

2015 IEEE International Parallel and Distributed Processing Symposium

(IPDPS 2015)

**Hyderabad, India
25-29 May 2015**

Pages 1-590



IEEE Catalog Number: CFP15023-POD
ISBN: 978-1-4799-8650-7

2015 IEEE 29th International Parallel and Distributed Processing Symposium

IPDPS 2015

Table of Contents

| | |
|---|-------|
| Message from the General Co-Chairs | xvi |
| Message from the Program Chair..... | xviii |
| Message from the Steering Co-Chairs..... | xx |
| Message from the Workshops Chair..... | xxi |
| IPDPS 2015 Organization..... | xxiii |
| IPDPS 2015 Reviewers..... | xxx |
| IPDPS 2015 Technical Program | xxxi |

Keynote 1

| | |
|--|---|
| Big Data: Scale Down, Scale Up, Scale Out | 3 |
| <i>Phillip B. Gibbons</i> | |

Session 1: Graph and Social Analytics

| | |
|---|----|
| Balanced Coloring for Parallel Computing Applications | 7 |
| <i>Hao Lu, Mahantesh Halappanavar, Daniel Chavarría-Miranda, Assefaw Gebremedhin, and Ananth Kalyanaraman</i> | |
| High-Performance Graph Analytics on Manycore Processors | 17 |
| <i>George M. Slota, Sivasankaran Rajamanickam, and Kamesh Madduri</i> | |
| Scalable Community Detection with the Louvain Algorithm | 28 |
| <i>Xinyu Que, Fabio Checconi, Fabrizio Petrini, and John A. Gunnels</i> | |
| Cooperative Computing for Autonomous Data Centers | 38 |
| <i>Jonathan Berry, Michael Collins, Aaron Kearns, Cynthia A. Phillips, Jared Saia, and Randy Smith</i> | |

Session 2: Numerical Linear Algebra

| | |
|--|----|
| Divide and Conquer Symmetric Tridiagonal Eigensolver for Multicore Architectures | 51 |
| <i>Grégoire Pichon, Azzam Haidar, Mathieu Faverge, and Jakub Kurzak</i> | |
| SPLATT: Efficient and Parallel Sparse Tensor-Matrix Multiplication | 61 |
| <i>Shaden Smith, Niranjay Ravindran, Nicholas D. Sidiropoulos, and George Karypis</i> | |
| A Sparse Direct Solver for Distributed Memory Xeon Phi-Accelerated Systems | 71 |
| <i>Piyush Sao, Xing Liu, Richard Vuduc, and Xiaoye Li</i> | |
| Malleable Sorting | 82 |
| <i>Patrick Flick, Peter Sanders, and Jochen Speck</i> | |

Session 3: High Performance Networks and Congestion Management

| | |
|--|-----|
| GASOLIN: Global Arbitration for Streams of Data in Optical Links | 93 |
| <i>Jiwei Liu, Jun Yang, and Rami Melhem</i> | |
| Contention-Based Nonminimal Adaptive Routing in High-Radix Networks | 103 |
| <i>Pablo Fuentes, Enrique Vallejo, Marina García, Ramón Beivide, Germán Rodríguez, Cyriel Minkenberg, and Mateo Valero</i> | |
| Identifying the Culprits Behind Network Congestion | 113 |
| <i>Abhinav Bhatele, Andrew R. Titus, Jayaraman J. Thiagarajan, Nikhil Jain, Todd Gamblin, Peer-Timo Bremer, Martin Schulz, and Laxmikant V. Kale</i> | |
| Embedding Nonblocking Multicast Virtual Networks in Fat-Tree Data Centers | 123 |
| <i>Jun Duan, Zhiyang Guo, and Yuanyuan Yang</i> | |

Session 4: Software for Heterogeneous Many-Core Systems

| | |
|---|-----|
| Cashmere: Heterogeneous Many-Core Computing | 135 |
| <i>Pieter Hijma, Ceriel J.H. Jacobs, Rob V. van Nieuwpoort, and Henri E. Bal</i> | |
| A Scheduling and Runtime Framework for a Cluster of Heterogeneous Machines with Multiple Accelerators | 146 |
| <i>Tarun Beri, Sorav Bansal, and Subodh Kumar</i> | |
| Hierarchical DAG Scheduling for Hybrid Distributed Systems | 156 |
| <i>Wei Wu, Aurelien Bouteiller, George Bosilca, Mathieu Faverge, and Jack Dongarra</i> | |
| Pushing the Performance Envelope of Modular Exponentiation Across Multiple Generations of GPUs | 166 |
| <i>Niall Emmart and Charles Weems</i> | |

