

2009 IEEE International Conference on Cluster Computing and Workshops

(CLUSTER 2009)

**New Orleans, Louisiana, USA
31 August – 4 September 2009**



**IEEE Catalog Number: CFP09235-PRT
ISBN: 978-1-4244-5011-4**

TABLE OF CONTENTS

Improving I/O Performance Using Soft-QoS-Based Dynamic Storage Cache Partitioning	1
<i>C.M. Patrick, R. Garg, S.W. Son, M. Kandemir</i>	
Analytical Modeling and Optimization for Affinity Based Thread Scheduling on Multicore Systems	11
<i>F. Song, S. Moore, J. Dongarra</i>	
Finding a Tradeoff Between Host Interrupt Load and MPI Latency Over Ethernet	21
<i>B. Goglin, N. Furmento</i>	
Comparing Map-Reduce and FREERIDE for Data-Intensive Applications	30
<i>W. Jiang, V.T. Ravi, G. Agrawal</i>	
Computing Median Values in a Cloud Environment Using GridBatch and MapReduce	40
<i>H. Liu, D. Orban</i>	
A Global Operating System for HPC Clusters	50
<i>E. Betti, M. Cesati, R. Gioiosa, F. Piermaria</i>	
On-Demand File Staging System for Linux Clusters	60
<i>A. Hori, Y. Kamoshida, H. Matsuba, K. Ohta, T. Yasui, S. Sumimoto, Y. Ishikawa</i>	
Utility Analysis for Internet-Oriented Server Consolidation in VM-Based Data Centers	70
<i>Y. Song, Y. Zhang, Y. Sun, W. Shi</i>	
Using Clouds for Metagenomics: A Case Study	80
<i>J. Wilkening, A. Wilke, N. Desai, F. Meyer</i>	
Scalable I/O Forwarding Framework for High-Performance Computing Systems	86
<i>N. Ali, P. Carns, K. Iskra, D. Kimpe, S. Lang, R. Latham, R. Ross, L. Ward, P. Sadayappan</i>	
Cluster Fault-Tolerance: An Experimental Evaluation of Checkpointing and MapReduce through Simulation	96
<i>T.C. Bressoud, M.A. Kozuch</i>	
Batch Queue Resource Scheduling for Workflow Applications	106
<i>Y. Zhang, C. Koelbel, K. Cooper</i>	
VCCP: A Transparent, Coordinated Checkpointing System for Virtualization-based Cluster Computing	116
<i>H. Ong, N. Saragol, K. Chanchio, C. Leangsuksun</i>	
How Reliable are Parallel Disk Systems When Energy-Saving Schemes are Involved?	126
<i>S. Yin, X. Ruan, A. Manzanares, X. Qin</i>	
On-line Automatic Resource Selection in Distributed Computing	135
<i>H. Liu, S. Sorensen, A. Nazir</i>	
Power-Aware Scheduling of Virtual Machines in DVFS-enabled Clusters	144
<i>G. Von Laszewski, L. Wang, A.J. Younge, X. He</i>	
Topics on Measuring Real Power Usage on High Performance Computing Platforms	154
<i>J.H. Laros III, K.T. Pedretti, S.M. Kelly, J.P. Vandyke, K.B. Ferreira, C.T. Vaughan, M. Swan</i>	
Using a Cluster as a Memory Resource: A Fast and Large Virtual Memory on MPI	162
<i>H. Midorikawa, K. Saito, M. Sato, T. Boku</i>	
Reliability-Aware Scalability Models for High Performance Computing	172
<i>Z. Zheng, Z. Lan</i>	
SJMR:Parallelizing Spatial Join with MapReduce on Clusters	181
<i>S. Zhang, J. Han, Z. Liu, K. Wang, Z. Xu</i>	
Data Mining Analysis to Validate Performance Tuning Practices for HPL	189
<i>T.Z. Tan, R.S.M. Goh, V. March, S. See</i>	
Integrating Software Distributed Shared Memory and Message Passing Programming	197
<i>H.J. Wong, A.P. Rendell</i>	
Overlapping Computation and Communication in SMT Clusters with Commodity Interconnects	207
<i>G. Goumas, N. Anastopoulos, N. Koziris, N. Ioannou</i>	
HPMR: Prefetching and Pre-shuffling in Shared MapReduce Computation Environment	217
<i>S. Seo, I. Jang, K. Woo, I. Kim, J. Kim, S. Maeng</i>	
DCR: A Fully Transparent Checkpoint/Restart Framework for Distributed Systems	225
<i>C. Ma, Z. Huo, J. Cai, D. Meng</i>	
A Scalable and Generic Task Scheduling System for Communication Libraries	235
<i>F. Trahay, A. Denis</i>	
Live Virtual Machine Migration with Adaptive Memory Compression	243
<i>H. Jin, L. Deng, S. Wu, X. Shi, X. Pan</i>	

