

2012 IEEE 48th International Parallel and Distributed Processing Symposium (IPIDPS 2012)

**Shanghai, China
21 – 25 May 2012**

Pages 1-679



**IEEE Catalog Number: CFP12023-PRT
ISBN: 978-1-4673-0975-2**

2012 IEEE 26th International Parallel and Distributed Processing Symposium

IPDPS 2012

Table of Contents

Message from the General Co-Chairs.....	xv
Message from the Program Co-Chairs.....	xvii
Message from the Steering Co-Chairs.....	xix
IPDPS 2012 Organization.....	xx
IPDPS 2012 Technical Program.....	xxiv
IPDPS 2012 Reviewers.....	xxvii

Keynote 1

Large-Scale Visual Data Analysis	1
<i>Chris Johnson</i>	

Session 1: Parallel Linear Algebra Algorithms I

A Predictive Model for Solving Small Linear Algebra Problems in GPU Registers	2
<i>Michael J. Anderson, David Sheffield, and Kurt Keutzer</i>	
A Parallel Tiled Solver for Dense Symmetric Indefinite Systems on Multicore Architectures	14
<i>Marc Baboulin, Dulceneia Becker, and Jack Dongarra</i>	
A Comprehensive Study of Task Coalescing for Selecting Parallelism Granularity in a Two-Stage Bidiagonal Reduction	25
<i>Azzam Haidar, Hatem Ltaief, Piotr Luszczek, and Jack Dongarra</i>	
Improving the Performance of Dynamical Simulations Via Multiple Right-Hand Sides	36
<i>Xing Liu, Edmond Chow, Karthikeyan Vaidyanathan, and Mikhail Smelyanskiy</i>	

Session 2: Bioinformatics and Performance Modeling

High-Performance Interaction-Based Simulation of Gut Immunopathologies with ENteric Immunity Simulator (ENISI)	48
<i>Keith Bisset, Md. Maksudul Alam, Josep Bassaganya-Riera, Adria Carbo, Stephen Eubank, Raquel Hontecillas, Stefan Hoops, Yongguo Mei, Katherine Wendelsdorf, Dawen Xie, Jae-Seung Yeom, and Madhav V. Marathe</i>	
A Parallel Algorithm for Spectrum-based Short Read Error Correction	60
<i>Ankit R. Shah, Sriram Chockalingam, and Srinivas Aluru</i>	

Enhancing the Scalability of Consistency-based Progressive Multiple Sequences Alignment Applications	71
<i>M. Orobotg, F. Cores, F. Guirado, C. Kemena, C. Notredame, and A. Ripoll</i>	
An Accurate GPU Performance Model for Effective Control Flow Divergence Optimization	83
<i>Zheng Cui, Yun Liang, Kyle Rupnow, and Deming Chen</i>	
Session 3: Dynamic Pipeline and Transactional Memory Optimizations	
SEL-TM: Selective Eager-Lazy Management for Improved Concurrency in Transactional Memory	95
<i>Lihang Zhao, Woojin Choi, and Jeff Draper</i>	
Robust SIMD: Dynamically Adapted SIMD Width and Multi-Threading Depth	107
<i>Jiayuan Meng, Jeremy W. Sheaffer, and Kevin Skadron</i>	
Dynamic Operands Insertion for VLIW Architecture with a Reduced Bit-width Instruction Set	119
<i>Jongwon Lee, Jonghee M. Youn, Jihoon Lee, Minwook Ahn, and Yunheung Paek</i>	
SUV: A Novel Single-Update Version-Management Scheme for Hardware Transactional Memory Systems	131
<i>Zhichao Yan, Hong Jiang, Dan Feng, Lei Tian, and Yujuan Tan</i>	
Session 4: Software Scheduling	
Heterogeneous Task Scheduling for Accelerated OpenMP	144
<i>Thomas R.W. Scogland, Barry Rountree, Wu-chun Feng, and Bronis R. de Supinski</i>	
A Source-aware Interrupt Scheduling for Modern Parallel I/O Systems	156
<i>Hongbo Zou, Xian-He Sun, Siyuan Ma, and Xi Duan</i>	
ExPERT: Pareto-Efficient Task Replication on Grids and a Cloud	167
<i>Orna Agmon Ben-Yehuda, Assaf Schuster, Artyom Sharov, Mark Silberstein, and Alexandru Iosup</i>	
Scheduling Closed-Nested Transactions in Distributed Transactional Memory	179
<i>Junwhan Kim and Binoy Ravindran</i>	
Session 5: Multicore Algorithms	
Power-aware Manhattan Routing on Chip Multiprocessors	189
<i>Anne Benoit, Rami Melhem, Paul Renaud-Goud, and Yves Robert</i>	
Efficient Resource Oblivious Algorithms for Multicores with False Sharing	201
<i>Richard Cole and Vijaya Ramachandran</i>	
Competitive Cache Replacement Strategies for Shared Cache Environments	215
<i>Anil Kumar Katti and Vijaya Ramachandran</i>	
A Novel Sorting Algorithm for Many-core Architectures Based on Adaptive Bitonic Sort	227
<i>Hagen Peters, Ole Schulz-Hildebrandt, and Norbert Luttenberger</i>	

