

2017 IEEE International Conference on Cluster Computing (CLUSTER 2017)

**Honolulu, Hawaii, USA
5 – 8 September 2017**



**IEEE Catalog Number: CFP17235-POD
ISBN: 978-1-5386-2327-5**

**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:	CFP17235-POD
ISBN (Print-On-Demand):	978-1-5386-2327-5
ISBN (Online):	978-1-5386-2326-8
ISSN:	1552-5244

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

2017 IEEE International Conference on Cluster Computing

CLUSTER 2017

Table of Contents

Message from the General Co-Chairs.....	xvi
Message from the Program Chairs.....	xvii
FTS Workshop Message.....	xix
HPCMASPA 2017 Workshop Message.....	xx
WRAp Workshop Message.....	xxii
REV-A Workshop Message.....	xxiii
Committees.....	xxiv
Technical Program Committee.....	xxvi
Keynotes.....	xxxi
Sponsors.....	xxxiv

Session 1: Best Papers 1

Towards Practical and Robust Labeled Pattern Matching in Trillion-Edge Graphs	1
<i>Tahsin Reza, Christine Klymko, Matei Ripeanu, Geoffrey Sanders, and Roger Pearce</i>	
Contention-Aware Kernel-Assisted MPI Collectives for Multi-/Many-Core Systems	13
<i>Sourav Chakraborty, Hari Subramoni, and Dhabaleswar K. Panda</i>	

Session 2: Algorithms

Parallel and Efficient Sensitivity Analysis of Microscopy Image Segmentation	
Workflows in Hybrid Systems	25
<i>Willian Barreiros Jr., George Teodoro, Tahsin Kurc, Jun Kong, Alba C. M. A. Melo, and Joel Saltz</i>	
S-Aligner: Ultrascalable Read Mapping on Sunway Taihu Light	36
<i>Xiaohui Duan, Kai Xu, Yuandong Chan, Christian Hundt, Bertil Schmidt, Pavan Balaji, and Weiguo Liu</i>	
A Unified Optimization Approach for Sparse Tensor Operations on GPUs	47
<i>Bangtian Liu, Chengyao Wen, Anand D. Sarwate, and Maryam Mehri Dehnavi</i>	

Session 3: Big Data and HPC

Canopus: A Paradigm Shift Towards Elastic Extreme-Scale Data Analytics on HPC Storage	58
<i>Tao Lu, Eric Suchyta, Dave Pugmire, Jong Choi, Scott Klasky, Qing Liu, Norbert Podhorszki, Mark Ainsworth, and Matthew Wolf</i>	
TAPIOCA: An I/O Library for Optimized Topology-Aware Data Aggregation on Large-Scale Supercomputers	70
<i>François Tessier, Venkatram Vishwanath, and Emmanuel Jeannot</i>	
Could Blobs Fuel Storage-Based Convergence Between HPC and Big Data?	81
<i>Pierre Matri, Yevhen Alforov, Álvaro Brandon, Michael Kuhn, Philip Carns, and Thomas Ludwig</i>	
Eley: On the Effectiveness of Burst Buffers for Big Data Processing in HPC Systems	87
<i>Orcun Yıldız, Amelie Chi Zhou, and Shadi Ibrahim</i>	

Session 4: Performance and Energy Modeling and Analysis

Predicting the Energy-Consumption of MPI Applications at Scale Using Only a Single Node	92
<i>Franz Christian Heinrich, Tom Cornebize, Augustin Degomme, Arnaud Legrand, Alexandra Carpen-Amarie, Sascha Hunold, Anne-Cécile Orgerie, and Martin Quinson</i>	
Mira: A Framework for Static Performance Analysis	103
<i>Kewen Meng and Boyana Norris</i>	
Delay Spotter: A Tool for Spotting Scheduler-Caused Delays in Task Parallel Runtime Systems	114
<i>An Huynh and Kenjiro Taura</i>	