Session 5: Scheduling Algorithms

| | |
|---|-----|
| Federated Scheduling of Sporadic DAG Task Systems | 179 |
| <i>Sanjoy Baruah</i> | |
| Addressing Fairness in SMT Multicores with a Progress-Aware Scheduler | 187 |
| <i>Josué Feliu, Julio Sahuquillo, Salvador Petit, and José Duato</i> | |
| Fast and High Quality Topology-Aware Task Mapping | 197 |
| <i>Mehmet Deveci, Kamer Kaya, Bora Uçar, and Ümit V. Çatalyürek</i> | |
| Workload-Driven VM Consolidation in Cloud Data Centers | 207 |
| <i>Hao Lin, Xin Qi, Shuo Yang, and Samuel Midkiff</i> | |

Session 6: Concurrency in Memory Systems

| | |
|--|-----|
| Update Consistency for Wait-Free Concurrent Objects | 219 |
| <i>Matthieu Perrin, Achour Mostefaoui, and Claude Jard</i> | |
| Modeling Energy Consumption of Lock-Free Queue Implementations | 229 |
| <i>Aras Atalar, Anders Gidenstam, Paul Renaud-Goud, and Philippas Tsigas</i> | |
| A Consistency Framework for Iteration Operations in Concurrent Data Structures | 239 |
| <i>Yiannis Nikolakopoulos, Anders Gidenstam, Marina Papatriantafilou, and Philippas Tsigas</i> | |
| An Automated Framework for Decomposing Memory Transactions to Exploit Partial Rollback | 249 |
| <i>Aditya Dhone, Roberto Palmieri, and Binoy Ravindran</i> | |

Session 7: MapReduce Advances

| | |
|---|-----|
| Cracking Down MapReduce Failure Amplification through Analytics Logging and Migration | 261 |
| <i>Yandong Wang, Huansong Fu, and Weikuan Yu</i> | |
| Grouping Blocks for MapReduce Co-Locality | 271 |
| <i>Xiao Yu and Bo Hong</i> | |
| SMapReduce: Optimising Resource Allocation by Managing Working Slots at Runtime | 281 |
| <i>Feng Liang and Francis C.M. Lau</i> | |
| High-Performance Design of YARN MapReduce on Modern HPC Clusters with Lustre and RDMA | 291 |
| <i>Md. Wasi-ur-Rahman, Xiaoyi Lu, Nusrat Sharmin Islam, Raghunath Rajachandrasekar, and Dhabaleswar K. (DK) Panda</i> | |

Session 8: Performance and Energy Optimizations

| | |
|---|-----|
| High-Performance Energy-Efficient Recursive Dynamic Programming with Matrix-Multiplication-Like Flexible Kernels | 303 |
| <i>Jesmin Jahan Tithi, Pramod Ganapathi, Aakrati Talati, Sonal Aggarwal, and Rezaul Chowdhury</i> | |
| Compiler-Directed Transformation for Higher-Order Stencils | 313 |
| <i>Protonu Basu, Mary Hall, Samuel Williams, Brian Van Straalen, Leonid Oliker, and Phillip Colella</i> | |
| LUC: Limiting the Unintended Consequences of Power Scaling on Parallel Transaction-Oriented Workloads | 324 |
| <i>Hung-Ching Chang, Bo Li, Godmar Back, Ali R. Butt, and Kirk W. Cameron</i> | |
| PowerFCT: Power Optimization of Data Center Network with Flow Completion Time Constraints | 334 |
| <i>Kuangyu Zheng, Xiaodong Wang, and Xiaorui Wang</i> | |

Session 9: Dynamic Networks

| | |
|---|-----|
| Leader Election in Sparse Dynamic Networks with Churn | 347 |
| <i>John Augustine, Tejas Kulkarni, and Sumathi Sivasubramaniam</i> | |
| Online Top-k-Position Monitoring of Distributed Data Streams | 357 |
| <i>Alexander Mäcker, Manuel Malatyali, and Friedhelm Meyer auf der Heide</i> | |
| DSLR: A Distributed Schedule Length Reduction Algorithm for WSNs | 365 |
| <i>Ashutosh Bhatia and R.C. Hansdah</i> | |
| Logarithmic-Time Complete Visibility for Robots with Lights | 375 |
| <i>Ramachandran Vaidyanathan, Costas Busch, Jerry L. Trahan, Gokarna Sharma, and Suresh Rai</i> | |