Session 6: Scheduling and Load Balancing Algorithms I

Optimizing Busy Time on Parallel Machines	238
<i>George B. Mertzios, Mordechai Shalom, Ariella Voloshin, Prudence W.H. Wong, and Shmuel Zaks</i>	
WATS: Workload-Aware Task Scheduling in Asymmetric Multi-core Architectures	249
<i>Quan Chen, Yawen Chen, Zhiyi Huang, and Minyi Guo</i>	
Parametric Utilization Bounds for Fixed-Priority Multiprocessor Scheduling	261
<i>Nan Guan, Martin Stigge, Wang Yi, and Ge Yu</i>	
Minimizing Weighted Mean Completion Time for Malleable Tasks Scheduling	273
<i>Olivier Beaumont, Nicolas Bonichon, Lionel Eyraud-Dubois, and Loris Marchal</i>	

Session 7: Scientific Applications

Load Balancing of Dynamical Nucleation Theory Monte Carlo Simulations through Resource Sharing Barriers	285
<i>Humayun Arafat, P. Sadayappan, James Dinan, Sriram Krishnamoorthy, and Theresa L. Windus</i>	
Highly Efficient Performance Portable Tracking of Evolving Surfaces	296
<i>Wei Yu, Franz Franchetti, James C. Hoe, and Tsuhan Chen</i>	
Advancing Large Scale Many-Body QMC Simulations on GPU Accelerated Multicore Systems	308
<i>Andres Tomas, Chia-Chen Chang, Richard Scalettar, and Zhaojun Bai</i>	
Reducing Data Movement Costs: Scalable Seismic Imaging on Blue Gene	320
<i>Michael Perrone, Lurng-Kuo Liu, Ligang Lu, Karen Magerlein, Changhoan Kim, Irina Fedulova, and Artyom Semikhin</i>	

Session 8: MPI Debugging and Performance Optimization

Opportunistic Data-driven Execution of Parallel Programs for Efficient I/O Services	330
<i>Xuechen Zhang, Kei Davis, and Song Jiang</i>	
SyncChecker: Detecting Synchronization Errors between MPI Applications and Libraries	342
<i>Zhezhe Chen, Xinyu Li, Jau-Yuan Chen, Hua Zhong, and Feng Qin</i>	
Holistic Debugging of MPI Derived Datatypes	354
<i>Joachim Protze, Tobias Hilbrich, Andreas Knüpfer, Bronis R. de Supinski, and Matthias S. Müller</i>	
Hierarchical Local Storage: Exploiting Flexible User-Data Sharing Between MPI Tasks	366
<i>Marc Tchiboukdjian, Patrick Carribault, and Marc Pérache</i>	

