

2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS 2017)

**Orlando, Florida, USA
29 May – 2 June 2017**

Pages 1-594



**IEEE Catalog Number: CFP17023-POD
ISBN: 978-1-5386-3915-3**

**Copyright © 2017 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 is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP17023-POD
ISBN (Print-On-Demand):	978-1-5386-3915-3
ISBN (Online):	978-1-5386-3914-6
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

2017 IEEE International Parallel and Distributed Processing Symposium

IPDPS 2017

Table of Contents

Message from the General Chair.....	xvi
Message from the Program Chair.....	xviii
Technical Program.....	xix
Conference Organization.....	xxi

Keynote 1

Computational Challenges in Constructing the Tree of Life	1
<i>Tandy Warnow</i>	

Session 1: Graph Algorithms

Monitoring Properties of Large, Distributed, Dynamic Graphs	2
<i>Gal Yehuda, Daniel Keren, and Islam Akaria</i>	
Parallel Construction of Suffix Trees and the All-Nearest-Smaller-Values Problem.....	12
<i>Patrick Flick and Srinivas Aluru</i>	
The Reverse Cuthill-McKee Algorithm in Distributed-Memory	22
<i>Ariful Azad, Mathias Jacquelin, Aydin Buluç, and Esmond G. Ng</i>	
SlimSell: A Vectorizable Graph Representation for Breadth-First Search	32
<i>Maciej Besta, Florian Marending, Edgar Solomonik, and Torsten Hoefler</i>	

Session 2: Computational Biology

SWhybrid: A Hybrid-Parallel Framework for Large-Scale Protein Sequence Database Search	42
<i>Haidong Lan, Weiguo Liu, Yongchao Liu, and Bertil Schmidt</i>	
PUNAS: A Parallel Ungapped-Alignment-Featured Seed Verification Algorithm for Next-Generation Sequencing Read Alignment.....	52
<i>Yuandong Chan, Kai Xu, Haidong Lan, Weiguo Liu, Yongchao Liu, and Bertil Schmidt</i>	

Eliminating Irregularities of Protein Sequence Search on Multicore Architectures	62
<i>Jing Zhang, Sanchit Misra, Hao Wang, and Wu-chun Feng</i>	
Communication Optimization on GPU: A Case Study of Sequence Alignment Algorithms	72
<i>Jie Wang, Xinfeng Xie, and Jason Cong</i>	
Session 3: Caches	
Elastic-Cache: GPU Cache Architecture for Efficient Fine- and Coarse-Grained Cache-Line Management	82
<i>Bingchao Li, Jizhou Sun, Murali Annavaram, and Nam Sung Kim</i>	
Content-Aware Non-Volatile Cache Replacement	92
<i>Qi Zeng and Jih-Kwon Peir</i>	
DEFT-Cache: A Cost-Effective and Highly Reliable SSD Cache for RAID Storage	102
<i>Jiguang Wan, Wei Wu, Ling Zhan, Qing Yang, Xiaoyang Qu, and Changsheng Xie</i>	
Adaptive Software Caching for Efficient NVRAM Data Persistence	112
<i>Pengcheng Li, Dhruva R. Chakrabarti, Chen Ding, and Liang Yuan</i>	
Session 4: Cloud & OS	
Container-Based Cloud Platform for Mobile Computation Offloading	123
<i>Song Wu, Chao Niu, Jia Rao, Hai Jin, and Xiaohai Dai</i>	
Enhancing Datacenter Resource Management through Temporal Logic Constraints	133
<i>Hao He, Jiang Hu, and Dilma Da Silva</i>	
High-Performance Virtual Machine Migration Framework for MPI Applications on SR-IOV Enabled InfiniBand Clusters	143
<i>Jie Zhang, Xiaoyi Lu, and Dhabaleswar K. (DK) Panda</i>	
Argo NodeOS: Toward Unified Resource Management for Exascale	153
<i>Swann Perarnau, Judicael A. Zounmevo, Matthieu Dreher, Brian C. Van Essen, Roberto Gioiosa, Kamil Iskra, Maya B. Gokhale, Kazutomo Yoshii, and Pete Beckman</i>	
Session 5: Distributed Algorithms	
Rational Fair Consensus in the Gossip Model	163
<i>Andrea Clementi, Luciano Gualà, Guido Proietti, and Giacomo Scornavacca</i>	
Leader Election in a Smartphone Peer-to-Peer Network	172
<i>Calvin Newport</i>	

