2015 IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW 2015)

Hyderabad, India
25-29 May 2015

Pages 1-636

IEEE Catalog Number: CFP1551J-POD
ISBN: 978-1-4673-7685-3

2015 IEEE International Parallel and Distributed Processing Symposium Workshops

IPDPSW 2015

Table of Contents

Message from the General Co-Chairs xxi

Message from the Workshops Chair xxiii

Workshop 1: HCW—Heterogeneity in Computing Workshop

HCW Introduction 1
Shoukat Ali and Denis Trystram

Message from the HCW Steering Committee Chair 3
Behrooz Shirazi

Message from the HCW Program Committee Chair 4
Denis Trystram

HCW 2014 Keynote Talk 5
Andrew Grimshaw

Session 1: Scheduling and Load Balancing

Considerations on Distributed Load Balancing for Fully Heterogeneous Machines: Two Particular Cases 6
Nathanaël Cheriere and Erik Saule

ProSteal: A Proactive Work Stealer for Bulk Synchronous Tasks Distributed on a Cluster of Heterogeneous Machines with Multiple Accelerators 17
Tarun Beri, Sorav Bansal, and Subodh Kumar

Scheduling Tasks with Precedence Constraints on Hybrid Multi-core Machines 27
Safia Kedad-Sidhoum, Florence Monna, and Denis Trystram

Session 2: Applications

Bridging the Gap between Performance and Bounds of Cholesky Factorization on Heterogeneous Platforms 34
Emmanuel Agullo, Olivier Beaumont, Lionel Eyraud-Dubois, Julien Herrmann, Suraj Kumar, Loris Marchal, and Samuel Thibault

Efficient Message Logging to Support Process Replicas in a Volunteer Computing Environment 46
Md Tarikul Islam, Hien Nguyen, Jaspal Subhlok, and Edgar Gabriel

Early Multi-node Performance Evaluation of a Knights Corner (KNC) Based NASA Supercomputer 57
Subhash Saini, Haoqiang Jin, Dennis Jespersen, Samson Cheung, Jahed Djomehri, Johnny Chang, and Robert Hood

Workshop 2: RAW—Reconfigurable Architectures Workshop

RAW Introduction and Committees 68
Jürgen Becker, Ken Eguro, Diana Göhringer, Wayne Luk, Marco D. Santambrogio, Ramachandran Vaidyanathan, and Steve Wilton

RAW 2015 Keynote 70
Viktor K. Prasanna

Session 1 — Runtime and Tools for Partially Reconfigurable FPGA-Based Systems

Mini-NOVA: A Lightweight ARM-based Virtualization Microkernel Supporting Dynamic Partial Reconfiguration 71
Tian Xia, Jean-Christophe Prevotet, and Fabienne Nouvel

Real-Time Multiprocessor Architecture for Sharing Stream Processing Accelerators 81
Berend H.J. Dekens, Marco J.G. Bekooij, and Gerard J.M. Smit

Partial Region and Bitstream Cost Models for Hardware Multitasking on Partially Reconfigurable FPGAs 90
Aurelio Morales-Villanueva and Ann Gordon-Ross

Relocation-Aware Floorplanning for Partially-Reconfigurable FPGA-Based Systems 97
Marco Rabozzi, Riccardo Cattaneo, Tobias Becker, Wayne Luk, and Marco D. Santambrogio

Session 2 — Applications and Special Purpose Architectures with Reconfigurable Hardware

High-Throughput Online Hash Table on FPGA 105
Da Tong, Shijie Zhou, and Viktor K. Prasanna

GraphMMU: Memory Management Unit for Sparse Graph Accelerators 113
Nachiket Kapre, Han Jianglei, Andrew Bean, Pradeep Moorthy, and Siddhartha

Adaptive Recursive Doubling Algorithm for Collective Communication 121
Omer Arap, Martin Swany, Geoffrey Brown, and Bryce Himebaugh

Accelerating Large-Scale Single-Source Shortest Path on FPGA 129
Shijie Zhou, Charalampos Chelmis, and Viktor K. Prasanna

Session 3 — New Architectures and Performance Evaluation for Reconfigurable Computing