Session 9: Parallel Graph Algorithms I

Fast and Efficient Graph Traversal Algorithm for CPUs: Maximizing Single-Node Efficiency	378
<i>Jatin Chhugani, Nadathur Satish, Changkyu Kim, Jason Sewall, and Pradeep Dubey</i>	
SAHAD: Subgraph Analysis in Massive Networks Using Hadoop	390
<i>Zhao Zhao, Guanying Wang, Ali R. Butt, Maleq Khan, V.S. Anil Kumar, and Madhav V. Marathe</i>	
Accelerating Nearest Neighbor Search on Manycore Systems	402
<i>Lawrence Cayton</i>	
Optimizing Large-scale Graph Analysis on Multithreaded, Multicore Platforms	414
<i>Guojing Cong and Konstantin Makarychev</i>	

Session 10: High Performance Computing Algorithms

Low-Cost Parallel Algorithms for 2:1 Octree Balance	426
<i>Tobin Isaac, Carsten Burstedde, and Omar Ghattas</i>	
A Case Study of Designing Efficient Algorithm-based Fault Tolerant Application for Exascale Parallelism	438
<i>Erlin Yao, Rui Wang, Mingyu Chen, Guangming Tan, and Ninghui Sun</i>	
High Performance Non-uniform FFT on Modern X86-based Multi-core Systems	449
<i>Dhiraj D. Kalamkar, Joshua D. Trzaskoz, Srinivas Sridharan, Mikhail Smelyanskiy, Daehyun Kim, Armando Manduca, Yunhong Shu, Matt A. Bernstein, Bharat Kaul, and Pradeep Dubey</i>	
NUMA Aware Iterative Stencil Computations on Many-Core Systems	461
<i>Mohammed Shaheen and Robert Strzodka</i>	

Session 11: Parallel Numerical Computation

Algebraic Block Multi-Color Ordering Method for Parallel Multi-Threaded Sparse Triangular Solver in ICCG Method	474
<i>Takeshi Iwashita, Hiroshi Nakashima, and Yasuhito Takahashi</i>	
The Parallel Computation of Morse-Smale Complexes	484
<i>Attila Gyulassy, Valerio Pascucci, Tom Peterka, and Robert Ross</i>	
Hybrid Static/dynamic Scheduling for Already Optimized Dense Matrix Factorization	496
<i>Simplice Donfack, Laura Grigori, William D. Gropp, and Vivek Kale</i>	

Session 12: Architecture Modeling and Scheduling

Understanding Cache Hierarchy Contention in CMPs to Improve Job Scheduling	508
<i>Josué Feliu, Julio Sahuquillo, Salvador Petit, and José Duato</i>	
Optimization of Parallel Discrete Event Simulator for Multi-core Systems	520
<i>Deepak Jagtap, Nael Abu-Ghazaleh, and Dmitry Ponomarev</i>	

Using the Translation Lookaside Buffer to Map Threads in Parallel Applications Based on Shared Memory	532
<i>Eduardo H.M. Cruz, Matthias Diener, and Philippe O.A. Navaux</i>	

Session 13: GPU-Based Computing

Automatic Resource Scheduling with Latency Hiding for Parallel Stencil Applications on GPGPU Clusters	544
<i>Kumiko Maeda, Masana Murase, Munehiro Doi, Hideaki Komatsu, Shigeho Noda, and Ryutaro Himeno</i>	
Productive Programming of GPU Clusters with OmpSs	557
<i>Javier Bueno, Judit Planas, Alejandro Duran, Rosa M. Badia, Xavier Martorell, Eduard Ayguadé, and Jesús Labarta</i>	
Generating Device-specific GPU Code for Local Operators in Medical Imaging	569
<i>Richard Membarth, Frank Hannig, Jürgen Teich, Mario Körner, and Wieland Eckert</i>	
Performance Portability with the Chapel Language	582
<i>Albert Sidelnik, Saeed Maleki, Bradford L. Chamberlain, María J. Garzarán, and David Padua</i>	