Leader Election in Asymmetric Labeled Unidirectional Rings	182
<i>Karine Altisen, Ajoy K. Datta, Stéphane Devismes, Anaïs Durand, and Lawrence L. Larmore</i>	

Tight Load Balancing Via Randomized Local Search	192
<i>Petra Berenbrink, Peter Kling, Christopher Liaw, and Abbas Mehrabian</i>	

Session 6: Numerical Simulation

Large Scale Manycore-Aware PIC Simulation with Efficient Particle Binning	202
<i>Hiroshi Nakashima, Yoshiki Summura, Keisuke Kikura, and Yohei Miyake</i>	

Optimization and Parallelization of B-Spline Based Orbital Evaluations in QMC on Multi/Many-Core Shared Memory Processors	213
<i>Amrita Mathuriya, Ye Luo, Anouar Benali, Luke Shulenburger, and Jeongnim Kim</i>	

One-Way Wave Equation Migration at Scale on GPUs Using Directive Based Programming	224
<i>Kshitij Mehta, Maxime Hugues, Oscar Hernandez, David E. Bernholdt, and Henri Calandra</i>	

Towards Highly scalable Ab Initio Molecular Dynamics (AIMD) Simulations on the Intel Knights Landing Manycore Processor	234
<i>Mathias Jacquelin, Wibe De Jong, and Eric Bylaska</i>	

Session 7: Novel Architectures

General Purpose Task-Dependence Management Hardware for Task-Based Dataflow Programming Models	244
<i>Xubin Tan, Jaume Bosch, Miquel Vidal, Carlos Álvarez, Daniel Jiménez-González, Eduard Ayguadé, and Mateo Valero</i>	

Accelerating Graph and Machine Learning Workloads Using a Shared Memory Multicore Architecture with Auxiliary Support for In-hardware Explicit Messaging	254
<i>Halit Dogan, Farrukh Hijaz, Masab Ahmad, Brian Kahne, Peter Wilson, and Omer Khan</i>	

Respin: Rethinking Near-Threshold Multiprocessor Design with Non-volatile Memory	265
<i>Xiang Pan, Anys Bacha, and Radu Teodorescu</i>	

MOCHA: Morphable Locality and Compression Aware Architecture for Convolutional Neural Networks	276
<i>Syed Mohammad Asad Hassan Jafri, Ahmed Hemani, Kolin Paul, and Naeem Abbas</i>	

Session 8: Performance Modeling and Tuning

Autotuning Stencil Computations with Structural Ordinal Regression Learning	287
<i>Biagio Cosenza, Juan J. Durillo, Stefano Ermon, and Ben Juurlink</i>	
Capability Models for Manycore Memory Systems: A Case-Study with Xeon Phi KNL	297
<i>Sabela Ramos and Torsten Hoefler</i>	
Apollo: Reusable Models for Fast, Dynamic Tuning of Input-Dependent Code	307
<i>David Beckingsale, Olga Pearce, Ignacio Laguna, and Todd Gamblin</i>	
Generating Performance Models for Irregular Applications	317
<i>Ryan D. Friese, Nathan R. Tallent, Abhinav Vishnu, Darren J. Kerbyson, and Adolfo Hoisie</i>	

Session 9: Communication & Coordination

Bounded Reordering Allows Efficient Reliable Message Transmission	327
<i>Keishla D Ortiz-Lopez and Jennifer L. Welch</i>	
Dynamic Adaptation in Wireless Networks Under Comprehensive Interference via Carrier Sense	337
<i>Dongxiao Yu, Yuexuan Wang, Tigran Tonoyan, and Magnús M. Halldórsson</i>	
Fault-Tolerant Online Packet Scheduling on Parallel Channels	347
<i>Paweł Garncaiek, Tomasz Jurdziński, and Krzysztof Loryś</i>	
Corrected Gossip Algorithms for Fast Reliable Broadcast on Unreliable Systems	357
<i>Torsten Hoefler, Amnon Barak, Amnon Shiloh, and Zvi Drezner</i>	

