

2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS 2016)

**Chicago, Illinois, USA
23-27 May 2016**

Pages 576-1152



**IEEE Catalog Number: CFP16023-POD
ISBN: 978-1-5090-2141-3**

**Copyright © 2016 by the Institute of Electrical and Electronics Engineers, Inc
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

******This publication is a representation of what appears in the IEEE Digital Libraries. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP16023-POD
ISBN (Print-On-Demand):	978-1-5090-2141-3
ISBN (Online):	978-1-5090-2140-6
ISSN:	1530-2075

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

2016 IEEE International Parallel and Distributed Processing Symposium

IPDPS 2016

Table of Contents

Message from the General Chair.....	xvi
Message from the Program Chair.....	xviii
Message from the Steering Co-Chairs.....	xix
Message from the Workshops Chair.....	xx
IPDPS 2016 Organization.....	xxii
IPDPS 2016 Reviewers.....	xxviii
IPDPS 2016 Technical Program.....	xxix

Keynote 1

Disruptive Research and Innovation.....	1
<i>Kai Li</i>	

Session 1: Graph Algorithms

Subgraph Counting: Color Coding Beyond Trees.....	2
<i>Venkatesan T. Chakaravarthy, Michael Kapralov, Prakash Murali, Fabrizio Petrini, Xinyu Que, Yogish Sabharwal, and Baruch Schieber</i>	
A Practical Parallel Algorithm for Diameter Approximation of Massive Weighted Graphs.....	12
<i>Matteo Ceccarello, Andrea Pietracaprina, Geppino Pucci, and Eli Upfal</i>	
Rabbit Order: Just-in-Time Parallel Reordering for Fast Graph Analysis.....	22
<i>Junya Arai, Hiroaki Shiokawa, Takeshi Yamamuro, Makoto Onizuka, and Sotetsu Iwamura</i>	
Distributed-Memory Algorithms for Maximum Cardinality Matching in Bipartite Graphs.....	32
<i>Ariful Azad and Aydin Buluç</i>	

Session 2: Software Environments and Tools

Automatic Parallel Pattern Detection in the Algorithm Structure Design Space	43
<i>Zia Ul Huda, Rohit Atre, Ali Jannesari, and Felix Wolf</i>	
ARCHER: Effectively Spotting Data Races in Large OpenMP Applications	53
<i>Simone Atzeni, Ganesh Gopalakrishnan, Zvonimir Rakamarić, Dong H. Ahn, Ignacio Laguna, Martin Schulz, Gregory L. Lee, Joachim Protze, and Matthias S. Müller</i>	
Algorithm and Architecture Independent Benchmarking with SEAK	63
<i>Nathan R. Tallent, Joseph B. Manzano, Nitin A. Gawande, Seunghwa Kang, Darren J. Kerbyson, Adolfo Hoisie, and Joseph K. Cross</i>	
Design and Implementation of a Parallel Research Kernel for Assessing Dynamic Load-Balancing Capabilities	73
<i>Evangelos Georganas, Rob F. Van der Wijngaart, and Timothy G. Mattson</i>	

Session 3: Network Architecture

VNRE: Flexible and Efficient Acceleration for Network Redundancy Elimination	83
<i>Xiongzi Ge, Yi Liu, Chengtao Lu, Jim Diehl, David H. C. Du, Liang Zhang, and Jian Chen</i>	
Analyzing Network Health and Congestion in Dragonfly-Based Supercomputers	93
<i>Abhinav Bhatele, Nikhil Jain, Yarden Livnat, Valerio Pascucci, and Peer-Timo Bremer</i>	
Random Regular Graph and Generalized De Bruijn Graph with k-Shortest Path Routing	103
<i>Peyman Faizian, Md Atiqul Mollah, Xin Yuan, Scott Pakin, and Michael Lang</i>	
Deflection Containment for Bufferless Network-on-Chips	113
<i>Xi-Yue Xiang and Nian-Feng Tzeng</i>	