Keynote 2

Exascale System Software for the Year of the Dragon	595
<i>Pete Beckman</i>	

Session 14: Parallel Matrix Factorizations

Mapping Dense LU Factorization on Multicore Supercomputer Nodes	596
<i>Jonathan Lifflander, Phil Miller, Ramprasad Venkataraman, Anshu Arya, Laxmikant Kale, and Terry Jones</i>	
Hierarchical QR Factorization Algorithms for Multi-core Cluster Systems	607
<i>Jack Dongarra, Mathieu Faverge, Thomas Herault, Julien Langou, and Yves Robert</i>	
New Scheduling Strategies and Hybrid Programming for a Parallel Right-looking Sparse LU Factorization Algorithm on Multicore Cluster Systems	619
<i>Ichitaro Yamazaki and Xiaoye S. Li</i>	
ShyLU: A Hybrid-Hybrid Solver for Multicore Platforms	631
<i>Sivasankaran Rajamanickam, Erik G. Boman, and Michael A. Heroux</i>	

Session 15: Distributed Computing and Programming Models

MATE-CG: A Map Reduce-Like Framework for Accelerating Data-Intensive Computations on Heterogeneous Clusters	644
<i>Wei Jiang and Gagan Agrawal</i>	
Automated and Agile Server Parameter Tuning with Learning and Control	656
<i>Yanfei Guo, Palden Lama, and Xiaobo Zhou</i>	

A Self-tuning Failure Detection Scheme for Cloud Computing Service	668
<i>Naixue Xiong, Athanasios V. Vasilakos, Jie Wu, Y. Richard Yang, Andy Rindos, Yuezhi Zhou, Wen-Zhan Song, and Yi Pan</i>	
PGAS for Distributed Numerical Python Targeting Multi-core Clusters	680
<i>Mads Ruben Burgdorff Kristensen, Yili Zheng, and Brian Vinter</i>	
Session 16: Memory Architectures	
Miss-Correlation Folding: Encoding Per-Block Miss Correlations in Compressed DRAM for Data Prefetching	691
<i>Gang Liu, Jih-Kwon Peir, and Victor Lee</i>	
On the Role of NVRAM in Data-intensive Architectures: An Evaluation	703
<i>Brian Van Essen, Roger Pearce, Sasha Ames, and Maya Gokhale</i>	
iTransformer: Using SSD to Improve Disk Scheduling for High-performance I/O	715
<i>Xuechen Zhang, Kei Davis, and Song Jiang</i>	
Switching Optically-Connected Memories in a Large-Scale System	727
<i>Abhirup Chakraborty, Eugen Schenfeld, and Dilma Da Silva</i>	
Session 17: High Performance Communication and Networking	
Supporting the Global Arrays PGAS Model Using MPI One-Sided Communication	739
<i>James Dinan, Pavan Balaji, Jeff R. Hammond, Sriram Krishnamoorthy, and Vinod Tipparaju</i>	
A uGNI-based Asynchronous Message-driven Runtime System for Cray Supercomputers with Gemini Interconnect	751
<i>Yanhua Sun, Gengbin Zheng, Laximant V. Kalé, Terry R. Jones, and Ryan Olson</i>	
PAMI: A Parallel Active Message Interface for the Blue Gene/Q Supercomputer	763
<i>Sameer Kumar, Amith R. Mamidala, Daniel A. Faraj, Brian Smith, Michael Blocksome, Bob Cernohous, Douglas Miller, Jeff Parker, Joseph Ratterman, Philip Heidelberger, Dong Chen, and Burkhard Steinmacher-Burrow</i>	
High-Performance Design of HBase with RDMA over InfiniBand	774
<i>Jian Huang, Xiangyong Ouyang, Jithin Jose, Md. Wasi-ur-Rahman, Hao Wang, Miao Luo, Hari Subramoni, Chet Murthy, and Dhabaleswar K. Panda</i>	
Session 18: Scheduling and Load Balancing Algorithms II	
Virtual Machine Resource Allocation for Service Hosting on Heterogeneous Distributed Platforms	786
<i>Mark Stillwell, Frédéric Vivien, and Henri Casanova</i>	
Consistency-aware Partitioning Algorithm in Multi-server Distributed Virtual Environments	798
<i>Yusen Li and Wentong Cai</i>	
Optimal Resource Rental Planning for Elastic Applications in Cloud Market	808
<i>Han Zhao, Miao Pan, Xinxin Liu, Xiaolin Li, and Yuguang Fang</i>	