Experiences with Compiler Support for Processors with Exposed Pipelines 137
Nicklas Bo Jensen, Pascal Schleuniger, Andreas Hindborg, Maxwell Walter, and Sven Karlsson

An Architecture for Configuring an Effcient Scan Path for a Subset of Elements 144
Arash Ashrafi and Ramachandran Vaidyanathan

Performance Modeling of Matrix Multiplication on 3D Memory Integrated FPGA 154
Shreyas G. Singapura, Anand Panangadan, and Viktor K. Prasanna

Enhancing Speedups for FPGA Accelerated SPICE through Frequency Scaling and Precision Reduction 163
Lim Hui Hui and Nachiket Kapre

Short Papers

An Automated High-Level Design Framework for Partially Reconfigurable FPGAs 170
Rohit Kumar and Ann Gordon-Ross

Intermediate-Level Synthesis of a Gauss-Jordan Elimination Linear Solver 176
Marc-Andre Daigneault and Jean Pierre David

K-Ways Partitioning of Polyhedral Process Networks: A Multi-level Approach 182
Riccardo Cattaneo, Mahdi Badie Moradmand, Donatella Sciuto, and Marco D. Santambrogio

Estimation of Non-functional Properties for Embedded Hardware with Application to Image Processing 190
Christian Herglotz, Jürgen Seiler, André Kaup, Arne Hendricks, Marc Reichenbach, and Dietmar Fey

Adaptive Reconfigurable Architecture for Image Denoising 196
Kartik V. Hegde, Vadiraj Kulkarni, R. Harshavardhan, and S. Sumam David

Workshop 3: HIPS—Workshop on High-Level Parallel Programming Models and Supportive Environments and LSPP—Workshop on Large-Scale Parallel Processing

HIPS-LSPP Introduction and Committees 202
Sriram Krishnamoorthy, Tobias Hilbrich, Darren J. Kerbyson, Ram Rajamony, and Charles Weems

HIPS-LSPP Keynotes 204
Torsten Hoefler and Laxmikant V. Kalé

Session I: Performance Analysis and Optimization

Folding Methods for Event Timelines in Performance Analysis 205
Matthias Weber, Ronald Geisler, Holger Brunst, and Wolfgang E. Nagel

Performance Analysis for Target Devices with the OpenMP Tools Interface 215
Tim Cramer, Robert Dietrich, Christian Terboven, Matthias S. Müller, and Wolfgang E. Nagel

High-Performance Coarray Fortran Support with MVAPICH2-X: Initial Experience and Evaluation 225
Jian Lin, Khaled Hamidouche, Xiaoyi Lu, Mingzhe Li, and Dhabaleswar K. Panda

On-demand Connection Management for OpenSHMEM and OpenSHMEM+MPI 235
Sourav Chakraborty, Hari Subramoni, Jonathan Perkins, Ammar A. Awan, and Dhabaleswar K. Panda

Session II: Parallelization

Speculative Runtime Parallelization of Loop Nests: Towards Greater Scope and Efficiency 245
Aravind Sukumaran Rajam, Luis Esteban Campostrini, Juan Manuel Martinez Caamaño, and Philippe Clauss

Session III: Application-Specific Studies

On the Impact of Execution Models: A Case Study in Computational Chemistry 255
Daniel Chavarría-Miranda, Mahantesh Halappanavar, Sriram Krishnamoorthy, Joseph Manzano, Abhinav Vishnu, and Adolfy Hoisie

Computing the Pseudo-Inverse of a Graph's Laplacian Using GPUs 265
Nishant Saurabh, Ana Lucia Varbanescu, and Gyan Ranjan

Workshop 4: NIDISC—Workshop on Nature Inspired Distributed Computing

NIDISC Introduction and Committees 275
Pascal Bouvry, Grégoire Danoy, Franciszek Seredyński, El-Ghazali Talbi, and Albert Y. Zomaya

Session 1: Applications of Bio-Inspired Algorithms

Dynamic Job Scheduling in the Cloud Using Slowdown Optimization and Sandpile Cellular Automata Model 276
Jakub Gąsior and Franciszek Seredyński