Session 10: Applications on GPUs

| | |
|--|-----|
| Indexing of Spatiotemporal Trajectories for Efficient Distance Threshold Similarity Searches on the GPU | 387 |
| <i>Michael Gowanlock and Henri Casanova</i> | |
| Efficient Selection Algorithm for Fast k-NN Search on GPUs | 397 |
| <i>Xiaoxin Tang, Zhiyi Huang, David Evers, Steven Mills, and Minyi Guo</i> | |
| Optimizing Sparse Matrix Operations on GPUs Using Merge Path | 407 |
| <i>Steven Dalton, Sean Baxter, Duane Merrill, Luke Olson, and Michael Garland</i> | |
| Performance Engineering of the Kernel Polynomial Method on Large-Scale CPU-GPU Systems | 417 |
| <i>Moritz Kreutzer, Andreas Pieper, Georg Hager, Gerhard Wellein, Andreas Alvermann, and Holger Fehske</i> | |

Session 11: Scheduling on Clusters

| | |
|--|-----|
| A Batch System with Efficient Adaptive Scheduling for Malleable and Evolving Applications | 429 |
| <i>Suraj Prabhakaran, Marcel Neumann, Sebastian Rinke, Felix Wolf, Abhishek Gupta, and Laxmikant V. Kale</i> | |
| Improving Batch Scheduling on Blue Gene/Q by Relaxing 5D Torus Network Allocation Constraints | 439 |
| <i>Zhou Zhou, Xu Yang, Zhiling Lan, Paul Rich, Wei Tang, Vitali Morozov, and Narayan Desai</i> | |
| Quiet Neighborhoods: Key to Protect Job Performance Predictability | 449 |
| <i>Ana Jokanovic, Jose Carlos Sancho, German Rodriguez, Alejandro Lucero, Cyriel Minkenberg, and Jesus Labarta</i> | |
| Stratified Sampling for Even Workload Partitioning Applied to IDA* and Delaunay Algorithms | 460 |
| <i>Jeeva Paudel, Levi H. S. Lelis, and José Nelson Amaral</i> | |

Session 12: Debugging and Verification

| | |
|---|-----|
| A Scalable Prescriptive Parallel Debugging Model | 473 |
| <i>Nicklas Bo Jensen, Niklas Quarfat Nielsen, Gregory L. Lee, Sven Karlsson, Matthew Legendre, Martin Schulz, and Dong H. Ahn</i> | |
| An Efficient Data-Dependence Profiler for Sequential and Parallel Programs | 484 |
| <i>Zhen Li, Ali Jannesari, and Felix Wolf</i> | |
| Decentralized Runtime Verification of LTL Specifications in Distributed Systems | 494 |
| <i>Menna Mostafa and Borzoo Bonakdarpour</i> | |
| Fast Proof Generation for Verifying Cloud Search | 504 |
| <i>Jingyu Zhou, Jiannong Cao, Bin Yao, and Minyi Guo</i> | |

Keynote 2

| | |
|---|-----|
| Julia: A Fresh Approach to Parallel Programming | 517 |
| <i>Alan Edelman</i> | |

Session 13: Randomized Algorithms

| | |
|---|-----|
| On the Influence of Graph Density on Randomized Gossiping | 521 |
| <i>Robert Elsässer and Dominik Kaaser</i> | |
| Distinct Random Sampling from a Distributed Stream | 532 |
| <i>Srikanta Tirthapura</i> | |

| | |
|---|-----|
| Randomized Renaming in Shared Memory Systems | 542 |
| <i>Petra Berenbrink, André Brinkmann, Robert Elsässer, Tom Friedetzky, and Lars Nagel</i> | |
| Threshold Load Balancing with Weighted Tasks | 550 |
| <i>Petra Berenbrink, Tom Friedetzky, Frederik Mallmann-Trenn, Sepehr Meshkinfamfar, and Chris Wastell</i> | |

Session 14: Scientific Applications I

| | |
|---|-----|
| merAligner: A Fully Parallel Sequence Aligner | 561 |
| <i>Evangelos Georganas, Aydin Buluç, Jarrod Chapman, Leonid Oliker, Daniel Rokhsar, and Katherine Yelick</i> | |
| An Algebraic Parallel Treecode in Arbitrary Dimensions | 571 |
| <i>William B. March, Bo Xiao, Chenhan D. Yu, and George Biros</i> | |
| 3D Cartesian Transport Sweep for Massively Parallel Architectures with PaRSEC | 581 |
| <i>Salli Moustafa, Mathieu Faverge, Laurent Pagne, and Pierre Ramet</i> | |
| A Pattern Specification and Optimizations Framework for Accelerating Scientific Computations on Heterogeneous Clusters | 591 |
| <i>Linchuan Chen, Xin Huo, and Gagan Agrawal</i> | |