Session 5: Resource Management and Runtime Systems

Automating the Application Data Placement in Hybrid Memory Systems	126
<i>Harald Servat, Antonio J. Peña, Germán Llort, Estanislao Mercadal, Hans-Christian Hoppe, and Jesús Labarta</i>	
QoS- and Contention- Aware Resource Provisioning in a Stream Processing Engine	137
<i>MohammadReza HoseinyFarahabady, Albert Y. Zomaya, and Zahir Tari</i>	
Runtime Techniques for Programming with Fast and Slow Memory	147
<i>Xiang Ni, Nikhil Jain, Kavitha Chandrasekar, and Laxmikant V. Kale</i>	

Session 6: Memory and Networks

Utility-Based Hybrid Memory Management	152
<i>Yang Li, Saugata Ghose, Jongmoo Choi, Jin Sun, Hui Wang, and Onur Mutlu</i>	

Performance Modeling for Optimal Data Placement on GPU with Heterogeneous Memory Systems	166
---	-----

Yingchao Huang and Dong Li

A Case for Uni-directional Network Topologies in Large-Scale Clusters	178
---	-----

*Michihiro Koibuchi, Tomohiro Totoki, Hiroki Matsutani, Hideharu Amano,
Fabien Chaix, Ikki Fujiwara, and Henri Casanova*

A Wait-Free Multi-word Atomic (1,N) Register for Large-Scale Data Sharing on Multi-core Machines	188
---	-----

Mauro Ianni, Alessandro Pellegrini, and Francesco Quaglia

Session 7: Visualization and I/O

Visual Analytics Techniques for Exploring the Design Space of Large-Scale High-Radix Networks	193
--	-----

*Jianping Kelvin Li, Misbah Mubarak, Robert B. Ross, Christopher D. Carothers,
and Kwan-Liu Ma*

Quantifying I/O and Communication Traffic Interference on Dragonfly Networks Equipped with Burst Buffers	204
---	-----

*Misbah Mubarak, Philip Carns, Jonathan Jenkins, Jianping Kelvin Li, Nikhil Jain,
Shane Snyder, Robert Ross, Christopher D. Carothers, Abhinav Bhatale,
and Kwan-Liu Ma*

Spatiotemporal Wavelet Compression for Visualization of Scientific Simulation Data	216
--	-----

Shaomeng Li, Sudhanshu Sane, Leigh Orf, Pablo Mininni, John Clyne, and Hank Childs

AUTOBAHN: Accelerating Concurrent, Durable File I/O via a Non-volatile Buffer	228
---	-----

*Hyeongwon Jang, Sang Youp Rhee, Jae Eun Kim, Sooyong Kang, Hyuck Han,
and Hyungsoo Jung*

Session 8: Best Papers 2

Justice: A Deadline-Aware, Fair-Share Resource Allocator for Implementing Multi-Analytics	233
--	-----

Stratos Dimopoulos, Chandra Krintz, and Rich Wolski

Accelerating a Burst Buffer Via User-Level I/O Isolation	245
--	-----

*Jaehyun Han, Donghun Koo, Glenn K. Lockwood, Jaehwan Lee, Hyeonsang Eom,
and Soonwook Hwang*

Session 9: High Performance Big Data Analytics

GraphH: High Performance Big Graph Analytics in Small Clusters	256
--	-----

Peng Sun, Yonggang Wen, Ta Nguyen Bin Duong, and Xiaokui Xiao

Parallel Multivariate Spatio-Temporal Clustering of Large Ecological Datasets on Hybrid Supercomputers	267
---	-----

*Sarat Sreepathi, Jitendra Kumar, Richard T. Mills, Forrest M. Hoffman,
Vamsi Sripathi, and William W. Hargrove*

A Comparison of Graph-Based Synthetic Data Generators for Benchmarking Next-Generation Intrusion Detection Systems	278
---	-----

*Stefano Iannucci, Hisham A. Kholidy, Amrita Dhakal Ghimire, Rui Jia,
Sherif Abdelwahed, and Ioana Banicescu*

Session 10: Virtualization

Dynamically Compiled Artifact Sharing for Clouds	290
--	-----

Panagiotis Patros, Dayal Dilli, Kenneth B. Kent, and Michael Dawson

ConVGPU: GPU Management Middleware in Container Based Virtualized Environment	301
--	-----