Extending I/O through High Performance Data Services	253
<i>H. Abbasi, J. Lofstead, F. Zheng, K. Schwan, M. Wolf, S. Klasky</i>	
An Efficient Hardware-Software Approach to Network Fault Tolerance with InfiniBand	263
<i>A. Vishnu, M. Krishnan, D.K. Panda</i>	
XCPU2-Distributed Seamless Desktop Extension	272
<i>L. Ionkov, E. Van Hensbergen</i>	
Increasing the Availability Provided by RADIC with Low Overhead	281
<i>G. Santos, L. Fialho, D. Rexachs, E. Luque</i>	
Dynamic Screen Division for Load Balancing the Raycasting of Irregular Data	289
<i>B.B. Labronici, C. Bentes, L.M.A. Drummond, R. Farias</i>	
Numerically Stable, Single-Pass, Parallel Statistics Algorithms	299
<i>J. Bennett, R. Grout, P. Pebay, D. Roe, D. Thompson</i>	
Reducing Network Contention with Mixed Workloads on Modern Multicore Clusters	307
<i>M.J. Koop, M. Luo, D.K. Panda</i>	
MDCSim: A Multi-tier Data Center Simulation Platform	317
<i>S. Lim, B. Sharma, G. Nam, E.K. Kim, C.R. Das</i>	
Granules: A Lightweight, Streaming Runtime for Cloud Computing With Support for Map-Reduce	326
<i>S. Pallickara, J. Ekanayake, G. Fox</i>	
Reasons for a Pessimistic or Optimistic Message Logging Protocol in MPI Uncoordinated Failure Recovery	336
<i>A. Bouteiller, T. Ropars, G. Bosilca, C. Morin, J. Dongarra</i>	
Fault-Aware, Utility-Based Job Scheduling on Blue Gene/P Systems	345
<i>W. Tang, Z. Lan, N. Desai, D. Buettner</i>	
Two-phase Load Distribution for Rendering Large 3D Models on a Graphics Cluster	355
<i>A. Beaudoin, D. Goswami, S. Mudur</i>	
Resource Selection and Allocation for Dynamic Adaptive Computing in Heterogeneous Clusters	365
<i>J.U. Duseelis, E.E. Cauich, R.K. Wang, I.D. Scherson</i>	
MITHRA: Multiple Data Independent Tasks on a Heterogeneous Resource Architecture	374
<i>R. Farivar, A. Verma, E.M. Chan, R.H. Campbell</i>	
A Distributed Device Paradigm for Commodity Applications	384
<i>E. Cauich, J.U. Duseelis, R. Wang, I.D. Scherson</i>	
Design Alternatives for Implementing Fence Synchronization in MPI-2 One-sided Communication for InfiniBand Clusters	394
<i>G. Santhanaraman, T. Gangadharappa, S. Narravula, A. Mamidala, D.K. Panda</i>	
Zone Clusters: A Virtual Cluster Based Upon Solaris Containers	403
<i>E. Roush, Z. Thanga</i>	
An Application Based MPI Message Throughput Benchmark	411
<i>B.W. Barrett, K.S. Hemmert</i>	
GridAtlas - A Grid Application and Resource Configuration Repository and Discovery Service	419
<i>E. Afgan, P. Bangalore, D. Duncan</i>	
Implementing WebGIS on Hadoop: A Case Study of Improving Small File I/O Performance on HDFS	429
<i>X. Liu, J. Han, Y. Zhong, C. Han, X. He</i>	
Coordinating the Use of GPU and CPU for Improving Performance of Compute Intensive Applications	437
<i>G. Teodoro, R. Sachetto, O. Sertel, M.N. Gurcan, W. Meira Jr., U. Catalyurek, R. Ferreira</i>	
Scheduling Strategies for HPC as a Service (HPCaaS)	447
<i>G. Shainer, T. Liu, J. Layton, J. Mora</i>	
Fast-Response Dynamic Routing Balancing for High-Speed Interconnection Networks	453
<i>D. Lugones, D. Franco, E. Luque</i>	
Distributed Control Plane Architecture of Next Generation IP Routers	462
<i>K. Nguyen, B. Jaumard</i>	
RDMA over Ethernet - A Preliminary Study	470
<i>H. Subramoni, P. Lai, M. Luo, D.K. Panda</i>	
Oblivious Routing Schemes in Extended Generalized Fat Tree Networks	479
<i>G. Rodriguez, C. Minkenber, R. Beivide, R.P. Luijten, J. Labarta, M. Valero</i>	
On Optimizing I/O Through InfiniBand RDMA for Commodity Clusters	487
<i>B. Allan, H. Chen, S. Cranford, R. Minnich, D. Rudish, L. Ward</i>	
The Event Data Store and I/O Framework for the ATLAS Experiment at the Large Hadron Collider	493
<i>P. Van Gemmeren, D. Malon</i>	
A Performance Evaluation of Scientific I/O Workloads on Flash-Based SSDs	501
<i>S. Park, K. Shen</i>	