A Multi-objective Evolutionary Algorithm for Cloud Platform Reconfiguration 286
François Legillon, Nouredine Melab, Didier Renard, and El-Ghazali Talbi

A Genetic Algorithm Approach for Adjusting Time Series Based Load Prediction 292
Raed Alkharboush, Robson Eduardo De Grande, and Azzedine Boukerche

Session 2: Parallel, Distributed, and Adaptive Algorithms

Differential Evolution on a GPGPU: The Influence of Parameters on Speedup and the Quality of Solutions 299
Omar Andres Carmona Cortes, Mônica Sakuray Pais, Filipo Novo Mór, Andrew Rau-Chaplin, and César Augusto Missio Marcon

Distributed Cellular Evolutionary Algorithms in a Byzantine Environment 307
Jakub Muszyński, Sébastien Varrette, Bernabé Dorronsoro, and Pascal Bouvry

Bayesian Based Metaheuristic for Large Scale Continuous Optimization 314
Amir Nakib, Bernard Thibault, and Patrick Siarry

Firefly Inspired Improved Distributed Proximity Algorithm for D2D Communication 323
Ajay Pratap and Rajiv Misra

Workshop 5: HiCOMB—Workshop on High Performance Computational Biology

HiCOMB Introduction and Committees 329
Sanguthevar Rajasekaran, Srinivas Aluru, and David A. Bader

HiCOMB 2015 Keynote and Invited Talks 331
Ramesh Hariharan, Ananth Kalyanaraman, Michela Taufer, Trilce Estrada, Pietro Cicotti, and Pavan Balaji

HiCOMB Session 1

Perfect Hashing Structures for Parallel Similarity Searches 332
Tuan Tu Tran, Mathieu Giraud, and Jean-Stéphane Varré

A Crossbar Interconnection Network in DNA 342
Basavaraj Talawar

Handling Heterogeneity for Efficient Implementations: A Case Study on Sequence Comparison 346
Denis Trystram

Phylogenetic Analysis Using MapReduce Programming Model 350
G.M. Siddesh, K.G. Srinivasa, Ishank Mishra, Abhinav Anurag, and Eklavya Uppal

HiCOMB Session 2

Towards Context-Aware DNA Sequence Compression for Efficient Data Exchange 357
Wajeeta Lohana, Jawwad A. Shamsi, Tahir Q. Syed, and Farrukh Hasan

HiCOMB Session 3

Generalised Implementation for Fixed-Length Approximate String Matching under Hamming Distance and Applications 367
Solon Pissis and Ahmad Retha

Fine-Grained Acceleration of HMMER 3.0 via Architecture-Aware Optimization on Massively Parallel Processors 375
Hanyu Jiang and Narayan Ganesan

Workshop 6: APDCM—Advances in Parallel and Distributed Computing Models

APDCM Introduction and Committees 384
Oscar H. Ibarra, Koji Nakano, Akihiro Fujiwara, and Susumu Matsumae

Session 1: Parallel Algorithms and Applications

Bulk GCD Computation Using a GPU to Break Weak RSA Keys 385
Toru Fujita, Koji Nakano, and Yasuaki Ito

A Simple Parallel Algorithm for Biconnected Components in Sparse Graphs 395
Meher Chaitanya and Kishore Kothapalli

Parallelizing a Real-Time Audio Application—A Case Study in Multithreaded Software Engineering 405
Marc Aurel Kiefer, Korbinian Molitorisz, Jochen Bieler, and Walter F. Tichy

Performance Modeling of Multi-tiered Web Applications with Varying Service Demands 415
Ajay Kattepur and Manoj Nambiar

Session 2: Parallel Computing Systems

Efficient Estimation of Non-stationary Traffic Parameters on Networks-on-Chip 425
Abhishek Bansal, Sambhav Gupta, and Turbo Majumder

A Methodology for Co-Location Aware Application Performance Modeling in Multicore Computing 434
Daniel Dauwe, Eric Jonardi, Ryan Friese, Sudeep Pasricha, Anthony A. Maciejewski, David A. Bader, and Howard Jay Siegel