Daeyoun Kang, Tae Joon Jun, Dohyeun Kim, Jaewook Kim, and Daeyoung Kim

Enabling Diverse Software Stacks on Supercomputers Using High Performance Virtual Clusters	310
---	-----

Andrew J. Younge, Kevin Pedretti, Ryan E. Grant, Brian L. Gaines, and Ron Brightwell

Session 11: Emerging Architectures and Parallel Processing

EclipseMR: Distributed and Parallel Task Processing with Consistent Hashing	322
---	-----

*Vicente A. B. Sanchez, Wonbae Kim, Youngmoon Eom, Kibeom Jin, Moohyeon Nam,
Deukyeon Hwang, Jik-Soo Kim, and Beomseok Nam*

Understanding the Role of GPGPU-Accelerated SoC-Based ARM Clusters	333
--	-----

Reza Azimi, Tyler Fox, and Sherief Reda

Effective Running of End-to-End HPC Workflows on Emerging Heterogeneous Architectures	344
--	-----

Kun Tang, Devesh Tiwari, Saurabh Gupta, Sudharshan S. Vazhkudai, and Xubin He

High Throughput and Low Latency on Hadoop Clusters Using Explicit Congestion Notification: The Untold Truth	349
--	-----

Renan Fischer e Silva and Paul M. Carpenter

A Scalable Network-Based Performance Analysis Tool for MPI on Large-Scale HPC Systems	354
--	-----

Hari Subramoni, Xiaoyi Lu, and Dhabaleswar K. Panda

Session 12: Data Storage and Processing

SoMeta: Scalable Object-Centric Metadata Management for High Performance Computing	359
<i>Houjun Tang, Suren Byna, Bin Dong, Jialin Liu, and Quincey Koziol</i>	
Automatic Data Filtering for In Situ Workflows	370
<i>Clément Mommessin, Matthieu Dreher, Bruno Raffin, and Tom Peterka</i>	
Task Allocation for Stream Processing with Recovery Latency Guarantee	379
<i>Hongliang Li, Jie Wu, Zhen Jiang, Xiang Li, and Xiaohui Wei</i>	
A Comparative Analysis of Materialized Views Selection and Concurrency Control Mechanisms in NoSQL Databases	384
<i>Ashish Tapdiya, Yuan Xue, and Daniel Fabbri</i>	

Session 13: Scheduling

Automatic, Abstracted and Portable Topology-Aware Thread Placement	389
<i>Jens Gustedt, Emmanuel Jeannot, and Farouk Mansouri</i>	
Dynamic Co-Scheduling Driven by Main Memory Bandwidth Utilization	400
<i>Jens Breitbart, Simon Pickartz, Stefan Lankes, Josef Weidendorfer, and Antonello Monti</i>	

Session 14: Performance Profiling

Tracking System Behavior from Resource Usage Data	410
<i>Niyazi Sorkunlu, Varun Chandola, and Abani Patra</i>	
Flexible Data Aggregation for Performance Profiling	419
<i>David Böhme, David Beckingsale, and Martin Schulz</i>	

Session 15: Leveraging Accelerated Systems

Implementing Lattice QCD Application with XcalableACC Language on Accelerated Cluster	429
<i>Masahiro Nakao, Hitoshi Murai, Hidetoshi Iwashita, Akihiro Tabuchi, Taisuke Boku, and Mitsuhsisa Sato</i>	
DH-Falcon: A Language for Large-Scale Graph Processing on Distributed Heterogeneous Systems	439
<i>Unnikrishnan Cheramangalath, Rupesh Nasre, and Y. N. Srikanth</i>	
Fast Failure Erasure Encoding Using Just in Time Compilation for CPUs, GPUs, and FPGAs	451
<i>David Rohr and Volker Lindenstruth</i>	