Session 10: Tools 1

DR-BW: Identifying Bandwidth Contention in NUMA Architectures with Supervised Learning	367
<i>Hao Xu, Shasha Wen, Alfredo Gimenez, Todd Gamblin, and Xu Liu</i>	
Data Centric Performance Measurement Techniques for Chapel Programs	377
<i>Hui Zhang and Jeffrey K. Hollingsworth</i>	
A Parallel FastTrack Data Race Detector on Multi-core Systems	387
<i>Young Wn Song and Yann-Hang Lee</i>	
Localized Fault Recovery for Nested Fork-Join Programs	397
<i>Gokcen Kestor, Sriram Krishnamoorthy, and Wenjing Ma</i>	

Session 11: Networks

Exploring DataVortex Systems for Irregular Applications	409
<i>Roberto Gioiosa, Antonino Tumeo, Jian Yin, Thomas Warfel, David Haglin, and Santiago Betelu</i>	
DC2-MTCP: Light-Weight Coding for Efficient Multi-Path Transmission in Data Center Network	419
<i>Jiyan Sun, Yan Zhang, Xin Wang, Shihan Xiao, Zhen Xu, Hongjing Wu, Xin Chen, and Yanni Han</i>	
A Scalable and Resilient Microarchitecture Based on Multiport Binding for High-Radix Router Design	429
<i>Yi Dai, Kefei Wang, Gang Qu, Liquan Xiao, Dezun Dong, and Xingyun Qi</i>	
Partitioning Low-Diameter Networks to Eliminate Inter-Job Interference	439
<i>Nikhil Jain, Abhinav Bhatele, Xiang Ni, Todd Gamblin, and Laxmikant V. Kale</i>	

Session 12: Libraries & Frameworks

Accelerating Spark Datasets by Inlining Deserialization	449
<i>Jan Wróblewski, Kazuaki Ishizaki, Hiroshi Inoue, and Moriyoshi Ohara</i>	
MRapid: An Efficient Short Job Optimizer on Hadoop	459
<i>Hong Zhang, Hai Huang, and Liqiang Wang</i>	
Accommodating Thread-Level Heterogeneity in Coupled Parallel Applications	469
<i>Samuel K. Gutiérrez, Kei Davis, Dorian C. Arnold, Randal S. Baker, Robert W. Robey, Patrick McCormick, Daniel Holladay, Jon A. Dahl, R. Joe Zerr, Florian Weik, and Christoph Junghans</i>	
Multi-GPU Graph Analytics	479
<i>Yuechao Pan, Yangzihao Wang, Yuduo Wu, Carl Yang, and John D. Owens</i>	

Industry Tutorial

NVIDIA Deep Learning Tutorial	491
<i>Julie Bernauer</i>	

Keynote 2

A Scalable System Architecture to Addressing the Next Generation of Predictive Simulation Workflows with Coupled Compute and Data Intensive Applications	492
<i>Mark Seager</i>	

Session 13: Motion Planning & Similarity Search

Fault-Tolerant Robot Gathering Problems on Graphs With Arbitrary Appearing Times	493
<i>Sergio Rajsbaum, Armando Castañeda, David Flores Peñaloza, and Manuel Alcántara</i>	
Distributed Vehicle Routing Approximation	503
<i>Akhil Krishnan, Mikhail Markov, and Borzoo Bonakdarpour</i>	
O(log N)-Time Complete Visibility for Asynchronous Robots with Lights	513
<i>Gokarna Sharma, Ramachandran Vaidyanathan, Jerry L. Trahan, Costas Busch, and Suresh Rai</i>	
Similarity Search on Automata Processors	523
<i>Vincent T. Lee, Justin Kotalik, Carlo C. del Mundo, Armin Alaghi, Luis Ceze, and Mark Oskin</i>	

Session 14: Applications

26 PFLOPS Stencil Computations for Atmospheric Modeling on Sunway TaihuLight	535
<i>Yulong Ao, Chao Yang, Xinliang Wang, Wei Xue, Haohuan Fu, Fangfang Liu, Lin Gan, Ping Xu, and Wenjing Ma</i>	
Image-Domain Gridding on Graphics Processors	545
<i>Bram Veenboer, Matthias Petschow, and John W. Romein</i>	
Aces4: A Platform for Computational Chemistry Calculations with Extremely Large Block-Sparse Arrays	555
<i>Beverly A. Sanders, Jason N. Byrd, Nakul Jindal, Victor F. Lotrich, Dmitry Lyakh, Ajith Perera, and Rodney J. Bartlett</i>	
PhiOpenSSL: Using the Xeon Phi Coprocessor for Efficient Cryptographic Calculations	565
<i>Shun Yao and Dantong Yu</i>	