Improved Bounds for Discrete Diffusive Load Balancing	820
<i>Clemens P.J. Adolphs and Petra Berenbrink</i>	

Session 19: Parallel Graph Algorithms II

Multi-core Spanning Forest Algorithms using the Disjoint-set Data Structure	827
<i>Md. Mostofa Ali Patwary, Peder Refsnes, and Fredrik Manne</i>	
Graph Partitioning for Reconfigurable Topology	836
<i>Deepak Ajwani, Shoukat Ali, and John P. Morrison</i>	
Multithreaded Clustering for Multi-level Hypergraph Partitioning	848
<i>Ümit V. Çatalyürek, Mehmet Deveci, Kamer Kaya, and Bora Uçar</i>	
Multithreaded Algorithms for Maximum Matching in Bipartite Graphs	860
<i>Ariful Azad, Mahantesh Halappanavar, Sivasankaran Rajamanickam, Erik G. Boman, Arif Khan, and Alex Pothén</i>	

Session 20: Data Intensive and Peer-to-Peer Computing

Multi-level Layout Optimization for Efficient Spatio-temporal Queries on ISABELA-compressed Data	873
<i>Zhenhuan Gong, Sriram Lakshminarasimhan, John Jenkins, Hemanth Kolla, Stephane Ethier, Jackie Chen, Robert Ross, Scott Klasky, and Nagiza F. Samatova</i>	
Evaluating Mesh-based P2P Video-on-Demand Systems	885
<i>Yingwu Zhu</i>	
Query Optimization and Execution in a Parallel Analytics DBMS	897
<i>Todd Eavis and Ahmad Taleb</i>	
Dynamic Message Ordering for Topic-Based Publish/Subscribe Systems	909
<i>Roberto Baldoni, Silvia Bonomi, Marco Platania, and Leonardo Querzoni</i>	

Session 21: Disk and Memory Software Optimization

iHarmonizer: Improving the Disk Efficiency of I/O-intensive Multithreaded Codes	921
<i>Yizhe Wang, Kei Davis, Yuehai Xu, and Song Jiang</i>	
Improving Parallel IO Performance of Cell-based AMR Cosmology Applications	933
<i>Yongen Yu, Douglas H. Rudd, Zhiling Lan, Nickolay Y. Gnedin, Andrey Kravtsov, and Jingjin Wu</i>	
Identifying Opportunities for Byte-Addressable Non-Volatile Memory in Extreme-Scale Scientific Applications	945
<i>Dong Li, Jeffrey S. Vetter, Gabriel Marin, Collin McCurdy, Cristian Cira, Zhuo Liu, and Weikuan Yu</i>	
NVMalloc: Exposing an Aggregate SSD Store as a Memory Partition in Extreme-Scale Machines	957
<i>Chao Wang, Sudharshan S. Vazhkudai, Xiaosong Ma, Fei Meng, Youngjae Kim, and Christian Engelmann</i>	