Session 4: Application Optimization

RUPS: Fixing Relative Distances among Urban Vehicles with Context-Aware Trajectories	123
<i>Hongzi Zhu, Shan Chang, Li Lu, and Wei Zhang</i>	
Hybrid Dynamic Trees for Extreme-Resolution 3D Sparse Data Modeling	132
<i>Mohammad M. Hossain, Thomas M. Tucker, Thomas R. Kurfess, and Richard W. Vuduc</i>	
Optimization of an Electromagnetics Code with Multicore Wavefront Diamond Blocking and Multi-dimensional Intra-Tile Parallelization	142
<i>Tareq M. Malas, Julian Hornich, Georg Hager, Hatem Ltaief, Christoph Pflaum, and David E. Keyes</i>	
Order-Invariant Real Number Summation: Circumventing Accuracy Loss for Multimillion Summands on Multiple Parallel Architectures	152
<i>Patrick E. Small, Rajiv K. Kalia, Aiichiro Nakano, and Priya Vashishta</i>	

Session 5: Linear Algebra and Solvers

INV-ASKIT: A Parallel Fast Direct Solver for Kernel Matrices	161
<i>Chenhan D. Yu, William B. March, Bo Xiao, and George Biros</i>	
A Fast Tridiagonal Solver for Intel MIC Architecture	172
<i>Xinliang Wang, Wei Xue, Jidong Zhai, Yangtong Xu, Weimin Zheng, and Haixiang Lin</i>	
A Relaxed Synchronization Approach for Solving Parallel Quadratic Programming Problems with Guaranteed Convergence	182
<i>Kooktae Lee, Raktim Bhattacharya, Jyotikrishna Dass, V. N. S. Prithvi Sakuru, and Rabi N. Mahapatra</i>	
Enhancing Scalability and Load Balancing of Parallel Selected Inversion via Tree-Based Asynchronous Communication	192
<i>Mathias Jacquelin, Lin Lin, Nathan Wichmann, and Chao Yang</i>	

Session 6: Fault Tolerance and Resilience

Optimal Resilience Patterns to Cope with Fail-Stop and Silent Errors	202
<i>Anne Benoit, Aurélien Cavelan, Yves Robert, and Hongyang Sun</i>	
Reducing Waste in Extreme Scale Systems through Introspective Analysis	212
<i>Leonardo Bautista-Gomez, Ana Gainaru, Swann Perarnau, Devesh Tiwari, Saurabh Gupta, Christian Engelmann, Franck Cappello, and Marc Snir</i>	
Fault Modeling of Extreme Scale Applications Using Machine Learning	222
<i>Abhinav Vishnu, Hubertus van Dam, Nathan R. Tallent, Darren J. Kerbyson, and Adolfo Hoisie</i>	
Efficient Checkpointing of Multi-threaded Applications as a Tool for Debugging, Performance Tuning, and Resiliency	232
<i>Max Grossman and Vivek Sarkar</i>	

Session 7: Modeling and Evaluation

X: A Comprehensive Analytic Model for Parallel Machines	242
<i>Ang Li, Shuaiwen Leon Song, Eric Brugel, Akash Kumar, Daniel Chavarría-Miranda, and Henk Corporaal</i>	
NiMC: Characterizing and Eliminating Network-Induced Memory Contention	253
<i>Taylor Groves, Ryan E. Grant, and Dorian Arnold</i>	
An Early Performance Study of Large-Scale POWER8 SMP Systems	263
<i>Xing Liu, Daniele Buono, Fabio Checconi, Jee W. Choi, Xinyu Que, Fabrizio Petrini, John A. Gunnels, and Jeff A. Stuecheli</i>	
A Methodology for Modeling Dynamic and Static Power Consumption for Multicore Processors	273
<i>Bhavishya Goel and Sally A. McKee</i>	