Session 15: Tools 2

Directive-Based Partitioning and Pipelining for Graphics Processing Units	575
<i>Xuewen Cui, Thomas R. W. Scogland, Bronis R. de Supinski, and Wu-chun Feng</i>	
ScalaIOExtrap: Elastic I/O Tracing and Extrapolation	585
<i>Xiaoqing Luo, Frank Mueller, Philip Carns, Jonathan Jenkins, Robert Latham, Robert Ross, and Shane Snyder</i>	
SimProf: A Sampling Framework for Data Analytic Workloads	595
<i>Jen-Cheng Huang, Lifeng Nai, Pranith Kumar, Hyojong Kim, and Hyesoon Kim</i>	
PaPar: A Parallel Data Partitioning Framework for Big Data Applications	605
<i>Hao Wang, Jing Zhang, Da Zhang, Sarunya Pumma, and Wu-chun Feng</i>	

Session 16: Data and Graph Analytics

swDNN: A Library for Accelerating Deep Learning Applications on Sunway TaihuLight	615
<i>Jiarui Fang, Haohuan Fu, Wenlai Zhao, Bingwei Chen, Weijie Zheng, and Guangwen Yang</i>	
Community Detection on the GPU	625
<i>Md. Naim, Fredrik Manne, Mahantesh Halappanavar, and Antonino Tumeo</i>	
Scalable Graph Traversal on Sunway TaihuLight with Ten Million Cores	635
<i>Heng Lin, Xiongchao Tang, Bowen Yu, Youwei Zhuo, Wenguang Chen, Jidong Zhai, Wanwang Yin, and Weimin Zheng</i>	
Partitioning Trillion-Edge Graphs in Minutes	646
<i>George M. Slota, Sivasankaran Rajamanickam, Karen Devine, and Kamesh Madduri</i>	

Session 17: Linear Algebra

Generating Families of Practical Fast Matrix Multiplication Algorithms	656
<i>Jianyuan Huang, Leslie Rice, Devin A. Matthews, and Robert A. van de Geijn</i>	
Bidiagonalization and R-Bidiagonalization: Parallel Tiled Algorithms, Critical Paths and Distributed-Memory Implementation	668
<i>Mathieu Faverge, Julien Langou, Yves Robert, and Jack Dongarra</i>	
Communication-Avoiding Parallel Algorithms for Solving Triangular Systems of Linear Equations	678
<i>Tobias Wicky, Edgar Solomonik, and Torsten Hoefler</i>	
A Work-Efficient Parallel Sparse Matrix-Sparse Vector Multiplication Algorithm	688
<i>Ariful Azad and Aydin Buluç</i>	

Session 18: Power Management

Power Efficient Sharing-Aware GPU Data Management	698
<i>Abdulaziz Tabbakh, Murali Annavaram, and Xuehai Qian</i>	
Fly-Over: A Light-Weight Distributed Power-Gating Mechanism for Energy-Efficient Networks-on-Chip	708
<i>Rahul Boyapati, Jiayi Huang, Ningyuan Wang, Kyung Hoon Kim, Ki Hwan Yum, and Eun Jung Kim</i>	
RCube: A Power Efficient and Highly Available Network for Data Centers	718
<i>Zhenhua Li and Yuanyuan Yang</i>	
Cooling-Aware Job Scheduling and Node Allocation for Overprovisioned HPC Systems	728
<i>Thang Cao, Wei Huang, Yuan He, and Masaaki Kondo</i>	

Session 19: Scheduling

Algorithms for Hierarchical and Semi-Partitioned Parallel Scheduling	738
<i>Vincenzo Bonifaci, Gianlorenzo D'Angelo, and Alberto Marchetti-Spaccamela</i>	
Efficient and Deterministic Scheduling for Parallel State Machine Replication	748
<i>Odorico M. Mendizabal, Rudá S. T. De Moura, Fernando Luís Dotti, and Fernando Pedone</i>	
Dynamic Memory-Aware Task-Tree Scheduling	758
<i>Guillaume Aupy, Clément Brasseur, and Loris Marchal</i>	
Approximation Proofs of a Fast and Efficient List Scheduling Algorithm for Task-Based Runtime Systems on Multicores and GPUs	768
<i>Olivier Beaumont, Lionel Eyraud-Dubois, and Suraj Kumar</i>	