Performance Constrained Static Energy Reduction Using Way-Sharing Target-Banks 444
Shounak Chakraborty, Shirshendu Das, and Hemangee K. Kapoor

Decoupling Contention with Victim Row-Buffer on Multicore Memory Systems 454
Ke Gao, Dongrui Fan, Jie Wu, and Zhiyong Liu

Session 3: Distributed Algorithms and Computing

Replicated Data Placement for Uncertain Scheduling 464
Manmohan Chaubey and Erik Saule

Scheduling Computational Workflows on Failure-Prone Platforms 473
Guillaume Aupy, Anne Benoit, Henri Casanova, and Yves Robert

A Generic Framework for Impossibility Results in Time-Varying Graphs 483
Nicolas Braud-Santoni, Swan Dubois, Mohamed-Hamza Kaaouachi, and Franck Petit

Enabling Ring Exploration with Myopic Oblivious Robots 490
Ajoy K. Datta, Anissa Lamani, Lawrence L. Larmore, and Franck Petit

Session 4: Wireless Networks and Distributed Systems

Implementing Uniform Reliable Broadcast in Anonymous Distributed Systems with Fair Lossy Channels 500
Jian Tang, Mikel Larrea, Sergio Arévalo, and Ernesto Jiménez

Causal Consistency for Geo-Replicated Cloud Storage under Partial Replication 509
Min Shen, Ajay D. Kshemkalyani, and Ta-Yuan Hsu

Asterisk PBX Capacity Evaluation 519
Lucas Rodrigues Costa, Lucas Saad N. Nunes, Jacir Luiz Bordim, and Koji Nakano

A Fair Randomized Contention Resolution Protocol for Wireless Nodes without Collision Detection Capabilities 525
Marcos Fagundes Caetano and Jacir Luiz Bordim

Workshop 7: HPBC—High Performance Big Data and Cloud Computing Workshop and HPDIC—High Performance Data Intensive Computing

HPBC Introduction and Committees 534
Eric Aubanel, Virendrakumar C. Bhavsar, and Michael Frumkin

HPBC Keynote 535
Tim Mattson

HPDIC Introduction and Committees 536
Christophe Cérin, R.K. Shyamasundar, Yuqing Gao, and Cong-Feng Jiang

Session 1: Big Data and Cloud Computing: Storage, Analytics and Data Transfer

Cache Support in a High Performance Fault-Tolerant Distributed Storage System for Cloud and Big Data 537
Lars Lundberg, Håkan Grahn, Dragos Ilie, and Christian Melander

A Complex Event Processing Toolkit for Detecting Technical Chart Patterns 547
Madhushi Niluka Bandara, Rajitha Madhushan Ranasinghe, Rashmi Woranga Mudugamuwa Arachchi, Channa Gayan Somathilaka, Srinath Perera, and Daya Chinthana Wimalasuriya

High-Performance Serverless Data Transfer over Wide-Area Networks 557
Eun-Sung Jung and Rajkumar Kettimuthu

Session 2: High Performance Data Intensive Computing

Improving Performance of Structured-Memory, Data-Intensive Applications on Multi-core Platforms via a Space-Filling Curve Memory Layout 565
E. Wes Bethel, David Camp, David Donofrio, and Mark Howison

Query Execution for RDF Data Using Structure Indexed Vertical Partitioning 575
Bhavik Shah, Trupti Padiya, and Minal Bhise

Storm Pub-Sub: High Performance, Scalable Content Based Event Matching System Using Storm 585
Medha Abhijeet Shah and Dinesh B. Kulkarni

Workshop 8: ASHES—Accelerators and Hybrid Exascale Systems

AsHES Introduction and Committees 591
James Dinan, Wenguang Chen, Xiaosong Ma, Pavan Balaji, Satoshi Matsuoka, Jiayuan Meng, and Yunquan Zhang

AsHES Keynote 593
Michela Taufer

Session 1: Accelerating Analytics

Towards a Combined Grouping and Aggregation Algorithm for Fast Query Processing in Columnar Databases with GPUs 594
Sina Meraji, John Keenleyside, Sunil Kamath, and Bob Blainey