Session 8: Graph Applications

Algorithmic Techniques for Solving Graph Problems on the Automata Processor	283
<i>Indranil Roy, Nagakishore Jammula, and Srinivas Aluru</i>	
A Case Study of Complex Graph Analysis in Distributed Memory: Implementation and Optimization	293
<i>George M. Slota, Sivasankaran Rajamanickam, and Kamesh Madduri</i>	
FastBFS: Fast Breadth-First Graph Search on a Single Server	303
<i>Shuhan Cheng, Guangyan Zhang, Jiwu Shu, Qingda Hu, and Weimin Zheng</i>	
GraphPad: Optimized Graph Primitives for Parallel and Distributed Platforms	313
<i>Michael J. Anderson, Narayanan Sundaram, Nadathur Satish, Md. Mostofa Ali Patwary, Theodore L. Willke, and Pradeep Dubey</i>	

Session 9: Cloud Resource Allocation

On First Fit Bin Packing for Online Cloud Server Allocation	323
<i>Xueyan Tang, Yusen Li, Runtian Ren, and Wentong Cai</i>	
Smoothed Online Resource Allocation in Multi-tier Distributed Cloud Networks	333
<i>Lei Jiao, Antonia Tulino, Jaime Llorca, Yue Jin, and Alessandra Sala</i>	
Dynamic Acceleration of Parallel Applications in Cloud Platforms by Adaptive Time-Slice Control	343
<i>Song Wu, Zhenjiang Xie, Haibao Chen, Sheng Di, Xinyu Zhao, and Hai Jin</i>	
Mystic: Predictive Scheduling for GPU Based Cloud Servers Using Machine Learning	353
<i>Yash Ukidave, Xiangyu Li, and David Kaeli</i>	

Session 10: Memory Management

TintMalloc: Reducing Memory Access Divergence via Controller-Aware Coloring	363
<i>Xing Pan, Ysaswini Jyothi Gownivaripalli, and Frank Mueller</i>	
Markov Chain-Based Adaptive Scheduling in Software Transactional Memory	373
<i>Pierangelo Di Sanzo, Marco Sannicandro, Bruno Ciciani, and Francesco Quaglia</i>	
MEMTUNE: Dynamic Memory Management for In-Memory Data Analytic Platforms	383
<i>Luna Xu, Min Li, Li Zhang, Ali R. Butt, Yandong Wang, and Zane Zhenhua Hu</i>	
High-Performance Hybrid Key-Value Store on Modern Clusters with RDMA Interconnects and SSDs: Non-blocking Extensions, Designs, and Benefits	393
<i>Dipti Shankar, Xiaoyi Lu, Nusrat Islam, Md. Wasi-Ur-Rahman, and Dhableswar K. (DK) Panda</i>	

Session 11: Scheduling and Resource Management

GreenMatch: Renewable-Aware Workload Scheduling for Massive Storage Systems	403
<i>Xiaoyang Qu, Jiguang Wan, Jun Wang, Liqiong Liu, Dan Luo, and Changsheng Xie</i>	
CATA: Criticality Aware Task Acceleration for Multicore Processors	413
<i>Emilio Castillo, Miquel Moreto, Marc Casas, Lluc Alvarez, Enrique Vallejo, Kallia Chronaki, Rosa Badia, Jose Luis Bosque, Ramon Beivide, Eduard Ayguade, Jesus Labarta, and Mateo Valero</i>	
TECfan: Coordinating Thermoelectric Cooler, Fan, and DVFS for CMP Energy Optimization	423
<i>Wenli Zheng, Kai Ma, and Xiaorui Wang</i>	
Utility Maximizing Thread Assignment and Resource Allocation	433
<i>Pan Lai, Rui Fan, Wei Zhang, and Fang Liu</i>	

Session 12: Scientific Applications (1)