Session 16: Fault Tolerance

Toward a General Theory of Optimal Checkpoint Placement	464
<i>Omer Subasi, Gokcen Kestor, and Sriram Krishnamoorthy</i>	
Algorithm-Directed Crash Consistence in Non-volatile Memory for HPC	475
<i>Shuo Yang, Kai Wu, Yifan Qiao, Dong Li, and Jidong Zhai</i>	
Checkpointing Workflows for Fail-Stop Errors	487
<i>Li Han, Louis-Claude Canon, Henri Casanova, Yves Robert, and Frédéric Vivien</i>	

Session 17: Numerical Methods and Libraries

Exploring On-Node Parallelism with Neutral, a Monte Carlo Neutral Particle Transport Mini-App	498
<i>Matt Martineau and Simon McIntosh-Smith</i>	
Manala: A Flexible Flow Control Library for Asynchronous Task Communication	509
<i>Matthieu Dreher, Kiran Sasikumar, Subramanian Sankaranarayanan, and Tom Peterka</i>	
Distributed Affine-Invariant MCMC Sampler	520
<i>Balazs Nemeth, Tom Haber, Jori Liesenborgs, and Wim Lamotte</i>	
A Stencil Framework to Realize Large-Scale Computations Beyond Device Memory Capacity on GPU Supercomputers	525
<i>Takashi Shimokawabe, Toshio Endo, Naoyuki Onodera, and Takayuki Aoki</i>	

Session 18: Programming and Systems Software

Trade-Off Between Prediction Accuracy and Underestimation Rate in Job Runtime Estimates	530
<i>Yuping Fan, Paul Rich, William E. Allcock, Michael E. Papka, and Zhiling Lan</i>	
CLIP: Cluster-Level Intelligent Power Coordination for Power-Bounded Systems	541
<i>Pengfei Zou, Tyler Allen, Claude H. Davis IV, Xizhou Feng, and Rong Ge</i>	
Pure Functions in C: A Small Keyword for Automatic Parallelization	552
<i>Tim Süß, Lars Nagel, Marc-André Vef, André Brinkmann, Dustin Feld, and Thomas Sodemann</i>	
The Effect of Resource Allocation and System Events on VM Consolidation	557
<i>Maruf Ahmed and Albert Y. Zomaya</i>	

Session 19: Algorithms and Tools for I/O and Big Data Management

Extending Skel to Support the Development and Optimization of Next Generation I/O Systems	563
<i>Jeremy Logan, Jong Youl Choi, Matthew Wolf, George Ostrouchov, Lipeng Wan, Norbert Podhorszki, William Godoy, Scott Klasky, Erich Lohrmann, Greg Eisenhauer, Chad Wood, and Kevin Huck</i>	
keybin: Key-Based Binning for Distributed Clustering	572
<i>Xinyu Chen, Jeremy Benson, and Trilce Estrada</i>	
Optimizing the Datapath for Key-value Middleware with NVMe SSDs over RDMA Interconnects	582
<i>Zhongqi An, Zhengyu Zhang, Qiang Li, Jing Xing, Hao Du, Zhan Wang, Zhigang Huo, and Jie Ma</i>	
TGE: Machine Learning Based Task Graph Embedding for Large-Scale Topology Mapping	587
<i>Jong Youl Choi, Jeremy Logan, Matthew Wolf, George Ostrouchov, Tahsin Kurc, Qing Liu, Norbert Podhorszki, Scott Klasky, Melissa Romanus, Qian Sun, Manish Parashar, Randy Michael Churchill, and CS Chang</i>	

Session 20: Silent Data Corruption

Detection of Silent Data Corruption in Adaptive Numerical Integration Solvers	592
<i>Pierre-Louis Guhur, Emil Constantinescu, Debojyoti Ghosh, Tom Peterka, and Franck Cappello</i>	
Evaluating the Viability of Using Compression to Mitigate Silent Corruption of Read-Mostly Application Data	603
<i>Scott Levy, Kurt B. Ferreira, and Patrick G. Bridges</i>	
A Gaussian Process Approach for Effective Soft Error Detection	608
<i>Omer Subasi and Sriram Krishnamoorthy</i>	