GraphReduce: Large-Scale Graph Analytics on Accelerator-Based HPC Systems 604
Dipanjan Sengupta, Kapil Agarwal, Shuaiwen Leon Song, and Karsten Schwan

Graph Coloring on the GPU and Some Techniques to Improve Load Imbalance 610
Shuai Che, Gregory Rodgers, Brad Beckmann, and Steve Reinhardt

Session 2: Algorithm Design for Heterogeneous Systems

GPU-based Parallel R-tree Construction and Querying 618
Sushil K. Prasad, Michael McDermott, Xi He, and Satish Puri

Fast Burrows Wheeler Compression Using All-Cores 628
Aditya Deshpande and P.J. Narayanan

A Novel Heterogeneous Algorithm for Multiplying Scale-Free Sparse Matrices 637
Kiran Raj Ramamoorthy, Dip Sankar Banerjee, Kannan Srinathan, and Kishore Kothapalli

Implementation of CG Method on GPU Cluster with Proprietary Interconnect TCA for GPU Direct Communication 647
Kazuya Matsumoto, Toshihiro Hanawa, Yuetsu Kodama, Hisafumi Fujii, and Taisuke Boku

Workshop 9: PLC—Programming Models, Languages, and Compilers for Manycore and Heterogeneous Architectures

PLC Introduction and Committees 656
Sunita Chandrasekaran

PLC Keynote 658
Michael Gschwind

Session I: Programming and Compilation Techniques for Heterogeneous and Multicore Systems

Implementing Cross-Device Atomics in Heterogeneous Processors 659
Meghana Gupta, Dibyendu Das, Prakash Raghavendra, Tony Tye, Leonid Lobachev, Amit Agarwal, and Ravish Hegde

A Novel Heterogeneous Framework for Local Dependency Dynamic Programming Problems 669
Rajesh Kumar and Kishore Kothapalli

OpenMP-MCA: Leveraging Multiprocessor Embedded Systems Using Industry Standards 679
Peng Sun, Sunita Chandrasekaran, and Barbara Chapman

Session II: Parallel Programming Experiences and Lessons Learned

Performance Portable Applications for Hardware Accelerators: Lessons Learned from SPEC ACCEL 689
Guido Juckeland, Alexander Grund, and Wolfgang E. Nagel

Understanding Performance Portability of OpenACC for Supercomputers 699
Suttinee Sawadsitang, James Lin, Simon See, Francois Bodin, and Satoshi Matsuoka

Session III: Novel Approaches for Emerging Platforms

Heterogeneous Habanero-C (H2C): A Portable Programming Model for Heterogeneous Processors 708
Deepak Majeti and Vivek Sarkar

Streamlining Whole Function Vectorization in C Using Higher Order Vector Semantics 718
Gil Rapaport, Ayal Zaks, and Yosi Ben-Asher

Workshop 10: EduPar—NSF/TCPP Workshop on Parallel and Distributed Computing Education

EduPar Introduction and Committees 728
Andrew Lumsdaine, Sushil K. Prasad, and Martina Barnas

EduPar Keynote 730
Geoffrey Charles Fox

Session 1: Methods and Tools

Teaching Simulations and High Performance Computing at Secondary Schools in the German State of Baden-Württemberg 731
Jörg Hilpert, Rüdiger Berlich, Peter Lürßen, Almut Zwölfer, and Jochen Barwind

The Active classroom: Students and Instructors Parallel Programming… in Parallel 739
Nasser Giacaman, Simar Kalra, and Oliver Sinnen

Introducing Tetra: An Educational Parallel Programming System 746
Ian Finlayson, Jerome Mueller, Shehan Rajapakse, and Daniel Easterling

Patternlets: A Teaching Tool for Introducing Students to Parallel Design Patterns 752
Joel C. Adams

Session 2: Course Design

A Research-Oriented Course on Advanced Multicore Architecture 760
Julio Sahuquillo, Salvador Petit, Vicent Selfa, and María Engracia Gómez

Updating an Introductory Performance Course with PDC Topics 766
Karen L. Karavanic and Daniel Leblanc