A Hybrid Decomposition Parallel Algorithm for Multi-scale Simulation of Viscoelastic Fluids	443
<i>Xiao-Wei Guo, Xin-Hai Xu, Qian Wang, Hao Li, Xiao-Guang Ren, Liyang Xu, and Xue-Jun Yang</i>	
A Hartree-Fock Application Using UPC++ and the New DArray Library	453
<i>David Ozog, Amir Kamil, Yili Zheng, Paul Hargrove, Jeff R. Hammond, Allen Malony, Wibe de Jong, and Kathy Yelick</i>	
A Fast Selected Inversion Algorithm for Green's Function Calculation in Many-Body Quantum Monte Carlo Simulations	463
<i>Chengming Jiang, Zhaojun Bai, and Richard Scalettar</i>	

Keynote 2

Memory, Storage and Processing in Future Parallel and Distributed Processing Systems	473
<i>Thomas Pawlowski</i>	

Session 13: Clustering and Partitioning

A New Approximation Algorithm for Matrix Partitioning in Presence of Strongly Heterogeneous Processors	474
<i>Olivier Beaumont, Lionel Eyraud-Dubois, and Thomas Lambert</i>	
Structural Clustering: A New Approach to Support Performance Analysis at Scale	484
<i>Matthias Weber, Ronny Brendel, Tobias Hilbrich, Kathryn Mohror, Martin Schulz, and Holger Brunst</i>	

PANDA: Extreme Scale Parallel K-Nearest Neighbor on Distributed Architectures	494
<i>Md. Mostofa Ali Patwary, Nadathur Rajagopalan Satish, Narayanan Sundaram, Jialin Liu, Peter Sadowski, Evan Racah, Suren Byna, Craig Tull, Wahid Bhimji, Prabhat, and Pradeep Dubey</i>	

DataNet: A Data Distribution-Aware Method for Sub-Dataset Analysis on Distributed File Systems	504
<i>Jun Wang, Jiangling Yin, Jian Zhou, Xuhong Zhang, and Ruijun Wang</i>	

Session 14: Accelerated Computing

Synchronization Trade-Offs in GPU Implementations of Graph Algorithms	514
<i>Rashid Kaleem, Anand Venkat, Sreepathi Pai, Mary Hall, and Keshav Pingali</i>	

Eliminating Intra-Warp Load Imbalance in Irregular Nested Patterns via Collaborative Task Engagement	524
<i>Farzad Khorasani, Bryan Rowe, Rajiv Gupta, and Laxmi N. Bhuyan</i>	

Compiler-Assisted Workload Consolidation for Efficient Dynamic Parallelism on GPU	534
<i>Hancheng Wu, Da Li, and Michela Becchi</i>	

OpenACC to FPGA: A Framework for Directive-Based High-Performance Reconfigurable Computing	544
<i>Seyong Lee, Jungwon Kim, and Jeffrey S. Vetter</i>	

Session 15: Memory Hierarchy

Architecting and Programming a Hardware-Incoherent Multiprocessor Cache Hierarchy	555
<i>Wooil Kim, Sanket Tavarageri, P. Sadayappan, and Josep Torrellas</i>	

Refree: A Refresh-Free Hybrid DRAM/PCM Main Memory System	566
<i>Bahareh Pourshirazi and Zhichun Zhu</i>	

Re-NUCA: A Practical NUCA Architecture for ReRAM Based Last-Level Caches	576
<i>Jagadish B. Kotra, Mohammad Arjomand, Diana Guttman, Mahmut T. Kandemir, and Chita R. Das</i>	

Evaluating and Improving Thread-Level Speculation in Hardware Transactional Memories	586
<i>Juan Salamanca, José Nelson Amaral, and Guido Araujo</i>	

Session 16: Optimization Techniques

System Noise Revisited: Enabling Application Scalability and Reproducibility with SMT	596
<i>Edgar A. León, Ian Karlin, and Adam T. Moody</i>	

