Zhang, and Vojislav Stojkovic</i>	

Session 3: Algorithms

Xtorus: An Extended Torus Topology for On-Chip Massive Data Communication	2061
<i>Liu Yu-hang, Zhu Ming-fa, Wang Jue, Xiao Li-min, and Gong Tao</i>	
Network Resource Control for Data Intensive Applications in Heterogeneous Infrastructures	2069
<i>Zhiming Zhao, Cosmin Dumitru, Paola Grosso, and Cees de Laat</i>	
A Scheduling Strategy Supporting OpenMP Task on Heterogeneous Multicore	2077
<i>Qian Cao and Min Zuo</i>	
Towards Parallel Spatial Query Processing for Big Spatial Data	2085
<i>Yunqin Zhong, Jizhong Han, Tieying Zhang, Zhenhua Li, Jinyun Fang, and Guihai Chen</i>	
A Server-Level Adaptive Data Layout Strategy for Parallel File Systems	2095
<i>Huaiming Song, Hui Jin, Jun He, Xian-He Sun, and Rajeev Thakur</i>	
The Chunk-Locality Index: An Efficient Query Method for Climate Datasets	2104
<i>Cheng Chen, Xiaomeng Huang, Haohuan Fu, and Guangwen Yang</i>	

Session 4: Cloud, Grid, Virtualization, and Miscellaneous

A Resource Auction Based Allocation Mechanism in the Cloud Computing Environment	2111
<i>Xingwei Wang, Jijia Sun, Min Huang, Chuan Wu, and Xueyi Wang</i>	
A Fault-Tolerant Target-Tracking Strategy Based on Unreliable Sensing in Wireless Sensor Networks	2116
<i>Yi Xie, Guoming Tang, Daifei Wang, Weidong Xiao, Daquan Tang, and Jiuyang Tang</i>	
One Double-Reduct Approach to Get Key Rules and the Experiment in Prison Computer Information Security	2126
<i>Ly Hanfei</i>	
A Pareto Frontier for Optimizing Data Transfer and Job Execution in Grids	2130
<i>Javid Taheri and Albert Y. Zomaya</i>	
Placement Strategy of Virtual Machines Based on Workload Characteristics	2140
<i>Jian Wan, Fei Pan, and Congfeng Jiang</i>	
Fine-grained Access Control and Revocation for Sharing Data on Clouds	2146
<i>Shan-shan Tu, Shao-zhang Niu, Hui Li, Yun Xiao-ming, and Meng-jiao Li</i>	
Component Interface Testing in Virtual Experiment for Visualization of Material Property Data	2156
<i>Chuanzhi Liu, Chunping Ouyang, and Yongbin Liu</i>	
Optimize Block-Level Cloud Storage System with Load-Balance Strategy	2162
<i>Li Zhou, Yi-Cheng Wang, Ji-Lin Zhang, Jian Wan, and Yong-Jian Ren</i>	
Online Scheduling with Migration Cost	2168
<i>Shuangquan Yang</i>	

DSDC: A Domain Scientific Data Cloud Based on Virtual Dataspaces	2176
<i>Zhenyu Liu, Changjun Hu, Yang Li, and Jingyuan Hu</i>	
Non-Cooperative Gaming and Bidding Model Based Resource Allocation in Virtual Machine Environment	2183
<i>Jian Wan, Dechuan Deng, and Congfeng Jiang</i>	
Workflow Models, Systems, Services and Applications in the Cloud - CloudFlow	
CloudFlow Introduction	2189
<i>Yong Zhao, Cui Lin, and Shiyong Lu</i>	
Session 1: Workflow	
A Framework for Nonrepudiable and Scalable Cross-Enterprise Workflow Management Systems in the Cloud	2191
<i>Gwan-Hwan Hwang, Yu-Cheng Hsiao, Yi-Chan Kao, and Heng-Yi Lin</i>	
Extending the Assurance Point (AP) Approach to Process Recovery for Use with Flow Groups	2201
<i>Le Gao, Susan D. Urban, Zev Friedman, and Jonathan Rodriguez</i>	
GreenPipe: A Hadoop Based Workflow System on Energy-efficient Clouds	2211
<i>Yaokuan Mao, Wenjun Wu, Hui Zhang, and Liang Luo</i>	
Adaptive Data Refinement for Parallel Dynamic Programming Applications	2220
<i>Shanjiang Tang, Ce Yu, Bu-Sung Lee, Chao Sun, and Jizhou Sun</i>	
Session 2: Cloud Computing	
Reliable Migration Module in Trusted Cloud Based on Security Level - Design and Implementation	2230
<i>Ying Chen, Qingni Shen, Pengfei Sun, Yangwei Li, Zhong Chen, and Sihan Qing</i>	
Cloud Services Gateway: A Tool for Exposing Private Services to the Public Cloud with Fine-grained Control	2237
<i>Srinath Perera, Rajika Kumarasiri, Supun Kamburugamuva, Senaka Fernando, Sanjiva Weerawarana, and Paul Fremantle</i>	
Improving the QoS of Web Applications across Multiple Virtual Machines in Cloud Computing Environment	2247
<i>Weizhe Zhang, Hongli Zhang, Huixiang Chen, Qizhen Zhang, and Albert M.K. Cheng</i>	
Session 3: Application	
A QoS-Aware Service Selection Method for Cloud Service Composition	2254
<i>Huihui Bao and Wanchun Dou</i>	
MapReduce Skyline Query Processing with a New Angular Partitioning Approach	2262
<i>Liang Chen, Kai Hwang, and Jian Wu</i>	