Session 20: Code Optimization

Automatic Collapsing of Non-Rectangular Loops	778
<i>Philippe Clauss, Ervin Altintas, and Matthieu Kuhn</i>	
HOMP: Automated Distribution of Parallel Loops and Data in Highly Parallel Accelerator-Based Systems	788
<i>Yonghong Yan, Jiawen Liu, Kirk W. Cameron, and Mariam Umar</i>	
Multigrain Parallelism: Bridging Coarse-Grain Parallel Programs and Fine-Grain Event-Driven Multithreading	799
<i>Jaime Arteaga Molina, Stéphane Zuckerman, and Guang R. Gao</i>	
Improving the Integration of Task Nesting and Dependencies in OpenMP	809
<i>Josep M. Perez, Vicenç Beltran, Jesus Labarta, and Eduard Ayguadé</i>	

Keynote 3

Runtime Aware Architectures	819
<i>Mateo Valero</i>	

Best Papers

Reducing Pagerank Communication via Propagation Blocking	820
<i>Scott Beamer, Krste Asanović, and David Patterson</i>	
Clustering Throughput Optimization on the GPU	832
<i>Michael Gowanlock, Cody M. Rude, David M. Blair, Justin D. Li, and Victor Pankratius</i>	
FlexVC: Flexible Virtual Channel Management in Low-Diameter Networks	842
<i>Pablo Fuentes, Enrique Vallejo, Ramón Beivide, Cyriel Minkenberg, and Mateo Valero</i>	

Relaxations for High-Performance Message Passing on Massively Parallel SIMT Processors	855
<i>Benjamin Klenk, Holger Fröening, Hans Eberle, and Larry Dennison</i>	

Session 21: Algorithms

The SEPO Model of Computation to Enable Larger-Than-Memory Hash Tables for GPU-Accelerated Big Data Analytics	866
<i>Reza Mokhtari and Michael Stumm</i>	
Elastic Consistent Hashing for Distributed Storage Systems	876
<i>Wei Xie and Yong Chen</i>	
An N log N Parallel Fast Direct Solver for Kernel Matrices	886
<i>Chenhan D. Yu, William B. March, and George Biros</i>	
A Robust Parallel Preconditioner for Indefinite Systems Using Hierarchical Matrices and Randomized Sampling	897
<i>Pieter Ghysels, Xiaoye Sherry Li, Christopher Gorman, and François-Henry Rouet</i>	

Session 22: Coordination

FFQ: A Fast Single-Producer/Multiple-Consumer Concurrent FIFO Queue	907
<i>Sergei Arnautov, Pascal Felber, Christof Fetzer, and Bohdan Trach</i>	
Scalable Lock-Free Vector with Combining	917
<i>Ivan Walulya and Philippos Tsigas</i>	
Automatic-Signal Monitors with Multi-object Synchronization	927
<i>Wei-Lun Hung and Vijay K. Garg</i>	
Optimal Algorithms for a Mesh-Connected Computer with Limited Additional Global Bandwidth	937
<i>Yujie An and Quentin F. Stout</i>	

Session 23: Power Management 2

An Adaptive Core-Specific Runtime for Energy Efficiency	947
<i>Sridutt Bhalachandra, Allan Porterfield, Stephen L. Olivier, and Jan F. Prins</i>	
Production Hardware Overprovisioning: Real-World Performance Optimization Using an Extensible Power-Aware Resource Management Framework	957
<i>Ryuichi Sakamoto, Thang Cao, Masaaki Kondo, Koji Inoue, Masatsugu Ueda, Tapasya Patki, Daniel Ellsworth, Barry Rountree, and Martin Schulz</i>	
Co-Run Scheduling with Power Cap on Integrated CPU-GPU Systems	967
<i>Qi Zhu, Bo Wu, Xipeng Shen, Li Shen, and Zhiying Wang</i>	