Key/Value-Enabled Flash Memory for Complex Scientific Workflows with On-Line Analysis and Visualization	608
<i>Stefan Eilemann, Fabien Delalandre, Jon Bernard, Judit Planas, Felix Schuermann, John Biddiscombe, Costas Bekas, Alessandro Curioni, Bernard Metzler, Peter Kaltstein, Peter Morjan, Joachim Fenkes, Ralph Bellofatto, Lars Schneidenbach, T. J. Christopher Ward, and Blake G. Fitch</i>	
Fast Classification of MPI Applications Using Lamport's Logical Clocks	618
<i>Zhou Tong, Scott Pakin, Michael Lang, and Xin Yuan</i>	
Online-Autotuning of Parallel SAH kD-Trees	628
<i>Martin Tillmann, Philip Pfaffe, Christopher Kaag, and Walter F. Tichy</i>	
Session 17: Communication Efficiency and Avoidance Algorithms	
Polynomial-Time Construction of Optimal MPI Derived Datatype Trees	638
<i>Robert Ganian, Martin Kalany, Stefan Szeider, and Jesper Larsson Träff</i>	
Write-Avoiding Algorithms	648
<i>Erin Carson, James Demmel, Laura Grigori, Nicholas Knight, Penporn Koanantakool, Oded Schwartz, and Harsha Vardhan Simhadri</i>	
Communication Efficient Algorithms for Top-k Selection Problems	659
<i>Lorenz Hübschle-Schneider and Peter Sanders</i>	
Minimal Aggregated Shared Memory Messaging on Distributed Memory Supercomputers	669
<i>Benjamin F. Jamroz and John M. Dennis</i>	
Session 18: Distributed Algorithms	
Never Say Never — Probabilistic and Temporal Failure Detectors	679
<i>Dacfez Dzung, Rachid Guerraoui, David Kozhaya, and Yvonne-Anne Pignolet</i>	
Gathering a Closed Chain of Robots on a Grid	689
<i>Sebastian Abshoff, Andreas Cord-Landwehr, Matthias Fischer, Daniel Jung, and Friedhelm Meyer auf der Heide</i>	
On Competitive Algorithms for Approximations of Top-k-Position Monitoring of Distributed Streams	700
<i>Alexander Mäcker, Manuel Malatyali, and Friedhelm Meyer auf der Heide</i>	
Towards a Restrained Use of Non-Equivocation for Achieving Iterative Approximate Byzantine Consensus	710
<i>Chuanyou Li, Michel Hurfin, Yun Wang, and Lei Yu</i>	

Session 19: I/O and Storage

Storage-Optimized Data-Atomic Algorithms for Handling Erasures and Errors in Distributed Storage Systems	720
<i>Kishori M. Konwar, N. Prakash, Erez Kantor, Nancy Lynch, Muriel Médard, and Alexander A. Schwarzmann</i>	
Fast Error-Bounded Lossy HPC Data Compression with SZ	730
<i>Sheng Di and Franck Cappello</i>	
I/O Aware Power Shifting	740
<i>Lee Savoie, David K. Lowenthal, Bronis R. de Supinski, Tanzima Islam, Kathryn Mohror, Barry Rountree, and Martin Schulz</i>	
On the Root Causes of Cross-Application I/O Interference in HPC Storage Systems	750
<i>Orcun Yildiz, Matthieu Dorier, Shadi Ibrahim, Rob Ross, and Gabriel Antoniu</i>	

Session 20: Scientific Applications (2)

Exploiting Variant-Based Parallelism for Data Mining of Space Weather Phenomena	760
<i>Michael Gowanlock, David M. Blair, and Victor Pankratius</i>	
Solving Open MIP Instances with ParaSCIP on Supercomputers Using up to 80,000 Cores	770
<i>Yuji Shinano, Tobias Achterberg, Timo Berthold, Stefan Heinz, Thorsten Koch, and Michael Winkler</i>	
AAlign: A SIMD Framework for Pairwise Sequence Alignment on x86-Based Multi- and Many-Core Processors	780
<i>Kaixi Hou, Hao Wang, and Wu-Chun Feng</i>	
Mendel: A Distributed Storage Framework for Similarity Searching over Sequencing Data	790
<i>Cameron Tolooee, Sangmi Lee Pallickara, and Asa Ben-Hur</i>	