Job Scheduling Strategies for Parallel Processing - JSSPP

JSSPP Workshop Introduction	2271
<i>Walfredo Cirne, Narayan Desai, Eitan Frachtenberg, and Uwe Schwiegelshohn</i>	

Large Scale Distributed Service-oriented Systems - LSDSS

LSDSS Introduction	2272
<i>Evangelos Kotsovinos, Jian Cao, and Jinjun Chen</i>	

Session 1: Service Technology

Web Service Classification Based on Automatic Semantic Annotation and Ensemble Learning	2274
<i>Li Yuan-jie and Cao Jian</i>	
An Effective Dynamic Web Service Selection Strategy with Global Optimal QoS Based on Particle Swarm Optimization Algorithm	2280
<i>Guosheng Kang, Jianxun Liu, Mingdong Tang, and Yu Xu</i>	
A Novel Semantic Web Service Configuration Approach	2286
<i>Ming-rui Wang and Min Liu</i>	
A Business-driven Methodology for Service-Oriented Information System Development	2292
<i>Hongming Cai, Fenglin Bu, and Lihong Jiang</i>	

Session 2: Large Scale Service System

A New Data Reduction Approach over the Stream Processor Architecture	2300
<i>Qingkui Chen, Li Xiao, and Songlin Zhuang</i>	
A Preventing Fraud Trust Model in P2P Networks	2305
<i>Siming Liu, Yang Yu, Jiaying Xu, and Zhenguang Huang</i>	
Research on Context-aware Scheduling Algorithm Based on Correlation in Smart Home Environment	2312
<i>Wang Jingxiao, Zheng Hong, and Sun Nigang</i>	
Cloud Program with a Pricing Strategy for IaaS in Cloud Computing	2316
<i>Xing Wu, Ming Chao Wang, Wu Zhang, and Yike Guo</i>	

Session 3: Service Process and Application

Business Process Oriented Platform-as-a-Service Framework for Process Instances Intensive Applications	2320
<i>Yongqing Zheng, Jinshan Pang, Jian Li, and Lizhen Cui</i>	
Dependency-based Risk Evaluation for Robust Workflow Scheduling	2328
<i>Mingzhong Wang, Kotagiri Ramamohanarao, and Jinjun Chen</i>	

A Petri-Net Based Context-Aware Workflow System for Smart Home	2336
<i>Zhang Xing, Zheng Hong, and Liu Yulong</i>	
Research on the Pattern of Tourism E-commerce	2343
<i>Xiangxiang Xie, Tao Hu, and Baohong Li</i>	
Multicore and GPU Programming Models, Languages and Compilers - PLC	
PLC Introduction	2347
<i>Weimin Zheng and Jesse Fang</i>	
Session I: Compilers for Multicore-SIMD Processors and GPUs	
Compiling C/C++ SIMD Extensions for Function and Loop Vectorizaion on Multicore-SIMD Processors	2349
<i>Xinmin Tian, Hideki Saito, Milind Girkar, Serguei V. Preis, Sergey S. Kozhukhov, Alekssei G. Cherkasov, Clark Nelson, Nikolay Panchenko, and Robert Geva</i>	
Automatic Offloading C++ Expression Templates to CUDA Enabled GPUs	2359
<i>Jie Chen, Balint Joo, William Watson III, and Robert Edwards</i>	
Enabling Mixed OpenMP/MPI Programming on Hybrid CPU/GPU Computing Architecture	2369
<i>Tyng-Yeu Liang, Hung-Fu Li, and Jun-Yao Chiu</i>	
Session II: Runtime Implementation and Performance Predication	
A Highly Efficient Implementation of I/O Functions on GPU	2378
<i>Wei Wu, Feng Bin Qi, Wang Quan He, and Shan Shan Wang</i>	
Performance Estimation of GPUs with Cache	2384
<i>Arun Kumar Parakh, M. Balakrishnan, and Kolin Paul</i>	
Implementation of XcalableMP Device Acceleration Extention with OpenCL	2394
<i>Takuma Nomizu, Daisuke Takahashi, Jinpil Lee, Taisuke Boku, and Mitsuhsisa Sato</i>	
Dynamic Scheduling for Work Agglomeration on Heterogeneous Clusters	2404
<i>Jonathan Lifflander, G. Carl Evans, Anshu Arya, and Laxmikant V. Kale</i>	
Session III: Programming for MultiCore Processors and GPUs	
Parallel Algorithms for Approximate String Matching with k Mismatches on CUDA	2414
<i>Yu Liu, Longjiang Guo, Jinbao Li, Meirui Ren, and Keqin Li</i>	
Performance Study of SIMD Programming Models on Intel Multicore Processors	2423
<i>Peter Kristof, Hongtao Yu, Zhiyuan Li, and Xinmin Tian</i>	
Optimizing Data Warehousing Applications for GPUs Using Kernel Fusion/Fission	2433
<i>Haicheng Wu, Gregory Diamos, Jin Wang, Srihari Cadambi, Sudhakar Yalamanchili, and Srimat Chakradhar</i>	
Power-aware Programming with GPU Accelerators	2443
<i>Changyou Zhang, Kun Huang, Xiang Cui, and Yifeng Chen</i>	