Session 15: Storage Systems Architecture

| | |
|--|-----|
| D-Code: An Efficient RAID-6 Code to Optimize I/O Loads and Read Performance | 603 |
| <i>Yingxun Fu and Jiwu Shu</i> | |
| HAS: Heterogeneity-Aware Selective Data Layout Scheme for Parallel File Systems on Hybrid Servers | 613 |
| <i>Shuibing He, Xian-He Sun, and Adnan Haider</i> | |
| Opass: Analysis and Optimization of Parallel Data Access on Distributed File Systems | 623 |
| <i>Jiangling Yin, Jun Wang, Jian Zhou, Tyler Lukasiewicz, Dan Huang, and Junyao Zhang</i> | |
| Improving Storage Availability in Cloud-of-Clouds with Hybrid Redundant Data Distribution | 633 |
| <i>Bo Mao, Suzhen Wu, and Hong Jiang</i> | |

Session 16: MPI and Charm++ Advances

| | |
|--|-----|
| Efficient Process Replication for MPI Applications: Sharing Work between Replicas | 645 |
| <i>Thomas Ropars, Arnaud Lefray, Dohyun Kim, and André Schiper</i> | |
| Charm++ and MPI: Combining the Best of Both Worlds | 655 |
| <i>Nikhil Jain, Abhinav Bhatele, Jae-Seung Yeom, Mark F. Adams, Francesco Miniati, Chao Mei, and Laxmikant V. Kale</i> | |
| Casper: An Asynchronous Progress Model for MPI RMA on Many-Core Architectures | 665 |
| <i>Min Si, Antonio J. Peña, Jeff Hammond, Pavan Balaji, Masamichi Takagi, and Yutaka Ishikawa</i> | |
| Scalable Asynchronous Contact Mechanics Using Charm++ | 677 |
| <i>Xiang Ni, Laxmikant V. Kale, and Rasmus Tamstorf</i> | |

Session 17: Combinatorial Algorithms and Optimization

| | |
|--|-----|
| Association Rule Mining with the Micron Automata Processor | 689 |
| <i>Ke Wang, Yanjun Qi, Jeffrey J. Fox, Mircea R. Stan, and Kevin Skadron</i> | |
| Cichlid: Efficient Large Scale RDFS/OWL Reasoning with Spark | 700 |
| <i>Rong Gu, Shanyong Wang, Fangfang Wang, Chunfeng Yuan, and Yihua Huang</i> | |
| Parallel Strategies for Solving Large Unit Commitment Problems in the California ISO Planning Model | 710 |
| <i>Guojing Cong, Carol Meyers, Deepak Rajan, and Tiziano Parriani</i> | |

Session 18: Scientific Applications II

| | |
|--|-----|
| Exploring Shared-Memory Optimizations for an Unstructured Mesh CFD Application on Modern Parallel Systems | 723 |
| <i>Dheevatsa Mudigere, Srinivas Sridharan, Anand Deshpande, Jongsoo Park, Alexander Heinecke, Mikhail Smelyanskiy, Bharat Kaul, Pradeep Dubey, Dinesh Kaushik, and David Keyes</i> | |
| A Performance Analysis of SIMD Algorithms for Monte Carlo Simulations of Nuclear Reactor Cores | 733 |
| <i>David Ozog, Allen D. Malony, and Andrew R. Siegel</i> | |
| Generating Optimized Fourier Interpolation Routines for Density Functional Theory Using SPIRAL | 743 |
| <i>Doru Thom Popovici, Francis P. Russel, Karl Wilkinson, Chris-Kriton Skylaris, Paul H. J. Kelly, and Franz Franchetti</i> | |
| Parallel Hessian Assembly for Seismic Waveform Inversion Using Global Updates | 753 |
| <i>Scott French, Yili Zheng, Barbara Romanowicz, and Katherine Yelick</i> | |