Keynote 3

Unlocking the Mysteries of the Universe with Supercomputers	800
<i>Katrin Heitmann</i>	

Best Papers Session

ZNN — A Fast and Scalable Algorithm for Training 3D Convolutional Networks on Multi-core and Many-Core Shared Memory Machines	801
<i>Aleksandar Zlateski, Kisuk Lee, and H. Sebastian Seung</i>	

Stochastic Matrix-Function Estimators: Scalable Big-Data Kernels with High Performance	812
<i>Peter W. J. Staar, Panagiotis Kl. Barkoutsos, Roxana Istrate, A. Cristiano I. Malossi, Ivano Tavernelli, Nikolaj Moll, Heiner Giefers, Christoph Hagleitner, Costas Bekas, and Alessandro Curioni</i>	
Discrete Cache Insertion Policies for Shared Last Level Cache Management on Large Multicores	822
<i>Aswinkumar Sridharan and André Seznec</i>	
Massively Parallel First-Principles Simulation of Electron Dynamics in Materials	832
<i>Erik W. Draeger, Xavier Andrade, John A. Gunnels, Abhinav Bhatele, André Schleife, and Alfredo A. Correa</i>	
Session 21: Numerical Algorithms	
Communication-Avoiding Parallel Sparse-Dense Matrix-Matrix Multiplication	842
<i>Penporn Koanantakool, Ariful Azad, Aydin Buluç, Dmitriy Morozov, Sang-Yun Oh, Leonid Oliker, and Katherine Yelick</i>	
Petascale Local Time Stepping for the ADER-DG Finite Element Method	854
<i>Alexander Breuer, Alexander Heinecke, and Michael Bader</i>	
Asymptotic Optimality of Parallel Short Division	864
<i>Niall Emmart and Charles Weems</i>	
High Performance Parallel Stochastic Gradient Descent in Shared Memory	873
<i>Scott Sallinen, Nadathur Satish, Mikhail Smelyanskiy, Samantika S. Sury, and Christopher Ré</i>	
Session 22: Graphs and Tensors	
Optimal Algorithms for Graphs and Images on a Shared Memory Mesh	883
<i>Yujie An and Quentin F. Stout</i>	
Parallel Graph Coloring for Manycore Architectures	892
<i>Mehmet Deveci, Erik G Boman, Karen D. Devine, and Sivasankaran Rajamanickam</i>	
A Medium-Grained Algorithm for Sparse Tensor Factorization	902
<i>Shaden Smith and George Karypis</i>	
Parallel Tensor Compression for Large-Scale Scientific Data	912
<i>Woody Austin, Grey Ballard, and Tamara G. Kolda</i>	
Session 23: Runtime Systems	
GinFlow: A Decentralised Adaptive Workflow Execution Manager	923
<i>Javier Rojas Balderrama, Matthieu Simonin, and Cédric Tedeschi</i>	

Hierarchical Parallel Dynamic Dependence Analysis for Recursively Task-Parallel Programs	933
<i>Nikolaos Papakonstantinou, Foivos S. Zakkak, and Polyvios Pratikakis</i>	
MPMD Framework for Offloading Load Balance Computation	943
<i>Olga Pearce, Todd Gamblin, Bronis R. de Supinski, Martin Schulz, and Nancy M. Amato</i>	
Integrating Abstractions to Enhance the Execution of Distributed Applications	953
<i>Matteo Turilli, Feng Liu, Zhao Zhang, Andre Merzky, Michael Wilde, Jon Weissman, Daniel S. Katz, and Shantenu Jha</i>	

