

# **2014 IEEE 28th International Parallel and Distributed Processing Symposium**

**(IPDPS 2014)**

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

**Pages 1-646**



**IEEE Catalog Number: CFP14023-POD  
ISBN: ;9: /3/69; ;/5: 23/:"**

# 2014 IEEE 28th International Parallel & Distributed Processing Symposium

## IPDPS 2014

### Table of Contents

Message from the General Chair.....	xvi
Message from the Program Chair.....	xviii
Message from the Steering Co-chairs.....	xx
Message from the Workshops Chair.....	xxi
IPDPS 2014 Organization.....	xxii
IPDPS 2014 Reviewers.....	xxx
IPDPS 2014 Technical Program.....	xxxix

---

#### Keynote Speaker 1

Scalability--Centric HPC System Design.....	3
<i>Yutong Lu</i>	

#### Session 1: Algorithms for Resource Management and Awareness

Cost-Optimal Execution of Boolean Query Trees with Shared Streams.....	7
<i>Henri Casanova, Lipyeow Lim, Yves Robert, Frédéric Vivien, and Dounia Zaidouni</i>	
It's About Time: On Optimal Virtual Network Embeddings under Temporal Flexibilities.....	17
<i>Matthias Rost, Stefan Schmid, and Anja Feldmann</i>	
Exploiting Geometric Partitioning in Task Mapping for Parallel Computers.....	27
<i>Mehmet Deveci, Sivasankaran Rajamanickam, Vitus J. Leung, Kevin Pedretti, Stephen L. Olivier, David P. Bunde, Umit V. Çatalyürek, and Karen Devine</i>	
Communication-Efficient Distributed Variance Monitoring and Outlier Detection for Multivariate Time Series.....	37
<i>Mohse Gabel, Assaf Schuster, and Daniel Keren</i>	

#### Session 2: Big Data Processing

MobiStreams: A Reliable Distributed Stream Processing System for Mobile Devices.....	51
<i>Huayong Wang and Li-Shiuan Peh</i>	

MapReuse: Reusing Computation in an In-Memory MapReduce System .....	61
<i>Devesh Tiwari and Yan Solihin</i>	
PAGE: A Framework for Easy PArallelization of GENomic Applications .....	72
<i>Mucahid Kutlu and Gagan Agrawal</i>	
Pythia: Faster Big Data in Motion through Predictive Software-Defined Network Optimization at Runtime .....	82
<i>Marcelo Veiga Neves, César A.F. De Rose, Kostas Katrinis, and Hubertus Franke</i>	
<b>Session 3: GPU</b>	
A Case for a Flexible Scalar Unit in SIMT Architecture .....	93
<i>Yi Yang, Ping Xiang, Michael Mantor, Norman Rubin, Lisa Hsu, Qunfeng Dong, and Huiyang Zhou</i>	
Scalar Waving: Improving the Efficiency of SIMD Execution on GPUs .....	103
<i>Ayse Yilmazer, Zhongliang Chen, and David Kaeli</i>	
Power and Performance Characterization and Modeling of GPU-Accelerated Systems .....	113
<i>Yuki Abe, Hiroshi Sasaki, Shinpei Kato, Koji Inoue, Masato Edahiro, and Martin Peres</i>	
Energy Efficient HPC on Embedded SoCs: Optimization Techniques for Mali GPU .....	123
<i>Ivan Grasso, Petar Radojkovic, Nikola Rajovic, Isaac Gelado, and Alex Ramirez</i>	
<b>Session 4: I/O, Storage, and Networking</b>	
Bursting the Cloud Data Bubble: Towards Transparent Storage Elasticity in IaaS Clouds .....	135
<i>Bogdan Nicolae, Pierre Riteau, and Kate Keahey</i>	
Scibox: Online Sharing of Scientific Data via the Cloud .....	145
<i>Jian Huang, Xuechen Zhang, Greg Eisenhauer, Karsten Schwan, Matthew Wolf, Stephane Ethier, and Scott Klasky</i>	
CALCioM: Mitigating I/O Interference in HPC Systems through Cross-Application Coordination .....	155
<i>Mathieu Dorier, Gabriel Antoniu, Rob Ross, Dries Kimpe, and Shadi Ibrahim</i>	
Active Measurement of the Impact of Network Switch Utilization on Application Performance .....	165
<i>Marc Casas and Greg Bronevetsky</i>	
<b>Session 5: Multi-core Algorithms</b>	
Multi-resource Real-Time Reader/Writer Locks for Multiprocessors .....	177
<i>Bryan C. Ward and James H. Anderson</i>	
Remote Invalidation: Optimizing the Critical Path of Memory Transactions .....	187
<i>Ahmed Hassan, Roberto Palmieri, and Binoy Ravindran</i>	