Poster Papers

A Novel Hybrid Transactional Memory Based on Abort Prediction and Adaptive Retry Policy	613
<i>Young-Sung Shin, Yeon-Woo Jang, Moon-Hwan Kang, and Jae-Woo Chang</i>	
Mitigating the Write Amplification Problem of Write-Optimized File Systems on Flash Storage	615
<i>Shuo-Han Chen, Jun-Long Lin, Tseng-Yi Chen, Tsan-Sheng Hsu, Hsin-Wen Wei, and Wei-Kuan Shih</i>	
Acceleration of Turbulent Flow Simulations with Intel Xeon Phi(TM) Manycore Processors	617
<i>Ji-Hoon Kang and Hoon Ryu</i>	

Distributed Parallel Backprojection for Real-Time Stripmap SAR Imaging on GPU Clusters	619
<i>Masato Gocho, Noboru Oishi, and Atsuo Ozaki</i>	
Parallelized Recovery of Hundreds of Millions Small Data Objects	621
<i>Kevin Beineke, Stefan Nothaas, and Michael Schoettner</i>	
OmniGraph: A Scalable Hardware Accelerator for Graph Processing	623
<i>Chongchong Xu, Chao Wang, Lei Gong, Yuntao Lu, Fan Sun, Yiwei Zhang, Xi Li, and Xuehai Zhou</i>	
A Comparative Study of HDD and SSD RAIDs' Impact on Server Energy Consumption	625
<i>Erica Tomes and Nihat Altiparmak</i>	
Analyzing Hybrid Transactional Memory Performance Using Intel SDE	627
<i>Mohammad A. Qayum, Abdel-Hameed A. Badawy, and Jeanine Cook</i>	
A Power-Efficient Accelerator Based on FPGAs for LSTM Network	629
<i>Yiwei Zhang, Chao Wang, Lei Gong, Yuntao Lu, Fan Sun, Chongchong Xu, Xi Li, and Xuehai Zhou</i>	
A Power-Efficient Accelerator for Convolutional Neural Networks	631
<i>Fan Sun, Chao Wang, Lei Gong, Chongchong Xu, Yiwei Zhang, Yuntao Lu, Xi Li, and Xuehai Zhou</i>	
Investigating the Effect of Garbage Collection on Service Level Objectives of Clouds	633
<i>Panagiotis Patros, Kenneth B. Kent, and Michael Dawson</i>	
Performance of Large-Scale Electronic Structure Calculations on Built-in FPGA Systems	635
<i>Seungmin Lee, Dukyun Nam, and Hoon Ryu</i>	
Efficient Swap Protocol of Remote Memory Paging for Out-of-Core Multi-thread Applications	637
<i>Hiroko Midorikawa, Kenji Kitagawa, and Hikari Ohura</i>	
Evaluating Effect of Write Combining on PCIe Throughput to Improve HPC Interconnect Performance	639
<i>Mahesh Chaudhari, Kedar Kulkarni, Shreeya Badhe, and Vandana Inamdar</i>	
Preliminary Interference Study About Job Placement and Routing Algorithms in the Fat-Tree Topology for HPC Applications	641
<i>Peixin Qiao, Xin Wang, Xu Yang, Yuping Fan, and Zhiling Lan</i>	
A Preliminary Study of Intra-Application Interference on Dragonfly Network	643
<i>Xin Wang, Xu Yang, Misbah Mubarak, Robert B. Ross, and Zhiling Lan</i>	
A New Direction for Streaming Graph Analysis	645
<i>Eisha Nathan, E. Jason Riedy, Anita Zakrajewska, and Chunxing Yin</i>	
SharP Hash: A High-Performing Distributed Hash for Extreme-Scale Systems	647
<i>Zachary W. Parchman, Ferrol Aderholdt, and Manjunath Gorentla Venkata</i>	