Novelties in Teaching High Performance Computing 772
Jawwad A. Shamsi, Nouman M. Durrani, and Nadeem Kafi

Session 3: Curriculum Integration

Incorporating PDC Modules Into Computer Science Courses at Jackson State University 779
Ali Abu El Humos, Sungbum Hong, Jacqueline Jackson, Xuejun Liang, Tzusheng Pei, and Bernard Aldrich

Integrating Parallel and Distributed Computing Topics into an Undergraduate CS Curriculum at UESTC 782
Guoming Lu, Jie Xu, Jieyan Liu, Bo Dai, Shenglin Gui, and Siyu Zhan

Fault-Tolerant Parallel and Distributed Computing for Software Engineering Undergraduates 788
Ali Ebnenasir and Jean Mayo

Workshop 11: GABB—Graph Algorithms Building Blocks

GABB Introduction and Committees 795
Tim Mattson

GABB Session 1

Declarative Patterns for Imperative Distributed Graph Algorithms 796
Marcin Zalewski, Nicholas Edmonds, and Andrew Lumsdaine

Parallel Triangle Counting and Enumeration Using Matrix Algebra 804
Ariful Azad, Aydin Buluç, and John Gilbert

GABB Session 2

Graph Laplacians and Least Squares on Graphs 812
Anil N. Hirani, Kaushik Kalyanaraman, and Seth Watts

Graphulo: Linear Algebra Graph Kernels for NoSQL Databases 822
Vijay Gadepally, Jake Bolewski, Dan Hook, Dylan Hutchison, Ben Miller, and Jeremy Kepner

A Unifying Programming Model for Parallel Graph Algorithms 831
Jeremiah Willcock and Andrew Lumsdaine

Fast Sparse Matrix and Sparse Vector Multiplication Algorithm on the GPU 841
Carl Yang, Yangzihao Wang, and John D. Owens

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

HPPAC Introduction and Committees 848
Wu-Chun Feng and Barry Rountree

Session 1: Provisioning and Management

Analyzing Energy-Time Tradeoff in Power Overprovisioned HPC Data Centers 849
Akhil Langer, Harshit Dokania, Laxmikant V. Kalé, and Udatta S. Palekar

Energy-Aware Server Provisioning by Introducing Middleware-Level Dynamic Green Scheduling 855
Daniel Balouek-Thomert, Eddy Caron, and Laurent Lefèvre

Adaptive Resource and Job Management for Limited Power Consumption 863
Yiannis Georgiou, David Glesser, and Denis Trystram

Session 2: Measurement, Modeling, and Optimization

On the Performance, Energy, and Power of Data-Access Methods in Heterogeneous Computing Systems 871
Rubasri Kalidas, Mayank Daga, Konstantinos Krommydas, and Wu-Chun Feng

On the Greenness of In-Situ and Post-Processing Visualization Pipelines 880
Vignesh Adhinarayanan, Wu-Chun Feng, Jonathan Woodring, David Rogers, and James Ahrens

Energy Modeling and Optimization for Tiled Nested-Loop Codes 888
Nirmal Prajapati, Waruna Ranasinghe, Vamshi Tandrapati, Rumen Andonov, Hristo Djidjev, and Sanjay Rajopadhye

Session 3: Efficiency

An Energy Efficiency Feature Survey of the Intel Haswell Processor 896
Daniel Hackenberg, Robert Schöne, Thomas Ilsche, Daniel Molka, Joseph Schuchart, and Robin Geyer

Iso-Power-Efficiency: An Approach to Scaling Application Codes with a Power Budget 905
Rogelio Long, Shirley Moore, and Barry Rountree

Using Dynamic Duty Cycle Modulation to Improve Energy Efficiency in High Performance Computing 911
Sridutt Bhalachandra, Allan Porterfield, and Jan F. Prins

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

PDSEC Introduction and Committees 919
Peter E. Strazdins, Raphaël Couturier, Keita Teranishi, John O'Donnell, Thomas Raubery, Gudula Rünger, and Laurence T. Yang