Revisiting Asynchronous Linear Solvers: Provable Convergence Rate through Randomization .....	198
<i>H. Avron, A. Druinsky, and A. Gupta</i>	
Accelerating MPI Collective Communications through Hierarchical Algorithms Without Sacrificing Inter-Node Communication Flexibility .....	208
<i>Benjamin S. Parsons and Vijay S. Pai</i>	
<b>Session 6: Computational Biology</b>	
Enabling In-Situ Data Analysis for Large Protein-Folding Trajectory Datasets .....	221
<i>Boyu Zhang, Trilce Estrada, Pietro Cicotti, and Michela Taufer</i>	
Overcoming the Limitations Posed by TCR-beta Repertoire Modeling through a GPU-Based In-Silico DNA Recombination Algorithm .....	231
<i>Gregory Striemer, Harsha Krovi, Ali Akoglu, Benjamin Vincent, Ben Hopson, Jeffrey Frelinger, and Adam Buntzman</i>	
Parallel Mutual Information Based Construction of Whole-Genome Networks on the Intel (R) Xeon Phi (TM) Coprocessor .....	241
<i>Sanchit Misra, Kiran Pamnany, and Srinivas Aluru</i>	
cuBLASTP: Fine-Grained Parallelization of Protein Sequence Search on a GPU .....	251
<i>Jing Zhang, Hao Wang, Heshan Lin, and Wu-Chun Feng</i>	
<b>Session 7: Interconnection Network</b>	
Skywalk: A Topology for HPC Networks with Low-Delay Switches .....	263
<i>Ikki Fujiwara, Michihiro Koibuchi, Hiroki Matsutani, and Henri Casanova</i>	
LFTI: A New Performance Metric for Assessing Interconnect Designs for Extreme-Scale HPC Systems .....	273
<i>Xin Yuan, Santosh Mahapatra, Michael Lang, and Scott Pakin</i>	
An Improved Router Design for Reliable On-Chip Networks .....	283
<i>Pavan Poluri and Ahmed Louri</i>	
Energy-Efficient Time-Division Multiplexed Hybrid-Switched NoC for Heterogeneous Multicore Systems .....	293
<i>Jieming Yin, Pingqiang Zhou, Sachin S. Sapatnekar, and Antonia Zhai</i>	
<b>Session 8: System-Level Resource Management</b>	
Heterogeneity-Aware Workload Placement and Migration in Distributed Sustainable Datacenters .....	307
<i>Dazhao Cheng, Changjun Jiang, and Xiaobo Zhou</i>	
Online Server and Workload Management for Joint Optimization of Electricity Cost and Carbon Footprint Across Data Centers .....	317
<i>Zahra Abbasi, Madhurima Pore, and Sandeep K.S. Gupta</i>	