Session 19: Resilience

| | |
|--|-----|
| Design for a Soft Error Resilient Dynamic Task-Based Runtime | 765 |
| <i>Chongxiao Cao, Thomas Herault, George Bosilca, and Jack Dongarra</i> | |
| Recovering from Overload in Multicore Mixed-Criticality Systems | 775 |
| <i>Jeremy P. Erickson, Namhoon Kim, and James H. Anderson</i> | |
| Investigating the Interplay between Energy Efficiency and Resilience in High Performance Computing | 786 |
| <i>Li Tan, Shuaiwen Leon Song, Panruo Wu, Zizhong Chen, Rong Ge, and Darren J. Kerbyson</i> | |

Session 20: Graph Analytics

| | |
|---|-----|
| A Hybrid Approach to Processing Big Data Graphs on Memory-Restricted Systems | 799 |
| <i>Harshvardhan, Brandon West, Adam Fidel, Nancy M. Amato, and Lawrence Rauchwerger</i> | |
| Distributed Programming over Time-Series Graphs | 809 |
| <i>Yogesh Simmhan, Neel Choudhury, Charith Wickramaarachchi, Alok Kumbhare, Marc Frincu, Cauligi Raghavendra, and Viktor Prasanna</i> | |
| Efficient and Simplified Parallel Graph Processing over CPU and MIC | 819 |
| <i>Linchuan Chen, Xin Huo, Bin Ren, Surabhi Jain, and Gagan Agrawal</i> | |

Keynote 3

| | |
|--|-----|
| Assisting H1N1 and Ebola Outbreak Response through High Performance Networked Epidemiology | 831 |
| <i>Madhav Marathe</i> | |

Best Papers Session

| | |
|--|-----|
| Two-Level Main Memory Co-Design: Multi-threaded Algorithmic Primitives, Analysis, and Simulation | 835 |
| <i>Michael A. Bender, Jonathan Berry, Simon D. Hammond, K. Scott Hemmert, Samuel McCauley, Branden Moore, Benjamin Moseley, Cynthia A. Phillips, David Resnick, and Arun Rodrigues</i> | |
| CA-SVM: Communication-Avoiding Support Vector Machines on Distributed Systems | 847 |
| <i>Yang You, James Demmel, Kenneth Czechowski, Le Song, and Richard Vuduc</i> | |
| Filtering, Reductions and Synchronization in the Anton 2 Network | 860 |
| <i>J.P. Grossman, Brian Towles, Brian Greskamp, and David E. Shaw</i> | |

| | |
|---|-----|
| Notified Access: Extending Remote Memory Access Programming Models for Producer-Consumer Synchronization | 871 |
| <i>Roberto Belli and Torsten Hoefer</i> | |

Session 21: Algorithms for Fault Tolerance

| | |
|---|-----|
| 2W-FD: A Failure Detector Algorithm with QoS | 885 |
| <i>Alejandro Tomsic, Pierre Sens, João Garcia, Luciana Arantes, and Julien Sopena</i> | |
| Stabilizing Byzantine-Fault Tolerant Storage | 894 |
| <i>Silvia Bonomi, Maria Potop-Butucaru, and Sébastien Tixeuil</i> | |
| Making BFT Protocols Really Adaptive | 904 |
| <i>Jean-Paul Bahsoun, Rachid Guerraoui, and Ali Shoker</i> | |
| Exploration of Lossy Compression for Application-Level Checkpoint/Restart | 914 |
| <i>Naoto Sasaki, Kento Sato, Toshio Endo, and Satoshi Matsuoka</i> | |

Session 22: Scheduling and Load Balancing

| | |
|--|-----|
| Load-Balanced Local Time Stepping for Large-Scale Wave Propagation | 925 |
| <i>Max Rietmann, Daniel Peter, Olaf Schenk, Bora Uçar, and Marcus Grote</i> | |
| Towards Balance-Affinity Tradeoff in Concurrent Subgraph Traversals | 936 |
| <i>Yinglong Xia, Lifeng Nai, and Jui-Hsin Lai</i> | |
| Controlled Contention: Balancing Contention and Reservation in Multicore Application Scheduling | 946 |
| <i>Jingjing Wang, Nael Abu-Ghazaleh, and Dmitry Ponomarev</i> | |
| Resource and Deadline-Aware Job Scheduling in Dynamic Hadoop Clusters | 956 |
| <i>Dazhao Cheng, Jia Rao, Changjun Jiang, and Xiaobo Zhou</i> | |