Parameterized Verification of GPU Kernel Programs	2450
<i>Guodong Li and Ganesh Gopalakrishnan</i>	
PhD Forum	
PhD Forum Introduction	2460
<i>Luc Bougé and Bo Hong</i>	
Communication-Optimal Parallel N-body Solvers	2462
<i>Aparna Chandramowlishwaran and Richard W. Vuduc</i>	
Modeling and Analysis for Performance and Power	2466
<i>Jee Whan Choi and Richard W. Vuduc</i>	
Fixed-Priority Multiprocessor Scheduling: Critical Instant, Response Time and Utilization Bound	2470
<i>Nan Guan and Wang Yi</i>	
Privacy Preserving Techniques for Location Based Services in Mobile Networks	2474
<i>Xinxin Liu and Xiaolin Li</i>	
Energy-aware Scheduling: Models and Complexity Results	2478
<i>Guillaume Aupy</i>	
Fault Tolerance in P2P-Grid Environments	2482
<i>Wang Huan and Nakazato Hidenori</i>	
A Fast Repair Code Based on Regular Graphs for Distributed Storage Systems	2486
<i>Yan Wang and Xin Wang</i>	
Inference of Huge Trees under Maximum Likelihood	2490
<i>Fernando Izquierdo-Carrasco and Alexandros Stamatakis</i>	
Multithreaded Algorithms for Matching in Graphs with Application to Data Analysis in Flow Cytometry	2494
<i>Ariful Azad and Alex Pothen</i>	
Sequence Alignment on Massively Parallel Heterogeneous Systems	2498
<i>Aleksandr Drozd, Naoya Maruyama, and Satoshi Matsuoka</i>	
Subgraph Querying in Relational Networks: A MapReduce Approach	2502
<i>Zhao Zhao</i>	
MapReduce Framework Optimization via Performance Modeling	2506
<i>Lijie Xu</i>	
Integrated Parallelization of Computations and Visualization for Large-scale Applications	2510
<i>Preeti Malakar, Vijay Natarajan, and Sathish S. Vadhiyar</i>	
Identity Based Schemes for Securing Mobile Ad Hoc Networks	2514
<i>Uttam Ghosh</i>	
QoS-Oriented Data Dissemination in VANETs	2518
<i>Lifeng Zhang and Beihong Jin</i>	

Coverage-aware Geocast Routing in Urban Vehicular Networks	2522
<i>Ruobing Jiang and Yanmin Zhu</i>	
Towards Modelling Parallelism and Energy Performance of Multicore Systems	2526
<i>Bogdan Marius Tudor and Yong Meng Teo</i>	
Parallel Circuit Simulation on Multi/Many-core Systems	2530
<i>Xiaoming Chen, Yu Wang, and Huazhong Yang</i>	
On the Correctness of Mixing Lazy and Eager Version Management in Transactions	2534
<i>Lihang Zhao and Jeff Draper</i>	
Characterizing Load and Communication Imbalance in Large-Scale Parallel Applications	2538
<i>David Böhme, Felix Wolf, and Markus Geimer</i>	
Autonomous and Energy-Aware Management of Large-Scale Cloud Infrastructures	2542
<i>Eugen Feller and Christine Morin</i>	
Minimalistic Adaptive Resource Management for Multi-tier Applications Hosted on Clouds	2546
<i>Waheed Iqbal</i>	
Designing Flexible Resource Rental Models for Implementing HPC-as-a-Service in Cloud	2550
<i>Han Zhao and Xiaolin Li</i>	
Generalizing the Utility of GPUs in Large-Scale Heterogeneous Computing Systems	2554
<i>Shucui Xiao and Wu-chun Feng</i>	

Author Index