Cost-Efficient and Resilient Job Life-Cycle Management on Hybrid Clouds .....	327
<i>Hsuan-Yi Chu and Yogesh Simmhan</i>	
A Coprocessor Sharing-Aware Scheduler for Xeon Phi-Based Compute Clusters .....	337
<i>Giuseppe Coviello, Srihari Cadambi, and Srimat Chakradhar</i>	
<b>Session 9: GPU Algorithms</b>	
Work-Efficient Parallel GPU Methods for Single-Source Shortest Paths .....	349
<i>Andrew Davidson, Sean Baxter, Michael Garland, and John D. Owens</i>	
Efficient Multi-GPU Computation of All-Pairs Shortest Paths .....	360
<i>Hristo Djidjev, Sunil Thulasidasan, Guillaume Chapuis, Rumen Andonov, and Dominique Lavenier</i>	
An Efficient GPU General Sparse Matrix-Matrix Multiplication for Irregular Data .....	370
<i>Weifeng Liu and Brian Vinter</i>	
Improving the Performance of CA-GMRES on Multicores with Multiple GPUs .....	382
<i>Ichitaro Yamazaki, Hartwig Anzt, Stanimire Tomov, Mark Hoemmen, and Jack Dongarra</i>	
<b>Session 10: Graph and Network Processing</b>	
How Well Do Graph-Processing Platforms Perform? An Empirical Performance Evaluation and Analysis .....	395
<i>Yong Guo, Marcin Biczak, Ana Lucia Varbanescu, Alexandru Iosup, Claudio Martella, and Theodore L. Willke</i>	
Complex Network Analysis Using Parallel Approximate Motif Counting .....	405
<i>George M. Slota and Kamesh Madduri</i>	
Finding Motifs in Biological Sequences Using the Micron Automata Processor .....	415
<i>Indranil Roy and Srinivas Aluru</i>	
Traversing Trillions of Edges in Real Time: Graph Exploration on Large-Scale Parallel Machines .....	425
<i>Fabio Checconi and Fabrizio Petrini</i>	
<b>Session 11: Modeling, Simulation, and Reliability</b>	
TBPoint: Reducing Simulation Time for Large-Scale GPGPU Kernels .....	437
<i>Jen-Cheng Huang, Lifeng Nai, Hyesoon Kim, and Hsien-Hsin S. Lee</i>	
Algorithmic Time, Energy, and Power on Candidate HPC Compute Building Blocks .....	447
<i>Jee Choi, Marat Dukhan, Xing Liu, and Richard Vuduc</i>	
Characterization of Impact of Transient Faults and Detection of Data Corruption Errors in Large-Scale N-Body Programs Using Graphics Processing Units .....	458
<i>Keun Soo Yim</i>	
Analytically Modeling Application Execution for Software-Hardware Co-design .....	468
<i>Jichi Guo, Jiayuan Meng, Qing Yi, Vitali Morozov, and Kalyan Kumaran</i>	

## Session 12: Accelerator Application Development and Optimization

Interactive Program Debugging and Optimization for Directive-Based, Efficient GPU Computing .....	481
<i>Seyong Lee, Dong Li, and Jeffrey S. Vetter</i>	
Unified Development for Mixed Multi-GPU and Multi-coprocessor Environments Using a Lightweight Runtime Environment .....	491
<i>Azzam Haidar, Chongxiao Cao, Asim Yarkhan, Piotr Luszczek, Stanimire Tomov, Khairul Kabir, and Jack Dongarra</i>	
Nitro: A Framework for Adaptive Code Variant Tuning .....	501
<i>Saurav Muralidharan, Manu Shantharam, Mary Hall, Michael Garland, and Bryan Catanzaro</i>	

## Keynote Speaker 2

Reading the Tea-Leaves: How Architecture Has Evolved at the High End .....	515
<i>Peter Kogge</i>	

## Session 13: Combinatorial Algorithms

New Effective Multithreaded Matching Algorithms .....	519
<i>Fredrik Manne and Mahantesh Halappanavar</i>	
A Medium-Grain Method for Fast 2D Bipartitioning of Sparse Matrices .....	529
<i>Daniël M. Pelt and Rob H. Bisseling</i>	
Bipartite Matching Heuristics with Quality Guarantees on Shared Memory Parallel Computers .....	540
<i>Fanny Dufossé, Kamer Kaya, and Bora Uçar</i>	
BFS and Coloring-Based Parallel Algorithms for Strongly Connected Components and Related Problems .....	550
<i>George M. Slota, Sivasankaran Rajamanickam, and Kamesh Madduri</i>	

## Session 14: Large Scale Scientific Applications

Large-Scale Hydrodynamic Brownian Simulations on Multicore and Manycore Architectures .....	563
<i>Xing Liu and Edmond Chow</i>	
Using Load Balancing to Scalably Parallelize Sampling-Based Motion Planning Algorithms .....	573
<i>Adam Fidel, Sam Ade Jacobs, Shishir Sharma, Nancy M. Amato, and Lawrence Rauchwerger</i>	
Petascale Application of a Coupled CPU-GPU Algorithm for Simulation and Analysis of Multiphase Flow Solutions in Porous Medium Systems .....	583
<i>James E. McClure, Hao Wang, Jan F. Prins, Cass T. Miller, and Wu-Chun Feng</i>	

A Spatio-temporal Coupling Method to Reduce the Time-to-Solution of Cardiovascular Simulations .....	593
<i>Amanda Randles and Efthimios Kaxiras</i>	