Keynote 3

Building Billion-Threads Computer and Elastic Processor	969
<i>Guo-Jie Li</i>	

Best Papers

HierKNEM: An Adaptive Framework for Kernel-Assisted and Topology-Aware Collective Communications on Many-core Clusters	970
<i>Teng Ma, George Bosilca, Aurelien Bouteiller, and Jack Dongarra</i>	
BRISA: Combining Efficiency and Reliability in Epidemic Data Dissemination	983
<i>Miguel Matos, Valerio Schiavoni, Pascal Felber, Rui Oliveira, and Etienne Riviere</i>	
Locality Principle Revisited: A Probability-Based Quantitative Approach	995
<i>Saurabh Gupta, Ping Xiang, Yi Yang, and Huiyang Zhou</i>	
Evaluating the Impact of TLB Misses on Future HPC Systems	1010
<i>Alessandro Morari, Roberto Gioiosa, Robert W. Wisniewski, Bryan S. Rosenburg, Todd A. Inglett, and Mateo Valero</i>	

Session 22: Network Algorithms

Optimal Algorithms and Approximation Algorithms for Replica Placement with Distance Constraints in Tree Networks	1022
<i>A. Benoit, H. Larchevêque, and P. Renaud-Goud</i>	
On Nonblocking Multirate Multicast Fat-tree Data Center Networks with Server Redundancy	1034
<i>Zhiyang Guo and Yuanyuan Yang</i>	
Distributed Transactional Memory for General Networks	1045
<i>Gokarna Sharma, Costas Busch, and Srivathsan Srinivasagopalan</i>	
On λ -Alert Problem	1057
<i>Marek Klonowski and Dominik Pajak</i>	

Session 23: GPU Acceleration

Efficient Quality Threshold Clustering for Parallel Architectures	1068
<i>Anthony Danalis, Collin McCurdy, and Jeffrey S. Vetter</i>	
A Highly Parallel Reuse Distance Analysis Algorithm on GPUs	1080
<i>Huimin Cui, Qing Yi, Jingling Xue, Lei Wang, Yang Yang, and Xiaobing Feng</i>	
Accelerating Large Scale Image Analyses on Parallel, CPU-GPU Equipped Systems	1093
<i>George Teodoro, Tahsin M. Kurck, Tony Pan, Lee A.D. Cooper, Jun Kong, Patrick Widener, and Joel H. Saltz</i>	
Radio Astronomy Beam Forming on Many-Core Architectures	1105
<i>Alessio Sclocco, Ana Lucia Varbanescu, Jan David Mol, and Rob V. van Nieuwpoort</i>	

Session 24: Interconnection Networks

Cross-layer Energy and Performance Evaluation of a Nanophotonic Manycore Processor System Using Real Application Workloads	1117
<i>George Kurian, Chen Sun, Chia-Hsin Owen Chen, Jason E. Miller, Jurgen Michel, Lan Wei, Dimitri A. Antoniadis, Li-Shiuan Peh, Lionel Kimerling, Vladimir Stojanovic, and Anant Agarwal</i>	
Exploring the Scope of the InfiniBand Congestion Control Mechanism	1131
<i>Ernst Gunnar Gran, Sven-Arne Reinemo, Olav Lysne, Tor Skeie, Eitan Zahavi, and Gilad Shainer</i>	
DCAF - A Directly Connected Arbitration-Free Photonic Crossbar for Energy-Efficient High Performance Computing	1144
<i>Christopher Nitta, Matthew Farrens, and Venkatesh Akella</i>	
Designing Non-blocking Allreduce with Collective Offload on InfiniBand Clusters: A Case Study with Conjugate Gradient Solvers	1156
<i>K. Kandalla, U. Yang, J. Keasler, T. Kolev, A. Moody, H. Subramoni, K. Tomko, J. Vienne, Bronis R. de Supinski, and Dhabaleswar K. Panda</i>	