Characterizing and Modeling Power and Energy for Extreme-Scale In-Situ Visualization	978
<i>Vignesh Adhinarayanan, Wu-chun Feng, David Rogers, James Ahrens, and Scott Pakin</i>	

Session 24: MPI

Application Level Reordering of Remote Direct Memory Access Operations	988
<i>Wim Lavrijsen and Costin Iancu</i>	
Toucan — A Translator for Communication Tolerant MPI Applications	998
<i>Sergio M. Martin, Marsha J. Berger, and Scott B. Baden</i>	
Memory Compression Techniques for Network Address Management in MPI	1008
<i>Yanfei Guo, Charles J. Archer, Michael Blocksome, Scott Parker, Wesley Bland, Ken Raffanetti, and Pavan Balaji</i>	
Transparent Caching for RMA Systems	1018
<i>Salvatore Di Girolamo, Flavio Vella, and Torsten Hoefler</i>	

Session 25: ML & Tensors

When Neurons Fail	1028
<i>El Mahdi El Mhamdi and Rachid Guerraoui</i>	
On Optimizing Distributed Tucker Decomposition for Dense Tensors	1038
<i>Venkatesan T. Chakaravarthy, Jee W. Choi, Douglas J. Joseph, Xing Liu, Prakash Murali, Yogish Sabharwal, and Dheeraj Sreedhar</i>	
Model-Driven Sparse CP Decomposition for Higher-Order Tensors	1048
<i>Jiajia Li, Jee Choi, Ioakeim Perros, Jimeng Sun, and Richard Vuduc</i>	
Sparse Tensor Factorization on Many-Core Processors with High-Bandwidth Memory	1058
<i>Shaden Smith, Jongsoo Park, and George Karypis</i>	

Session 26: Resource Management

Proximity-Aware Balanced Allocations in Cache Networks	1068
<i>Ali Pourmiri, Mahdi Jafari Siavoshani, and Seyed Pooya Shariatpanahi</i>	
Addressing Performance Heterogeneity in MapReduce Clusters with Elastic Tasks	1078
<i>Wei Chen, Jia Rao, and Xiaobo Zhou</i>	
Autonomic Resource Management for Program Orchestration in Large-Scale Data Analysis	1088
<i>Masahiro Tanaka, Kenjiro Taura, and Kentaro Torisawa</i>	

Mimir: Memory-Efficient and Scalable MapReduce for Large Supercomputing Systems	1098
<i>Tao Gao, Yanfei Guo, Boyu Zhang, Pietro Cicotti, Yutong Lu, Pavan Balaji, and Michela Taufer</i>	

Session 27: Compression & Memoization

Elastic Data Compression with Improved Performance and Space Efficiency for Flash-Based Storage Systems	1109
<i>Bo Mao, Hong Jiang, Suzhen Wu, Yaodong Yang, and Zaifa Xi</i>	
E ² MC: Entropy Encoding Based Memory Compression for GPUs	1119
<i>Sohan Lal, Jan Lucas, and Ben Juurlink</i>	
Significantly Improving Lossy Compression for Scientific Data Sets Based on Multidimensional Prediction and Error-Controlled Quantization	1129
<i>Dingwen Tao, Sheng Di, Zizhong Chen, and Franck Cappello</i>	
ATM: Approximate Task Memoization in the Runtime System	1140
<i>Iulian Brumar, Marc Casas, Miquel Moreto, Mateo Valero, and Gurindar S. Sohi</i>	

Session 28: Persistent Memory

Design and Implementation of Papyrus: Parallel Aggregate Persistent Storage	1151
<i>Jungwon Kim, Kittisak Sajjapongse, Seyong Lee, and Jeffrey S. Vetter</i>	
Language-Based Optimizations for Persistence on Nonvolatile Main Memory Systems	1163
<i>Joel Edward Denny, Seyong Lee, and Jeffrey S. Vetter</i>	
MetaKV: A Key-Value Store for Metadata Management of Distributed Burst Buffers	1174
<i>Teng Wang, Adam Moody, Yue Zhu, Kathryn Mohror, Kento Sato, Tanzima Islam, and Weikuan Yu</i>	
Parallelism and Garbage Collection Aware I/O Scheduler with Improved SSD Performance	1184
<i>Jiayang Guo, Yiming Hu, Bo Mao, and Suzhen Wu</i>	

Author Index