PDSEC Keynote 921
Naoya Maruyama

Session 1: Best Paper

Energy Consumption Reduction with DVFS for Message Passing Iterative Applications on Heterogeneous Architectures 922
Jean-Claude Charr, Raphaël Couturier, Ahmed Fanfakh, and Arnaud Giersch

Session 2: Performance

Quantifying the Effects of Contention on Parallel File Systems 932
Steven A. Wright and Stephen A. Jarvis

Highly Scalable Algorithms for the Sparse Grid Combination Technique 941
Peter E. Strazdins, Md. Mohsin Ali, and Brendan Harding

Predicting Optimal Power Allocation for CPU and DRAM Domains 951
Ananta Tiwari, Martin Schulz, and Laura Carrington

Session 3: Linear Algebra

Performance Evaluation of the Eigen Exa Eigensolver on Oakleaf-FX: Tridiagonalization Versus Pentadiagonalization 960
Takeshi Fukaya and Toshiyuki Imamura

A Resilient Framework for Iterative Linear Algebra Applications in X10 970
Sara S. Hamouda, Josh Milthorpe, Peter E. Strazdins, and Vijay Saraswat

Combining Backward and Forward Recovery to Cope with Silent Errors in Iterative Solvers 980
Massimiliano Fasi, Yves Robert, and Bora Uçar

TSIRM: A Two-Stage Iteration with Least-Squares Residual Minimization Algorithm to Solve Large Sparse Linear Systems 990
Raphaël Couturier, Lilia Ziane Khodja, and Christophe Guyeux

Session 4: GPUs and Manycore

Modeling Cooperative Threads to Project GPU Performance for Adaptive Parallelism 998
Jiayuan Meng, Thomas D. Uram, Vitali Morozov, Venkatram Vishwanath, and Kalyan Kumaran

GPU Accelerated Molecular Dynamics with Method of Heterogeneous Load Balancing 1008
Takuro Udagawa and Masakazu Sekijima

Parallel Methods for Optimizing High Order Constellations on GPUs 1014
Paolo Spallaccini, Farbod Kayhan, Stefano Chinnici, and Guido Montorsi

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

DPDNS Introduction and Committees 1024
Dimiter R. Avresky, Erik Maehle, Nectarios Koziris, and Anastassios Nanos

Session 1: Reliability and Threat-Detection

Performance and Energy Efficient Asymmetrically Reliable Caches for Multicore Architectures 1025
Sanem Arslan, Haluk Rahmi Topcuoglu, Mahmut Taylan Kandemir, and Oguz Tosun

Distributed Scheduling Algorithm for Highly Available Component Based Applications 1033
Marc Eduard Frincu

Optimizing Defensive Investments in Energy-Based Cyber-Physical Systems 1042
Paul Wood, Saurabh Bagchi, and Alefiya Hussain

Session 2: Fault Tolerance

Towards Detecting Patterns in Failure Logs of Large-Scale Distributed Systems 1052
Nentawe Gurumdimma, Arshad Jhumka, Maria Liakata, Edward Chuah, and James Browne

Communication Pattern-Based Distributed Snapshots in Large-Scale Systems 1062
Salem Saker and Adnan Agbaria

A Machine Learning-Based Framework for Building Application Failure Prediction Models 1072
Alessandro Pellegrini, Pierangelo Di Sanzo, and Dimiter R. Avresky

Session 3: Algorithms, Protocols, and Topologies

Trapezoid Quorum Protocol Dedicated to Erasure Resilient Coding Based Schemes 1082
Théodore Jean Richard Relaza, Jacques Jorda, and Abdelaziz M'Zoughi

A Distributed Greedy Heuristic for Computing Voronoi Tessellations with Applications Towards Peer-to-Peer Networks 1089
Brendan Benshoof, Andrew Rosen, Anu G. Bourgeois, and Robert W. Harrison

Dependability Modeling and Assessment of Complex Adaptive Networked Systems 1097
Kaliappa Ravindran

Workshop 15: PCO—Parallel Computing and Optimization

PCO Introduction and Committees 1106
Didier El Baz and Bora Uçar

PCO Keynote 1108
Alex Pothen

Session 1: Optimization Techniques for Parallel or Distributed Architectures

The Promethee Method for Cloud Brokering with Trust and Assurance Criteria 1109
Christian Toinard, Timothée Ravier, Christophe Cérin, and Yanik Ngoko

