2015 IEEE International Conference on Cluster Computing (Cluster 2015)

Chicago, Illinois, USA
8 – 11 September 2015

IEEE Catalog Number: CFP15235-POD
ISBN: 978-1-4673-6599-4

2015 IEEE International Conference on Cluster Computing

CLUSTER 2015

Table of Contents

CLUSTER 2015 General Co-Chairs Welcome Message xviii

CLUSTER 2015 Message from the Program Chair xxi

CLUSTER 2015 Organizing Committee xxiii

CLUSTER 2015 Program Committee xxiv

CLUSTER 2015 External Reviewers xxvii

CLUSTER 2015 Sponsors xxviii

HUCAA 2015 Workshop Welcome Message xxix

HUCAA 2015 Program Committee xxx

HPCMASPA 2015 Workshop Welcome Message xxxi

HPCMASPA 2015 Program Committee xxxii

HiPINEB 2015 Workshop Welcome Message xxxiii

HiPINEB 2015 Organizing Committees xxxiv

Campus Bridging Workshop Overview xxxv

Session 1: Best Paper Candidates

Parallel Modularity-Based Community Detection on Large-Scale Graphs 1
Jianping Zeng and Hongfeng Yu

Optimizing Explicit Hydrodynamics for Power, Energy, and Performance 11
Edgar A. León, Ian Karlin, and Ryan E. Grant

Machines Tuning Machines: Configuring Distributed Stream Processors with Bayesian Optimization 22
Lorenz Fischer, Shen Gao, and Abraham Bernstein

Workload-Aware Resource Reservation for Multi-tenant NoSQL 32
Jiaan Zeng and Beth Plale

Session 2: Task Parallel Computing

Automatic Command Queue Scheduling for Task-Parallel Workloads in OpenCL 42
Ashwin Mandayam Aji, Antonio J. Peña, Pavan Balaji, and Wu-chun Feng

Session 3: Big Data Processing

Taming Non-local Stragglers Using Efficient Prefetching in MapReduce 52
Ze Yu, Min Li, Xin Yang, Han Zhao, and Xiaolin Li

IOSIG+: On the Role of I/O Tracing and Analysis for Hadoop Systems 62
Bo Feng, Xi Yang, Kun Feng, Yanlong Yin, and Xian-He Sun

Exploring Memory Hierarchy to Improve Scientific Data Read Performance 66
Wenzhao Zhang, Houjun Tang, Xiaocheng Zou, Steven Harenberg, Qing Liu, Scott Klasky, and Nagiza F. Samatova

A Case Study of Optimizing Big Data Analytical Stacks Using Structured Data Shuffling 70
Dixin Tang, Taoying Liu, Rubao Lee, Hong Liu, and Wei Li

An SSD-HDD Integrated Storage Architecture for Write-Once-Read-Once Applications on Clusters 74
Cailiang Xu, Wei Wang, Deng Zhou, and Tao Xie

Session 4: GPU Computing

Exploiting GPUDirect RDMA in Designing High Performance OpenSHMEM for NVIDIA GPU Clusters 78
Khaled Hamidouche, Akshay Venkatesh, Ammar Ahmad Awan, Hari Subramoni, Ching-Hsiang Chu, and Dhabaleswar K. (DK) Panda

Improving Strong-Scaling on GPU Cluster Based on Tightly Coupled Accelerators Architecture 88
Toshihiro Hanawa, Hisafumi Fujii, Norihisa Fujita, Tetsuya Odajima, Kazuya Matsumoto, Yuetsu Kodama, and Taisuke Boku

Exploring the Suitability of Remote GPGPU Virtualization for the OpenACC Programming Model Using rCUDA 92
Adrián Castelló, Antonio J. Peña, Rafael Mayo, Pavan Balaji, and Enrique S. Quintana-Ortí

PLB-HeC: A Profile-Based Load-Balancing Algorithm for Heterogeneous CPU-GPU Clusters 96
Luis Sant'Ana, Daniel Cordeiro, and Raphael Camargo