Session 23: Heterogeneous Systems

| | |
|--|------|
| Mitigating the Susceptibility of GPGPUs Register File to Process Variations | 969 |
| <i>Jingweijia Tan and Xin Fu</i> | |
| PRO: Progress Aware GPU Warp Scheduling Algorithm | 979 |
| <i>Jayvant Anantpur and R. Govindarajan</i> | |
| Performance Impact of Batching Web-Application Requests Using Hot-Spot Processing on GPUs | 989 |
| <i>Tobias Fjälling and Per Stenström</i> | |
| An Approach for Energy Efficient Execution of Hybrid Parallel Programs | 1000 |
| <i>Lavanya Ramapantulu, Dumitrel Loghin, and Yong Meng Teo</i> | |

Session 24: I/O Optimizations

| | |
|---|------|
| Scheduling the I/O of HPC Applications Under Congestion | 1013 |
| <i>Ana Gainaru, Guillaume Aupy, Anne Benoit, Franck Cappello, Yves Robert, and Marc Snir</i> | |
| Leveraging Naturally Distributed Data Redundancy to Reduce Collective I/O Replication Overhead | 1023 |
| <i>Bogdan Nicolae</i> | |
| Exploring Data Staging Across Deep Memory Hierarchies for Coupled Data Intensive Simulation Workflows | 1033 |
| <i>Tong Jin, Fan Zhang, Qian Sun, Hoang Bui, Melissa Romanus, Norbert Podhorszki, Scott Klasky, Hemanth Kolla, Jacqueline Chen, Robert Hager, Choong-Seock Chang, and Manish Parashar</i> | |
| Reducing Vector I/O for Faster GPU Sparse Matrix-Vector Multiplication | 1043 |
| <i>Pham Nguyen Quang Anh, Rui Fan, and Yonggang Wen</i> | |

Session 25: Graph Algorithms

| | |
|--|------|
| Parallel Graph Partitioning for Complex Networks | 1055 |
| <i>Henning Meyerhenke, Peter Sanders, and Christian Schulz</i> | |
| A Self-Stabilizing Memory Efficient Algorithm for the Minimum Diameter Spanning Tree under an Omnipotent Daemon | 1065 |
| <i>Lélia Blin, Fadwa Boubeker, and Swan Dubois</i> | |
| A Parallel Tree Grafting Algorithm for Maximum Cardinality Matching in Bipartite Graphs | 1075 |
| <i>Ariful Azad, Aydin Buluç, and Alex Pothen</i> | |

Session 26: Resource Management

| | |
|--|------|
| Fair Resource Allocation for Heterogeneous Tasks | 1087 |
| <i>Koyel Mukherjee, Partha Dutta, Gurulingesh Raravi, Thangaraj Rajasubramaniam, Koustuv Dasgupta, and Atul Singh</i> | |
| Resources-Conscious Asynchronous High-Speed Data Transfer in Multicore Systems: Design, Optimizations, and Evaluation | 1097 |
| <i>Tan Li, Yufei Ren, Dantong Yu, and Shudong Jin</i> | |
| RISC: Robust Infrastructure over Shared Computing Resources through Dynamic Pricing and Incentivization | 1107 |
| <i>Tridib Mukherjee, Partha Dutta, Vinay G. Hegde, and Sujit Gujar</i> | |

Session 27: Architectural Support for Runtime and Thermal Management

| | |
|---|------|
| A Dual-Consistency Cache Coherence Protocol | 1119 |
| <i>Alberto Ros and Alexandra Jimboorean</i> | |
| Nexus#: A Distributed Hardware Task Manager for Task-Based Programming Models | 1129 |
| <i>Tamer Dallou, Nina Engelhardt, Ahmed Elhossini, and Ben Juurlink</i> | |
| Minimizing Thermal Variation Across System Components | 1139 |
| <i>Kaicheng Zhang, Seda Ogrenci-Memik, Gokhan Memik, Kazutomo Yoshii, Rajesh Sankaran, and Pete Beckman</i> | |

Session 28: Performance Monitoring and Prediction

| | |
|---|------|
| PCERE: Fine-Grained Parallel Benchmark Decomposition for Scalability | |
| Prediction | 1151 |
| <i>Mihail Popov, Chadi Akel, Florent Conti, William Jalby, and Pablo de Oliveira Castro</i> | |
| Matching Application Signatures for Performance Predictions Using a Single Execution | 1161 |
| <i>Anirudh Jayakumar, Prakash Murali, and Sathish Vadhiyar</i> | |
| Monitoring Large-Scale Location-Based Information Systems | 1171 |
| <i>Hammad Khan, Julien Gascon-Samson, Jörg Kienzle, and Bettina Kemme</i> | |

Author Index