Session 24: GPUs

cusFFT: A High-Performance Sparse Fast Fourier Transform Algorithm on GPUs	963
<i>Cheng Wang, Sunita Chandrasekaran, and Barbara Chapman</i>	
Balancing Scalar and Vector Execution on GPU Architectures	973
<i>Zhongliang Chen and David Kaeli</i>	
Exploiting Maximal Overlap for Non-Contiguous Data Movement Processing on Modern GPU-Enabled Systems	983
<i>C-H. Chu, K. Hamidouche, A. Venkatesh, D. S. Banerjee, H. Subramoni, and Dhabaleswar K. (DK) Panda</i>	
Online Algorithm-Based Fault Tolerance for Cholesky Decomposition on Heterogeneous Systems with GPUs	993
<i>Jieyang Chen, Xin Liang, and Zizhong Chen</i>	

Session 25: Scheduling

Reusable Resource Scheduling via Colored Interval Covering	1003
<i>Venkatesan T. Chakaravarthy, Sreyash Kenkre, Sakib A. Mondal, Vinayaka Pandit, and Yogish Sabharwal</i>	
Partitioned Feasibility Tests for Sporadic Tasks on Heterogeneous Machines	1013
<i>Shaurya Ahuja, Kefu Lu, and Benjamin Moseley</i>	
Are Static Schedules so Bad? A Case Study on Cholesky Factorization	1021
<i>Emmanuel Agullo, Olivier Beaumont, Lionel Eyraud-Dubois, and Suraj Kumar</i>	

Session 26: System Software

Optimization and Analysis of MPI Collective Communication on Fat-Tree Networks	1031
<i>Sameer Kumar, Sameh S. Sharkawi, and K. A. Nysal Jan</i>	
On the Scalability, Performance Isolation and Device Driver Transparency of the IHK/McKernel Hybrid Lightweight Kernel	1041
<i>Balazs Gerofi, Masamichi Takagi, Atsushi Hori, Gou Nakamura, Tomoki Shirasawa, and Yutaka Ishikawa</i>	

ZCCloud: Exploring Wasted Green Power for High-Performance Computing1051
Fan Yang and Andrew A. Chien

Agile Live Migration of Virtual Machines1061
*Umesh Deshpande, Danny Chan, Ten-Young Guh, James Edouard, Kartik Gopalan,
and Nilton Bila*

Session 27: Security and Fault Tolerance

Lazy Repair for Addition of Fault-Tolerance to Distributed Programs1071
Mohammad Roohitavaf, Yiyan Lin, and Sandeep S. Kulkarni

Security RBSG: Protecting Phase Change Memory with Security-Level Adjustable
Dynamic Mapping1081
*Fangting Huang, Dan Feng, Wen Xia, Wen Zhou, Yucheng Zhang, Min Fu,
Chuntao Jiang, and Yukun Zhou*

Mitigation of Denial of Service Attack with Hardware Trojans in NoC Architectures1091
Travis Boraten and Avinash Karanth Kodi

CRC-Based Memory Reliability for Task-Parallel HPC Applications1101
Omer Subasi, Osman Unsal, Jesus Labarta, Gulay Yalcin, and Adrian Cristal

Session 28: Data Streaming

Differentiated Scheduling of Response-Critical and Best-Effort Wide-Area Data
Transfers1113
Rajkumar Kettimuthu, Gagan Agrawal, P. Sadayappan, and Ian Foster

High Performance Pattern Matching Using the Automata Processor1123
Indranil Roy, Ankit Srivastava, Marziyeh Nourian, Michela Becchi, and Srinivas Aluru

GPU-Accelerated Outlier Detection for Continuous Data Streams1133
Chandima HewaNadungodage, Yuni Xia, and John Jaehwan Lee

NEPTUNE: Real Time Stream Processing for Internet of Things and Sensing
Environments1143
Thilina Buddhika and Shrideep Pallickara

Author Index