AMM: Scalable Memory Reuse Model to Predict the Performance of Physics Codes	649
<i>Gopinath Chennupati, Nandakishore Santhi, Stephan Eidenbenz, and Sunil Thulasidasan</i>	
A Probabilistic Monte Carlo Framework for Branch Prediction	651
<i>Bhargava Kalla, Nandakishore Santhi, Abdel-Hameed A. Badawy, Gopinath Chennupati, and Stephan Eidenbenz</i>	
HPC-Oriented Toolchain for Hardware Simulators	653
<i>Olivier Serres, Engin Kayraklıoglu, and Tarek El-Ghazawi</i>	
Introducing Weirs: An Abstraction for Next Generation Streaming Workflows	655
<i>Erich Lohrmann, Greg Eisenhauer, and Matthew Wolf</i>	
A Comparison of Parallel Graph Processing Implementations	657
<i>Samuel D. Pollard and Boyana Norris</i>	
Co-locating Graph Analytics and HPC Applications	659
<i>Kevin Brown and Satoshi Matsuoka</i>	

Re-Emergence of Vector Architectures Workshop (REV-A)

Accelerating Smith-Waterman Alignment Workload with Scalable Vector Computing	661
<i>Dong-Hyeon Park, Jonathan Beaumont, and Trevor Mudge</i>	
Halide Vectorization for Android Photography Applications — A Case Study	669
<i>Martin Johnson and Daniel Playne</i>	
Preliminary Performance Evaluation of Application Kernels Using ARM SVE with Multiple Vector Lengths	677
<i>Yuetsu Kodama, Tetsuya Odajima, Motohiko Matsuda, Miwako Tsuji, Jinpil Lee, and Mitsuhsisa Sato</i>	
Vectorization-Aware Loop Optimization with User-Defined Code Transformations	685
<i>Hiroyuki Takizawa, Thorsten Reimann, Kazuhiko Komatsu, Takashi Soga, Ryusuke Egawa, Akihiro Musa, and Hiroaki Kobayashi</i>	
Performance and Power Analysis of SX-ACE Using HP-X Benchmark Programs	693
<i>Ryusuke Egawa, Kazuhiko Komatsu, Yoko Isobe, Toshihiro Kato, Souya Fujimoto, Hiroyuki Takizawa, Akihiro Musa, and Hiroaki Kobayashi</i>	
Performance Evaluation of Quantum ESPRESSO on NEC SX-ACE	701
<i>Osamu Watanabe, Akihiro Musa, Hiroaki Hokari, Shivanshu Singh, Raghunandan Mathur, and Hiroaki Kobayashi</i>	

3rd International Workshop on Fault Tolerant Systems (FTS)

Assuming Failure Independence: Are We Right to be Wrong?	709
<i>Guillaume Aupy, Yves Robert, and Frédéric Vivien</i>	
MACORD: Online Adaptive Machine Learning Framework for Silent Error Detection	717
<i>Omer Subasi, Sheng Di, Prasanna Balaprakash, Osman Unsal, Jesus Labarta, Adrian Cristal, Sriram Krishnamoorthy, and Franck Cappello</i>	

cudaCR: An In-Kernel Application-Level Checkpoint/Restart Scheme for CUDA-Enabled GPUs	725
<i>Behnam Pourghassemi and Aparna Chandramowlishwaran</i>	
Application-Based Fault Tolerance Techniques for Fully Protecting Sparse Matrix Solvers	733
<i>Grzegorz Pawelczak, Simon McIntosh-Smith, James Price, and Matt Martineau</i>	
Performance Implications of Failures on MapReduce Applications	741
<i>Mohammad Tanvir Rahman, Edgar Gabriel, and Jaspal Subhlok</i>	
A Malleable and Fault-Tolerant Task Pool Framework for X10	749
<i>Marco Bungart and Claudia Fohry</i>	

Workshop on Monitoring and Analysis for High Performance Computing Systems Plus Applications (HPCMASPA)