## **Session 15: Multicore and Transactional Memory**

Mitigating the Mismatch between the Coherence Protocol and Conflict Detection in Hardware Transactional Memory .....	605
<i>Lihang Zhao, Lizhong Chen, and Jeffrey Draper</i>	
Performance and Energy Analysis of the Restricted Transactional Memory Implementation on Haswell .....	615
<i>Bhavishya Goel, Ruben Titos-Gil, Anurag Negi, Sally A. McKee, and Per Stenstrom</i>	
Runtime-Guided Cache Coherence Optimizations in Multi-core Architectures .....	625
<i>Madhavan Manivannan and Per Stenstrom</i>	
High Performance Alltoall and Allgather Designs for InfiniBand MIC Clusters .....	637
<i>Akshay Venkatesh, Sreeram Potluri, Raghunath Rajachandrasekar, Miao Luo, Khaled Hamidouche, and Dhabaleswar K. Panda</i>	

## **Session 16: HPC Operating Systems and Runtime Systems**

HPMMAP: Lightweight Memory Management for Commodity Operating Systems .....	649
<i>Brian Kocoloski and John Lange</i>	
Victim Selection and Distributed Work Stealing Performance: A Case Study .....	659
<i>Swann Perarnau and Mitsuhsa Sato</i>	
Power-Efficient Multiple Producer-Consumer .....	669
<i>Ramy Medhat, Borzoo Bonakdarpour, and Sebastian Fischmeister</i>	
Efficient Data Race Detection for C/C++ Programs Using Dynamic Granularity .....	679
<i>Young Wn Song and Yann-Hang Lee</i>	

## **Session 17: Algorithms for Distributed Computing**

Improved Time Bounds for Linearizable Implementations of Abstract Data Types .....	691
<i>Jiaqi Wang, Edward Talmage, Hyunyoung Lee, and Jennifer L. Welch</i>	
DEX: Self-Healing Expanders .....	702
<i>Gopal Pandurangan, Peter Robinson, and Amitabh Trehan</i>	
Fair Maximal Independent Sets .....	712
<i>Jeremy T. Fineman, Calvin Newport, Micah Sherr, and Tonghe Wang</i>	

## Session 18: Milestones at the Petascale

Balancing CPU-GPU Collaborative High-Order CFD Simulations on the Tianhe-1A Supercomputer .....	725
<i>Chuanfu Xu, Lilun Zhang, Xiaogang Deng, Jianbin Fang, Guangxue Wang, Wei Cao, Yonggang Che, Yongxian Wang, and Wei Liu</i>	
Shedding Light on Lithium/Air Batteries Using Millions of Threads on the BG/Q Supercomputer .....	735
<i>Valéry Weber, Costas Bekas, Teodoro Laino, Alessandro Curioni, Adam Bertsch, and Scott Futral</i>	
Enabling and Scaling a Global Shallow-Water Atmospheric Model on Tianhe-2 .....	745
<i>Wei Xue, Chao Yang, Haohuan Fu, Xinliang Wang, Yangtong Xu, Lin Gan, Yutong Lu, and Xiaoqian Zhu</i>	
Overcoming the Scalability Challenges of Epidemic Simulations on Blue Waters .....	755
<i>Jae-Seung Yeom, Abhinav Bhatele, Keith Bisset, Eric Bohm, Abhishek Gupta, Laxmikant V. Kale, Madhav Marathe, Dimitrios S. Nikolopoulos, Martin Schulz, and Lukasz Wesolowski</i>	

## Session 19: Storage and Reliability

POD: Performance Oriented I/O Deduplication for Primary Storage Systems in the Cloud .....	767
<i>Bo Mao, Hong Jiang, Suzhen Wu, and Lei Tian</i>	
Pipelined Compaction for the LSM-Tree .....	777
<i>Zigang Zhang, Yinliang Yue, Bingsheng He, Jin Xiong, Mingyu Chen, Lixin Zhang, and Ninghui Sun</i>	
EDM: An Endurance-Aware Data Migration Scheme for Load Balancing in SSD Storage Clusters .....	787
<i>Jiaxin Ou, Jiwu Shu, Youyou Lu, Letian Yi, and Wei Wang</i>	

## Session 20: Map/Reduce and Big Data

Characterization and Optimization of Memory-Resident MapReduce on HPC Systems .....	799
<i>Yandong Wang, Robin Goldstone, Weikuan Yu, and Teng Wang</i>	
MIC-SVM: Designing a Highly Efficient Support Vector Machine for Advanced Modern Multi-core and Many-Core Architectures .....	809
<i>Yang You, Shuaiwen Leon Song, Haohuan Fu, Andres Marquez, Maryam Mehri Dehnavi, Kevin Barker, Kirk W. Cameron, Amanda Peters Randles, and Guangwen Yang</i>	
BigKernel—High Performance CPU-GPU Communication Pipelining for Big Data-Style Applications .....	819
<i>Reza Mokhtari and Michael Stumm</i>	