Session 25: Software Reliability

Taming of the Shrew: Modeling the Normal and Faulty Behaviour of Large-scale HPC Systems	1168
<i>Ana Gainaru, Franck Cappello, and William Kramer</i>	
Meteor Shower: A Reliable Stream Processing System for Commodity Data Centers	1180
<i>Huayong Wang, Li-Shiuan Peh, Emmanouil Koukourmidis, Shao Tao, and Mun Choon Chan</i>	
Hybrid Transactions: Lock Allocation and Assignment for Irrevocability	1192
<i>Jaswanth Sreeram and Santosh Pande</i>	
Profiling-based Adaptive Contention Management for Software Transactional Memory	1204
<i>Zhengyu He, Xiao Yu, and Bo Hong</i>	

Session 26: Communication Protocols and Benchmarking Algorithms

HydEE: Failure Containment without Event Logging for Large Scale Send-Deterministic MPI Applications	1216
<i>Amina Guermouche, Thomas Ropars, Marc Snir, and Franck Cappello</i>	
Distributed Demand and Response Algorithm for Optimizing Social-Welfare in Smart Grid	1228
<i>Qifen Dong, Li Yu, Wen-Zhan Song, Lang Tong, and Shaojie Tang</i>	
Scalable Distributed Consensus to Support MPI Fault Tolerance	1240
<i>Darius Buntinas</i>	
ScalaBenchGen: Auto-Generation of Communication Benchmarks Traces	1250
<i>Xing Wu, Vivek Deshpande, and Frank Mueller</i>	

Session 27: Parallel Algorithms

A Self-Stabilization Process for Small-World Networks	1261
<i>Sebastian Kniesburges, Andreas Koutsopoulos, and Christian Scheideler</i>	
Self-organizing Particle Systems	1272
<i>Maximilian Drees, Martina Hüllmann, Andreas Koutsopoulos, and Christian Scheideler</i>	
PARDA: A Fast Parallel Reuse Distance Analysis Algorithm	1284
<i>Qingpeng Niu, James Dinan, Qingda Lu, and P. Sadayappan</i>	
A Lower Bound on Proximity Preservation by Space Filling Curves	1295
<i>Pan Xu and Srikanta Tirthapura</i>	

Session 28: Software Performance Analysis and Optimization

Modeling and Analyzing Key Performance Factors of Shared Memory MapReduce	1306
<i>Devesh Tiwari and Yan Solihin</i>	
Predicting Potential Speedup of Serial Code via Lightweight Profiling and Emulations with Memory Performance Model	1318
<i>Minjang Kim, Pranith Kumar, Hyesoon Kim, and Bevin Brett</i>	
Scalable Critical-Path Based Performance Analysis	1330
<i>David Böhme, Felix Wolf, Bronis R. de Supinski, Martin Schulz, and Markus Geimer</i>	
FractalMRC: Online Cache Miss Rate Curve Prediction on Commodity Systems	1341
<i>Lulu He, Zhibin Yu, and Hai Jin</i>	

Session 29: Performance Optimization Frameworks and Methods

Enabling In-situ Execution of Coupled Scientific Workflow on Multi-core Platform	1352
<i>Fan Zhang, Ciprian Docan, Manish Parashar, Scott Klasky, Norbert Podhorszki, and Hasan Abbasi</i>	
GTI: A Generic Tools Infrastructure for Event-Based Tools in Parallel Systems	1364
<i>Tobias Hilbrich, Matthias S. Müller, Bronis R. de Supinski, Martin Schulz, and Wolfgang E. Nagel</i>	
An Efficient Framework for Multi-dimensional Tuning of High Performance Computing Applications	1376
<i>Guojing Cong, Huifang Wen, I-hsin Chung, David Klepacki, Hiroki Murata, and Yasushi Negishi</i>	
An SMT-Selection Metric to Improve Multithreaded Applications' Performance	1388
<i>Justin R. Funston, Kaoutar El Maghraoui, Joefon Jann, Pratap Pattnaik, and Alexandra Fedorova</i>	

Author Index