Technical Papers

Big Data Meets HPC Log Analytics: Scalable Approach to Understanding Systems at Extreme Scale	758
<i>Byung H. Park, Saurabh Hukerikar, Ryan Adamson, and Christian Engelmann</i>	
Data Mining-Based Analysis of HPC Center Operations	766
<i>Jannis Klinkenberg, Christian Terboven, Stefan Lankes, and Matthias S. Müller</i>	
Job Storage Performance Monitoring on Sonexion with Project Caribou	774
<i>Nathan Schumann and Craig Flaskerud</i>	
LIKWID Monitoring Stack: A Flexible Framework Enabling Job Specific Performance monitoring for the masses	781
<i>Thomas Röhl, Jan Eitzinger, Georg Hager, and Gerhard Wellein</i>	
Monitoring Infrastructure: The Challenges of Moving Beyond Petascale	785
<i>Amanda Bonnie, Mike Mason, and Daniel Illescas</i>	
PFAnalyzer: A Toolset for Analyzing Application-Aware Dynamic Interconnects	789
<i>Keichi Takahashi, Susumu Date, Dashdavaa Khureltulga, Yoshiyuki Kido, and Shinji Shimojo</i>	

Work In Progress Papers

Holistic Measurement-Driven System Assessment	797
<i>Saurabh Jha, Jim Brandt, Ann Gentile, Zbigniew Kalbarczyk, Greg Bauer, Jeremy Enos, Michael Showerman, Larry Kaplan, Brett Bode, Annette Greiner, Amanda Bonnie, Mike Mason, Ravishankar K. Iyer, and William Kramer</i>	
lo2s — Multi-core System and Application Performance Analysis for Linux	801
<i>Thomas Ilsche, Robert Schöne, Mario Bielert, Andreas Gocht, and Daniel Hackenberg</i>	

Measuring Minimum Switch Port Metric Retrieval Time and Impact for Multi-layer InfiniBand Fabrics	805
<i>Michael Aguilar, Benjamin A. Allan, and Sergei Polevitzky</i>	
Understanding Performance Variability on the Aries Dragonfly Network	809
<i>Taylor Groves, Yizi Gu, and Nicholas J. Wright</i>	
YAViT (Yet Another Viz Tool): Raising the Level of Abstraction in End-User HPC Interactions	814
<i>Omar Aaziz, Ujjwal Panthi, and Jonathan Cook</i>	
Second International Workshop on Representative Applications (WRAp)	
Assessing Representativeness of Kernels Using Descriptive Statistics	818
<i>Youngsung Kim, John M. Dennis, and Christopher Kerr</i>	
A Performance Projection of Mini-Applications onto Benchmarks Toward the Performance Projection of Real-Applications	826
<i>Miwako Tsuji, William T. C. Kramer, and Mitsuhsisa Sato</i>	
Achieving Performance Portability for a Heat Conduction Solver Mini-Application on Modern Multi-core Systems	834
<i>Richard O. Kirk, Gihan R. Mudalige, Istvan Z. Reguly, Steven A. Wright, Matt J. Martineau, and Stephen A. Jarvis</i>	
TeaLeaf: A Mini-Application to Enable Design-Space Explorations for Iterative Sparse Linear Solvers	842
<i>Simon McIntosh-Smith, Matthew Martineau, Tom Deakin, Grzegorz Pawelczak, Wayne Gaudin, Paul Garrett, Wei Liu, Richard Smedley-Stevenson, and David Beckingsale</i>	
The Arch Project: Physics Mini-Apps for Algorithmic Exploration and Evaluating Programming Environments on HPC Architectures	850
<i>Matthew Martineau and Simon McIntosh-Smith</i>	
Thoughtful Precision in Mini-Apps	858
<i>Shane Fogerty, Siddhartha Bishnu, Yuliana Zamora, Laura Monroe, Steve Poole, Michael Lam, Joe Schoonover, and Robert Robey</i>	
Quicksilver: A Proxy App for the Monte Carlo Transport Code Mercury	866
<i>David F. Richards, Ryan C. Bleile, Patrick S. Brantley, Shawn A. Dawson, Michael Scott McKinley, and Matthew J. O'Brien</i>	
Pushing the Limits of Irregular Access Patterns on Emerging Network Architecture: A Case Study	874
<i>Roberto Gioiosa, Thomas Warfel, Antonino Tumeo, and Ryan Friese</i>	
Author Index	882