A TSQR Based Krylov Basis Computation Method on Hybrid GPU Cluster 106
Langshi Chen and Serge Petiton

Session 5: Machine Learning and Data Mining

Fast and Accurate Support Vector Machines on Large Scale Systems 110
Abhinav Vishnu, Jeyanthi Narasimhan, Lawrence Holder, Darren Kerbyson, and Adolfy Hoisie

A Machine-Learning Approach for Communication Prediction of Large-Scale Applications 120
Nikela Papadopoulou, Georgios Goumas, and Nectarios Koziris

An Efficient Parallel Approach of Parsing and Indexing for Large-Scale XML Datasets 124
Song Kunfang and Lu Hongwei

Collective I/O Tuning Using Analytical and Machine Learning Models 128
Florin Isaila, Prasanna Balaprakash, Stefan M. Wild, Dries Kimpe, Rob Latham, Rob Ross, and Paul Hovland

Large Scale Frequent Pattern Mining Using MPI One-Sided Model 138
Abhinav Vishnu and Khushbu Agarwal

Session 6: Resilience and Reliability

Fast Fault Injection and Sensitivity Analysis for Collective Communications 148
Kun Feng, Manjunath Gorentla Venkata, Dong Li, and Xian-He Sun

A Practical Approach for Handling Soft Errors in Iterative Applications 158
Jiaqi Liu, Mehmet Can Kurt, and Gagan Agrawal

Ensuring Data Durability with Increasingly Interdependent Content 162
Veronica Estrada Galinanes and Pascal Felber

On the Need for Reproducible Numerical Accuracy through Intelligent Runtime Selection of Reduction Algorithms at the Extreme Scale 166
Dylan Chapp, Travis Johnston, and Michela Taufer

Towards Building Resilient Scientific Applications: Resilience Analysis on the Impact of Soft Error and Transient Error Tolerance with the CLAMR Hydrodynamics Mini-App 176
Qiang Guan, Nathan DeBardeleben, Brian Artkinson, Robert Robey, and WIlliam M. Jones

DINO: Divergent Node Cloning for Sustained Redundancy in HPC 180
Arash Rezaei and Frank Mueller

Session 7: High Performance I/O

Dynamic Model-Driven Parallel I/O Performance Tuning 184
Babak Behzad, Surendra Byna, Stefan M. Wild, Prabhat, and Marc Snir

TRIO: Burst Buffer Based I/O Orchestration 194
Teng Wang, Sarp Oral, Michael Pritchard, Bin Wang, and Weikuan Yu

BPS: A Balanced Partial Stripe Write Scheme to Improve the Write Performance of RAID-6 204
Congjin Du, Chentao Wu, Jie Li, Minyi Guo, and Xubin He

RDMA-Based Direct Transfer of File Data to Remote Page Cache 214
Shin Sasaki, Kazushi Takahashi, Yoshihiro Oyama, and Osamu Tatebe

Session 8: MPI

High Performance MPI Datatype Support with User-Mode Memory Registration: Challenges, Designs, and Benefits 226
Mingzhe Li, Hari Subramoni, Khaled Hamidouche, Xiaoyi Lu, and Dhabaleswar K. (DK) Panda

Session 9: Distributed Data Processing

Overcoming Hadoop Scaling Limitations through Distributed Task Execution 236
Ke Wang, Ning Liu, Iman Sadooghi, Xi Yang, Xiaobing Zhou, Tonglin Li, Michael Lang, Xian-He Sun, and Ioan Raicu

SideWalk: A Facility of Lightweight Out-of-Band Communications for Augmenting Distributed Data Processing Flows 246
Yin Huai, Yuan Yuan, Rubao Lee, and Xiaodong Zhang

High-Performance, Distributed Dictionary Encoding of RDF Datasets 250
Alessandro Morari, Jesse Weaver, Oreste Villa, David Haglin, Antonino Tumeo, Vito Giovanni Castellana, and John Feo