Buffer Allocation Based On-Chip Memory Optimization for Many-Core Platforms 1119
Maximilian Odendahl, Andrés Goens, Rainer Leupers, Gerd Ascheid, and Tomas Henriksson

Session 2: Combinatorial Scientific Computing and Parallel Optimization Algorithms

Semi-two-dimensional Partitioning for Parallel Sparse Matrix-Vector Multiplication 1125
Enver Kayaaslan, Bora Uçar, and Cevdet Aykanat

Parallel Asynchronous Modified Newton Methods for Network Flows 1135
Didier El Baz and Moussa Elkihel

A Branch-and-Estimate Heuristic Procedure for Solving Nonconvex Integer Optimization Problems 1143
Prashant Palkar and Ashutosh Mahajan

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

ParLearning Introduction and Committees 1152
Sutanay Choudhury, Arindam Pal, Anand Panangadan, and Yinglong Xia

ParLearning Keynotes 1154
David A. Bader, Yihua Huang, and Ananth Kalyanaraman

LFRTrainer: Large-Scale Face Recognition Training System 1157
Tao Luo, Yin Liao, Yurong Chen, Jianguo Li, and Victor Lee

Empowering Fast Incremental Computation over Large Scale Dynamic Graphs 1166
Charith Wickramaarachchi, Charalampos Chelmis, and Viktor K. Prasanna

Scaling Up the Training of Deep CNNs for Human Action Recognition 1172
M. Sai Rajeswar, A. Ravi Sankar, Vineeth N. Balasubramaniam, and C.D. Sudheer

Scalable Task-Parallel SGD on Matrix Factorization in Multicore Architectures 1178
Yusuke Nishioka and Kenjiro Taura

Analysis of Subgraph-Centric Distributed Shortest Path Algorithm 1185
Ravikant Dindokar, Neel Choudhury, and Yogesh Simmhan

Cost-Driven Scheduling for Deadline-Constrained Workflow on Multi-clouds 1191
Bing Lin, Wenzhong Guo, Guolong Chen, Naixue Xiong, and Rongrong Li

Workshop 17: JSSPP—Workshop on Job Scheduling Strategies for Parallel Processing

JSSPP Introduction and Committees 1199
Walfredo Cirne and Narayan Desai

Workshop 18: iWAPT—International Workshop on Automatic Performance Tuning

iWAPT Introduction and Committees 1200
Yusaku Yamamoto and Weichung Wang

iWAPT Invited Talks 1202
Ponnuswamy Sadayappan and Ray-Bing Chen

iWAPT Session 1

Auto-tuning Non-blocking Collective Communication Operations 1204
Youcef Barigou, Vishwanath Venkatesan, and Edgar Gabriel

Improved Internode Communication for Tile QR Decomposition for Multicore Cluster Systems 1214
Tomohiro Suzuki

iWAPT Session 2

Directive-Based Auto-Tuning for the Finite Difference Method on the Xeon Phi 1221
Takahiro Katagiri, Satoshi Ohshima, and Masaharu Matsumoto

Machine Learning Based Auto-Tuning for Enhanced OpenCL Performance Portability 1231
Thomas L. Falch and Anne C. Elster

A Roofline-Based Performance Estimator for Distributed Matrix-Multiply on Intel CnC 1241
Martin Kong, Louis-Noël Pouchet, and Ponnuswamy Sadayappan

iWAPT Session 3

Energy Prediction of OpenMP Applications Using Random Forest Modeling Approach 1251
Shajulin Benedict, R.S. Rejitha, Philipp Gschwandtner, Radu Prodan, and Thomas Fahringer

Auto-Tuning the Java Virtual Machine 1261
Sanath Jayasena, Milinda Fernando, Tharindu Rusira, Chalitha Perera, and Chamara Philips

Workshop 19: Julia—Invited Workshop: A New Approach to High Performance Technical Computing

Julia Introduction 1271
Alan Edelman

Author Index