DataMPI: Extending MPI to Hadoop-Like Big Data Computing .....	829
<i>Xiaoyi Lu, Fan Liang, Bing Wang, Li Zha, and Zhiwei Xu</i>	

### **Session 21: Network Algorithms**

An Efficient Method for Stream Semantics over RDMA .....	841
<i>Patrick MacArthur and Robert D. Russell</i>	

Collaborative Network Configuration in Hybrid Electrical/Optical Data Center Networks .....	852
<i>Zhiyang Guo and Yuanyuan Yang</i>	

Optimizing Bandwidth Allocation in Flex-Grid Optical Networks with Application to Scheduling .....	862
<i>Hadas Shachnai, Ariella Voloshin, and Shmuel Zaks</i>	

Balancing On-Chip Network Latency in Multi-application Mapping for Chip-Multiprocessors .....	872
<i>Di Zhu, Lizhong Chen, Siyu Yue, Timothy M. Pinkston, and Massoud Pedram</i>	

### **Keynote Speaker 3**

Astrophysical Applications of Machine Learning at Scale and under Duress .....	885
<i>Joshua Bloom</i>	

### **Best Papers Session**

Scalable Single Source Shortest Path Algorithms for Massively Parallel Systems .....	889
<i>Venkatesan T. Chakaravarthy, Fabio Checconi, Fabrizio Petrini, and Yogish Sabharwal</i>	

A New Scalable Parallel Algorithm for Fock Matrix Construction .....	902
<i>Xing Liu, Aftab Patel, and Edmond Chow</i>	

ReDHiP: Recalibrating Deep Hierarchy Prediction for Energy Efficiency .....	915
<i>Xun Li, Diana Franklin, Ricardo Bianchini, and Frederic T. Chong</i>	

F2C2-STM: Flux-Based Feedback-Driven Concurrency Control for STMs .....	927
<i>Kaushik Ravichandran and Santosh Pande</i>	

### **Session 22: Performance Characterization and Optimization**

Identifying Code Phases Using Piece-Wise Linear Regressions .....	941
<i>Harald Servat, Germán Llort, Juan González, Judit Giménez, and Jesús Labarta</i>	

Auto-Tuning Dedispersion for Many-Core Accelerators .....	952
<i>Alessio Sclocco, Henri E. Bal, Jason Hessels, Joeri van Leeuwen, and Rob V. van Nieuwpoort</i>	

RCMP: Enabling Efficient Recomputation Based Failure Resilience for Big Data Analytics .....	962
<i>Florin Dinu and T.S. Eugene Ng</i>	

A Step towards Energy Efficient Computing: Redesigning a Hydrodynamic Application on CPU-GPU .....	972
<i>Tingxing Dong, Veselin Dobrev, Tzanio Kolev, Robert Rieben, Stanimire Tomov, and Jack Dongarra</i>	

### **Session 23: Multithreading and Concurrency**

Using Multiple Threads to Accelerate Single Thread Performance .....	985
<i>Zehra Sura, Kevin O'Brien, and Jose Brunheroto</i>	
Active Measurement of Memory Resource Consumption .....	995
<i>Marc Casas and Greg Bronevetsky</i>	
Locating Parallelization Potential in Object-Oriented Data Structures .....	1005
<i>Korbinian Molitorisz, Thomas Karcher, Alexander Bieleš, and Walter F. Tichy</i>	

### **Session 24: Numerical Algorithms**

An Accelerated Recursive Doubling Algorithm for Block Tridiagonal Systems .....	1019
<i>Sudip K. Seal</i>	
Designing LU-QR Hybrid Solvers for Performance and Stability .....	1029
<i>Mathieu Faverge, Julien Herrmann, Julien Langou, Bradley R. Lowery, Yves Robert, and Jack Dongarra</i>	
Effectively Exploiting Parallel Scale for All Problem Sizes in LU Factorization .....	1039
<i>Md Rakib Hasan and R. Clint Whaley</i>	
Anatomy of High-Performance Many-Threaded Matrix Multiplication .....	1049
<i>Tyler M. Smith, Robert van de Geijn, Mikhail Smelyanskiy, Jeff R. Hammond, and Field G. Van Zee</i>	