I/O-Aware Batch Scheduling for Petascale Computing Systems 254
Zhou Zhou, Xu Yang, Dongfang Zhao, Paul Rich, Wei Tang, Jia Wang, and Zhiling Lan

Session 10: Energy Efficiency

Performance-to-Power Ratio Aware Virtual Machine (VM) Allocation in Energy-Efficient Clouds 264
Xiaojun Ruan and Haiquan Chen

A Workload-Aware Energy Model for Virtual Machine Migration 274
Vincenzo De Maio, Gabor Kecskemeti, and Radu Prodan

Session 11: Graph Processing

GraphTrek: Asynchronous Graph Traversal for Property Graph-Based Metadata Management 284
Dong Dai, Philip Carns, Robert B. Ross, John Jenkins, Kyle Blauer, and Yong Chen

Towards Multi-site Metadata Management for Geographically Distributed Cloud Workflows 294
Luis Pineda-Morales, Alexandru Costan, and Gabriel Antoniu

Session 12: Application Acceleration

PaRSEC in Practice: Optimizing a Legacy Chemistry Application through Distributed Task-Based Execution 304
Anthony Danalis, Heike Jagode, George Bosilca, and Jack Dongarra

Optimizing I/O for Petascale Seismic Simulations on Unstructured Meshes 314
Sebastian Rettenberger and Michael Bader

LBM-HPC - An Open-Source Tool for Fluid Simulations. Case Study: Unified Parallel C (UPC-PGAS) 318
Pedro Valero-Lara and Johan Jansson

Scaling Data Intensive Physics Applications to 10k Cores on Non-dedicated Clusters with Lobster 322
Anna Woodard, Matthias Wolf, Charles Mueller, Nil Valls, Ben Tovar, Patrick Donnelly, Peter Ivie, Kenyi Hurtado Anampa, Paul Brenner, Douglas Thain, Kevin Lannon, and Michael Hildreth

RE-PAGE: Domain-Specific REplication and PArallel Processing of GEnomic Data 332
Mucahid Kutlu and Gagan Agrawal

Session 13: Network and High Performance Communication

Re-evaluating Network Onload vs. Offload for the Many-Core Era 342
Matthew G.F. Dosanjh, Ryan E. Grant, Patrick G. Bridges, and Ron Brightwell

Fast Calculation of Max-Min Fair Rates for Multi-commodity Flows in Fat-Tree Networks 351
Md Atiqul Mollah, Xin Yuan, Scott Pakin, and Michael Lang

Comparing Global Link Arrangements for Dragonfly Networks 361
Emily Hastings, David Rincon-Cruz, Marc Spehlmann, Sofia Meyers, Anda Xu, David P. Bunde, and Vitus J. Leung

Towards the InfiniBand SR-IOV vSwitch Architecture 371
Evangelos Tasoulas, Ernst Gunnar Gran, Bjørn Dag Johnsen, Kyrre Begnum, and Tor Skeie

Session 14: Parallel Algorithms

MPC: A Massively Parallel Compression Algorithm for Scientific Data 381
Annie Yang, Hari Mukka, Farbod Hesaaraki, and Martin Burtscher

Balancing Thread-Level and Task-Level Parallelism for Data-Intensive Workloads on Clusters and Clouds 390
Olivia Choudhury, Dinesh Rajan, Nicholas Hazekamp, Sandra Gesing, Douglas Thain, and Scott Emrich

LU Factorization: Towards Hiding Communication Overheads with a Lookahead-Free Algorithm 394
Tan Nguyen and Scott B. Baden

Distributed-Memory Algorithms for Maximal Cardinality Matching Using Matrix Algebra 398
Ariful Azad and Aydin Buluç

Session 15: Task and Process Scheduling

The Cost of Synchronizing Imbalanced Processes in Message Passing Systems 408
Ivy Bo Peng, Stefano Markidis, and Erwin Laure

An Approach to Selecting Thread + Process Mixes for Hybrid MPI + OpenMP Applications 418
Hormozd Gahvari, Martin Schulz, and Ulrike Meier Yang

On the Application Task Granularity and the Interplay with the Scheduling Overhead in Many-Core Shared Memory Systems 428
Dana Akhmetova, Gokcen Kestor, Roberto Gioiosa, Stefano Markidis, and Erwin Laure

Session 16: PGAS and Shared Memory Programming

OpenSHMEM as a Portable Communication Layer for PGAS Models: A Case Study with Coarray Fortran 438
Naveen Namashivayam, Deepak Eachempati, Dounia Khaldi, and Barbara Chapman

A Team-Based Methodology of Memory Hierarchy-Aware Runtime Support in Coarray Fortran 448
Dounia Khaldi, Deepak Eachempati, Shiyao Ge, Pierre Jouvelot, and Barbara Chapman

Optimizing Caching DSM for Distributed Software Speculation 452
Sai Charan Koduru, Keval Vora, and Rajiv Gupta

Empirical Comparison of Three Versioning Architectures 456
Hajime Fujita, Kamil Iskra, Pavan Balaji, and Andrew A. Chien

Detecting Thread-Safety Violations in Hybrid OpenMP/MPI Programs 460
Hongyi Ma, Liqiang Wang, and Krishanthan Krishnamoorthy

Session 17: Cluster Tools

Toward Rapid Understanding of Production HPC Applications and Systems 464
Anthony Agelastos, Benjamin Allan, Jim Brandt, Ann Gentile, Sophia Lefantzi, Steve Monk, Jeff Ogden, Mahesh Rajan, and Joel Stevenson

ObsCon: Integrated Monitoring and Control for Parallel, Real-Time Applications 474
Alan Nussbaum, Shwetha Mathangi Chandra Choodamani, and Karsten Schwan

VecMeter: Measuring Vectorization on the Xeon Phi 478
Joshua Peraza, Ananta Tiwari, William A. Ward, Jr., Roy Campbell, and Laura Carrington

Toward Interlanguage Parallel Scripting for Distributed-Memory Scientific Computing 482
Justin M. Wozniak, Timothy G. Armstrong, Ketan C. Maheshwari, Daniel S. Katz, Michael Wilde, and Ian T. Foster

Poster Papers

A Cache Management Scheme for Hiding Garbage Collection Latency in Flash-Based Solid State Drives 486
Wei Xie and Yong Chen

A Performance Comparison of CUDA Remote GPU Virtualization Frameworks 488
Carlos Reaño and Federico Silla

A Two-Tiered Approach to I/O Quality of Service in Docker Containers 490
Sean McDaniel, Stephen Herbein, and Michela Taufer

Accelerating Laue Depth Reconstruction Algorithm with CUDA 492
Ke Yue, Schwarz Nicholas, and Tischler Jonathan Z.

An FPGA-Based Accelerator for Neighborhood-Based Collaborative Filtering Recommendation Algorithms 494
Xiang Ma, Chao Wang, Qi Yu, Xi Li, and Xuehai Zhou

Can Cloud Service Get His Family? A Step Towards Service Family Detecting 496
Xinkui Zhao, Jianwei Yin, Chen Zhi, Pengxiang Lin, and Zuoning Chen

Design a Hash-Based Control Mechanism in vSwitch for Software-Defined Networking Environment 498
Shih-Wen Hsu, Tseng-Yi Chen, Yun-Chun Chang, Shuo-Han Chen, Han-Chieh Chao, Tsen-Yeh Lin, and Wei-Kuan Shih

Development of MapReduce and MPI Programs for Motif Search 500
Mejdl Safran, Saad Al-qahtani, Michelle Zhu, and Dunren Che

Distributed Modular Monitoring (DiMMon) Approach to Supercomputer Monitoring 502
Konstantin Stefanov and Vladimir Voevodin

Efficient Distributed Data Clustering on Spark 504
Jia Li, Dongsheng Li, and Yiming Zhang

Energy-Aware Job Management Approaches for Workflow in Cloud 506
Mustafa Khaleel and Michelle M. Zhu