### **Session 25: Performance Impacts of Hardware Acceleration**

Comparative Performance Analysis of Intel (R) Xeon Phi (TM), GPU, and CPU: A Case Study from Microscopy Image Analysis .....	1063
<i>George Teodoro, Tahsin Kurc, Jun Kong, Lee Cooper, and Joel Saltz</i>	
A Framework for Lattice QCD Calculations on GPUs .....	1073
<i>F.T. Winter, M.A. Clark, R.G. Edwards, and B. Joó</i>	
Improving Communication Performance and Scalability of Native Applications on Intel Xeon Phi Coprocessor Clusters .....	1083
<i>Karthikeyan Vaidyanathan, Kiran Pamnany, Dhiraj D. Kalamkar, Alexander Heinecke, Mikhail Smelyanskiy, Jongsoo Park, Daehyun Kim, Aniruddha Shet, G. Bharat Kaul, Balint Joo, and Pradeep Dubey</i>	
Computational Co-design of a Multiscale Plasma Application: A Process and Initial Results .....	1093
<i>Joshua Payne, Dana Knoll, Allen McPherson, William Taitano, Luis Chacon, Guangye Chen, and Scott Pakin</i>	

## Session 26: Programming Models and Tools

UPC++: A PGAS Extension for C++ .....	1105
<i>Yili Zheng, Amir Kamil, Michael B. Driscoll, Hongzhang Shan, and Katherine Yelick</i>	
An Evaluation of One-Sided and Two-Sided Communication Paradigms on Relaxed-Ordering Interconnect .....	1115
<i>Khaled Z. Ibrahim, Paul H. Hargrove, Costin Iancu, and Katherine Yelick</i>	
Scaling Irregular Applications through Data Aggregation and Software Multithreading .....	1126
<i>Alessandro Morari, Antonino Tumeo, Daniel Chavarria-Miranda, Oreste Villa, and Mateo Valero</i>	
Generalizing Run-Time Tiling with the Loop Chain Abstraction .....	1136
<i>Michelle Mills Strout, Fabio Luporini, Christopher D. Krieger, Carlo Bertolli, Gheorghe-Teodor Bercea, Catherine Olschanowsky, J. Ramanujam, and Paul H.J. Kelly</i>	

## Session 27: Algorithms for High Performance Computing

s-Step Krylov Subspace Methods as Bottom Solvers for Geometric Multigrid .....	1149
<i>Samuel Williams, Mike Lijewski, Ann Almgren, Brian Van Straalen, Erin Carson, Nicholas Knight, and James Demmel</i>	
Reconstructing Householder Vectors from Tall-Skinny QR .....	1159
<i>Grey Ballard, James Demmel, Laura Grigori, Mathias Jacquelin, Hong Diep Nguyen, and Edgar Solomonik</i>	
Petascale General Solver for Semidefinite Programming Problems with Over Two Million Constraints .....	1171
<i>Katsuki Fujisawa, Toshio Endo, Yuichiro Yasui, Hitoshi Sato, Naoki Matsuzawa, Satoshi Matsuoka, and Hayato Waki</i>	
Optimization of Multi-level Checkpoint Model for Large Scale HPC Applications .....	1181
<i>Sheng Di, Mohamed Slim Bouguerra, Leonardo Bautista-Gomez, and Franck Cappello</i>	

## Session 28: Scalable Algorithms

Evaluating the Impact of SDC on the GMRES Iterative Solver .....	1193
<i>James Elliott, Mark Hoemmen, and Frank Mueller</i>	
A Multi-core Parallel Branch-and-Bound Algorithm Using Factorial Number System .....	1203
<i>Mohand Mez maz, Rudi Leroy, Nouredine Melab, and Daniel Tuytens</i>	
Optimizing Sparse Matrix-Multiple Vectors Multiplication for Nuclear Configuration Interaction Calculations .....	1213
<i>Hasan Metin Aktulga, Aydin Buluç, Samuel Williams, and Chao Yang</i>	

## **Session 29: Resilience and Reliability**

FMI: Fault Tolerant Messaging Interface for Fast and Transparent Recovery .....	1225
<i>Kento Sato, Adam Moody, Kathryn Mohror, Todd Gamblin, Bronis R. de Supinski, Naoya Maruyama, and Satoshi Matsuoka</i>	
Designing Bit-Reproducible Portable High-Performance Applications .....	1235
<i>Andrea Arteaga, Oliver Fuhrer, and Torsten Hoefler</i>	
F-SEFI: A Fine-Grained Soft Error Fault Injection Tool for Profiling Application Vulnerability .....	1245
<i>Qiang Guan, Nathan Debardeleben, Sean Blanchard, and Song Fu</i>	
<b>Author Index</b> .....	1255