Evaluating R-Based Big Data Analytic Frameworks 508
Mei Liang, Cesar Trejo, Lavanya Muthu, Linh B. Ngo, Andre Luckow, and Amy W. Apon

Evaluating the Support of MTC Applications on Intel Xeon Phi Many-Core Accelerators 510
Poornima Nookala, Serapheim Dimitropoulos, Karl Stough, and Ioan Raicu

Flexible Error Recovery Using Versions in Global View Resilience 512
Nan Dun, Hajime Fujita, Aiman Fang, Yan Liu, Andrew A. Chien, Pavan Balaj, Kamil Iskra, Wesley Bland, and Andrew Siegel

GO-Docker: A Batch Scheduling System with Docker Containers 514
Olivier Sallou and Cyril Monjeaud

GRAPH/Z: A Key-Value Store Based Scalable Graph Processing System 516
Tonglin Li, Chaoqi Ma, Jiabao Li, Xiaobing Zhou, Ke Wang, Dongfang Zhao, Iman Sadooghi, and Ioan Raicu

Highly Scalable Parallel Search-Tree Algorithms: The Virtual Topology Approach 518
Faisal N. Abu-Khzam, Amer E. Mouawad, and Karim A. Jahed

HRDBMS: A NewSQL Database for Analytics 519
Jason Arnold, Boris Glavic, and Ioan Raicu

Minimizing Data Transmission Latency by Bipartite Graph in MapReduce 521
Jie Wei, Shangguang Wang, Lingyan Zhang, Ao Zhou, Qibo Sun, Ruisheng Shi, and Fangchun Yang

monBench: A Database Performance Benchmark for Cloud Monitoring System 523
Xinkui Zhao, Jianwei Yin, Chen Zhi, Pengxiang Lin, Shichun Feng, Hao Wu, and Zuoning Chen

Mutated Near Optimal Vertex Cover Algorithm (NOVCA) Visualization on a Tile Display 525
Sanjaya Gajurel and Roger Bielefeld

Network Quality of Service in Docker Containers 527
Ayush Dusia, Yang Yang, and Michela Taufer

Pallas: An Application-Driven Task and Network Simulation Framework 529
Yuming Ye, Ziyang Li, Dongsheng Li, Yiming Zhang, Feng Liu, and Yuxing Peng

Peer Comparison of XSEDE and NCAR Publication Data 531
Gregor von Laszewski, Fugang Wang, Geoffrey C. Fox, David L. Hart, Thomas R. Furlani, Robert L. DeLeon, and Steven M. Gallo

Performance of the NVIDIA Jetson TK1 in HPC 533
Yash Ukidave, David Kaeli, Umesh Gupta, and Kurt Keville.

Dynamic CPU Resource Allocation in Containerized Cloud Environments 535
Jose Monsalve, Aaron Landwehr, and Michela Taufer

Toward Auto-tuned Krylov Basis Computation for Different Sparse Matrix Formats and Interconnects on GPU Clusters 537
Langshi Chen and Serge Petition

Towards Building a Lightweight Key-Value Store on Parallel File System 539
Jiaan Zeng and Beth Plale

Understanding the Propagation of Error Due to a Silent Data Corruption in a Sparse Matrix Vector Multiply 541
Jon Calhoun, Marc Snir, Luke Olson, and Maria Garzaran

FTS 2015

Exploiting Spatial Information in Datasets to Enable Fault Tolerant Sparse Matrix Solvers 543
Rob Hunt and Simon McIntosh-Smith

Partial Differential Equations Preconditioner Resilient to Soft and Hard Faults 552
Francesco Rizzi, Karla Morris, Khachik Sargsyan, Paul Mycek, Cosmin Safta, Olivier LeMaitre, Omar Knio, and Bert Debusschere

Fault-Tolerant Protocol for Hybrid Task-Parallel Message-Passing Applications 563
Tatiana Martsinkevich, Omer Subasi, Osman Unsal, Franck Cappello, and Jesus Labarta

Programmer-Guided Reliability for Extreme-Scale Applications 571
David E. Bernholdt, Wael R. Elwasif, Christos Kartsaklis, Seyong Lee, and Tiffany M. Mintz

Building a Fault Tolerant Application Using the GASPI Communication Layer 580
Faisal Shahzad, Moritz Kreutzer, Thomas Zeiser, Rui Machado, Andreas Pieper, Georg Hager, and Gerhard Wellein

Stay Alive, Don't Give Up: DUE and SDC Reduction with Memory Repair 588
Dong Wan Kim and Mattan Erez

Detecting and Correcting Data Corruption in Stencil Applications through Multivariate Interpolation 595
Leonardo Arturo Bautista Gomez and Franck Cappello

HUCAA 2015

Pairwise Sequence Alignment with Gaps with GPU 603
Thomas C. Carroll, Jude-Thaddeus Ojiaku, and Prudence W.H. Wong

Scalable Relativistic High-Resolution Shock-Capturing for Heterogeneous Computing 611
Forrest Wolfgang Glines, Matthew Anderson, and David Neilsen

On the Execution of Computationally Intensive CPU-Based Libraries on Remote Accelerators for Increasing Performance: Early Experience with the OpenBLAS and FFTW Libraries 619
Santiago Mislata Valero and Federico Silla

Hybrid Communication with TCA and InfiniBand on a Parallel Programming Language XcalableACC for GPU Clusters 627
Tetsuya Odajima, Taisuke Boku, Toshihiro Hanawa, Hitoshi Murai, Masahiro Nakao, Akihiro Tabuchi, and Mitsuhisa Sato

Evaluation of FFT for GPU Cluster Using Tightly Coupled Accelerators Architecture 635
Toshihiro Hanawa, Hisafumi Fujii, Norihisa Fujita, Tetsuya Odajima, Kazuya Matsumoto, and Taisuke Boku

HPCMASPA 2015

Analysis of XDMoD/SUPReMM Data Using Machine Learning Techniques 642
Steven M. Gallo, Joseph P. White, Robert L. DeLeon, Thomas R. Furlani, Helen Ngo, Abani K. Patra, Matthew D. Jones, Jeffrey T. Palmer, Nikolay Simakov, Jeanette M. Sperhac, Martins Innus, Thomas Yearke, and Ryan Rathsam

Practical Resource Monitoring for Robust High Throughput Computing 650
Gideon Juve, Benjamin Tovar, Rafael Ferreira da Silva, Dariusz Król, Douglas Thain, Ewa Deelman, William Allcock, and Miron Livny

New Systems, New Behaviors, New Patterns: Monitoring Insights from System Standup 658
Jim Brandt, Ann Gentile, Cindy Martin, Jason Repik, and Narate Taerat

A LogP Extension for Modeling Tree Aggregation Networks 666
Taylor Groves, Samuel K. Gutierrez, and Dorian Arnold

Push Me Pull You: Integrating Opposing Data Transport Modes for Efficient HPC Application Monitoring 674
Omar Aaziz, Jonathan Cook, and Hadi Sharifi

The Performance Implication of Task Size for Applications on the HPX Runtime System 682
Patricia Grubel, Hartmut Kaiser, Jeanine Cook, and Adrian Serio

Comparison of Vendor Supplied Environmental Data Collection Mechanisms 690
Sean Wallace, Venkatram Vishwanath, Susan Coghlan, Zhiling Lan, and Michael E. Papka

WattProf: A Flexible Platform for Fine-Grained HPC Power Profiling 698
Mohammad Rashti, Gerald Sabin, David Vansickle, and Boyana Norris

Real Time Visualization of Monitoring Data for Large Scale HPC Systems 706
Michael Showerman

Evolution of Monitoring over the Lifetime of a High Performance Computing Cluster 710
Adam DeConinck and Kathleen Kelly

Monitoring High Performance Computing Systems for the End User 714
Christopher Lee Moore, Prabhu Singh Khalsa, Todd Alan Yilk, and Michael Mason

Extending LDMS to Enable Performance Monitoring in Multi-core Applications 717
Steven Feldman, Deli Zhang, Damian Dechev, and James Brandt

WRAp 2015

Performance Evaluation of Unstructured Mesh Physics on Advanced Architectures 721
Charles R. Ferenbaugh

Expressing Parallelism on Many-Core for Deterministic Discrete Ordinates Transport 729
Tom Deakin, Simon McIntosh-Smith, and Wayne Gaudin

Design and Development of Domain Specific Active Libraries with Proxy Applications 738
Istvan Zoltan Reguly, Gihan R. Mudalige, and Michael B. Giles

Enabling Tractable Exploration of the Performance of Adaptive Mesh Refinement 746
Courtenay T. Vaughan and Richard F. Barrett

Developing MiniApps on Modern Platforms Using Multiple Programming Models 753
O.E.B. Messer, E. D'Azevedo, J. Hill, W. Joubert, S. Laosooksathit, and A. Tharrington

Evaluation of Parallel Communication Models in Nekbone, a Nek5000 Mini-Application 760
Ilya Ivanov, Jing Gong, Dana Akhmetova, Ivy Bo Peng, Stefano Markidis, Erwin Laure, Rui Machado, Mirko Rahn, Valeria Bartsch, Alistair Hart, and Paul Fischer

Mini-App Driven Optimisation of Inertial Confinement Fusion Codes 768
R.F. Bird, P. Gillies, M.R. Bareford, J.A. Herdman, and S.A. Jarvis

Introducing and Exploiting Hierarchical Structural Information 777
Daniel Rubio Bonilla, Colin W. Glass, Jan Kuper, and Robert de Groote

CMT-bone: A Mini-App for Compressible Multiphase Turbulence Simulation Software 785
Nalini Kumar, Mrugesh Sringarpure, Tania Banerjee, Jason Hackl, S. Balachandar, Herman Lam, Alan George, and Sanjay Ranka

HiPINEB 2015

Fault-Tolerant Routing for Exascale Supercomputer: The BXI Routing Architecture 793
Pierre Vignéras and Jean-Noël Quintin

Throughput Unfairness in Dragonfly Networks under Realistic Traffic Patterns 801
Pablo Fuentes, Enrique Vallejo, Cristóbal Camarero, Ramón Beivide, and Mateo Valero

Modeling a Large Data-Acquisition Network in a Simulation Framework 809
Tommaso Colombo, Holger Fröning, Pedro Javier García, and Wainer Vandelli

Efficient Queuing Schemes for HoL-Blocking Reduction in Dragonfly Topologies with Minimal-Path Routing 817
Pedro Yébenes, Jesus Escudero-Sahuquillo, Pedro J. García, and Francisco J. Quiles

InfiniBand Verbs Optimizations for Remote GPU Virtualization 825
Carlos Reaño and Federico Silla

Multipath Load Balancing for M × N Communication Patterns on the Blue Gene/Q Supercomputer Interconnection Network 833
Huy Bui, Robert Jacob, Preeti Malakar, Venkatram Viswanath, Andrew Johnson, Micheal E. Papka, and Jason Leigh

VEF Traces: A Framework for Modelling MPI Traffic in Interconnection Network Simulators 841
Franisco J. Andújar, Juan A. Villar, José L. Sánchez, Francisco J. Alfaro, and Jesus Escudero-Sahuquillo

SlimUpdate: Minimal Routing Update for Performance-Based Reconfigurations in Fat-Trees 849
Feroz Zahid, Ernst Gunnar Gran, Bartosz Bogdanski, Bjørn Dag Johnsen, and Tor Skeie

Campus Bridging: Reducing Obstacles on the Path to Big Answers

XCBC and XNIT - Tools for Cluster Implementation and Management in Research and Training 857
Jeremy Fischer, Eric Coulter, Richard Knepper, Charles Peck, and Craig A. Stewart

Building Bridges from the Campus to XSEDE 865
Lee Liming, Ian Foster, and Steven Tuecke

Author Index 869