Analyzing Massive Astrophysical Datasets: Can Pig/Hadoop or a Relational DBMS Help?	506
<i>S. Loebman, D. Nunley, Y. Kwon, B. Howe, M. Balazinska, J.P. Gardner</i>	
24/7 Characterization of Petascale I/O Workloads	516
<i>P. Carns, R. Latham, R. Ross, K. Iskra, S. Lang, K. Riley</i>	
Workload-Adaptive Management of Energy-Smart Disk Storage Systems	526
<i>E. Otoo, D. Rotem, S. Tsao</i>	
Interfaces for Coordinated Access in the File System	537
<i>S. Lang, R. Latham, R. Ross, D. Kimpe</i>	
Combining I/O Operations for Multiple Array Variables in Parallel NetCDF	546
<i>K. Gao, W. Liao, A. Choudhary, R. Ross, R. Latham</i>	
Overlapped Checkpointing with Hardware Assist	556
<i>C. Mitchell, J. Nunez, J. Wang</i>	
An Innovative Application Execution Toolkit for Multicluster Grids	566
<i>Z. Yun, Z. Lei, G. Allen, D.S. Katz, T. Kosar, S. Jha, J. Ramanujam</i>	
Cross-Domain Metadata Management in Data Intensive Distributed Computing Environment	570
<i>X. Wang, I. Akturk, T. Kosar</i>	
Two-Phase Conflict Detection for Transactional Memory on Clusters	574
<i>R. Wang, K. Lu, X. Lu</i>	
Accelerating SIFT on Parallel Architectures	577
<i>S. Warn, W. Emenecker, J. Cothren, A. Apon</i>	
HPC Failure Prediction Proficiency Metrics	581
<i>N. Taerat, C. Leangsuksun</i>	
Parallel Implementation of a Tightly Coupled Ablation Prediction Code Using MPI	585
<i>N. Mullenix, A. Povitsky</i>	
WinBioinfTools: Bioinformatics Tools for Windows Cluster	589
<i>H. Mohamed, M. Abouelhoda</i>	
Architecting iSCSI-based I/O Systems for High Performance Computing Clusters	593
<i>J. Liberman, L. Ou, S. Chandok</i>	
A Symmetric O(n log n) Message Distributed Snapshot Algorithm for Large-Scale Systems	597
<i>A.D. Kshemkalyani</i>	
HPP-Controller: An Intra-node Controller Designed for Connecting Heterogeneous CPUs	601
<i>Q. Li, P. Zhang, N. Sun</i>	
An Assessment of Multi-Core for a Performance Prediction Model of Tomographic Reconstruction	605
<i>P. Fritzsche, R. Muresano, D. Rexachs, E. Luque</i>	
Resource Manager with Multi-core Support for Parallel Desktop	609
<i>J.R. Garcia, J.L. Lerida, P. Hernandez</i>	
How SPMD Applications Could Be Efficiently Executed on Multicore Environments?	613
<i>R. Muresano, D. Rexachs, E. Luque</i>	
Parallel Application Signature	617
<i>A. Wong, D. Rexachs, E. Luque</i>	
Improving Metadata Management for Small Files in HDFS	621
<i>G. Mackey, S. Sehrish, J. Wang</i>	
Experiences with Hybrid Clusters	625
<i>D. Jamsek, E. Van Hensbergen</i>	
A Compiler for High Performance Computing With Many-Core Accelerators	629
<i>N. Nakasato, J. Makino</i>	
GPU Clusters for High-Performance Computing	638
<i>V.V. Kindratenko, J.J. Enos, G. Shi, M.T. Showerman, G.W. Arnold, J.E. Stone, J.C. Phillips, W. Hwu</i>	
Message Passing for GPGPU Clusters: cudaMPI	646
<i>O.S. Lawlor</i>	
Performance Analysis of Memory Transfers and GEMM Subroutines on NVIDIA Tesla GPU Cluster	654
<i>V. Allada, T. Benjegerdes, B. Bode</i>	
Investigating the Use of GPU-Accelerated Nodes for SAR Image Formation	663
<i>T.D.R. Hartley, A.R. Fasih, C.A. Berdanier, F. Ozguner, U.V. Catalyurek</i>	
Integrating Web 2.0 Technologies with Scientific Simulation Codes for Real-Time Collaboration	671
<i>G. Allen, F. Loffler, T. Radke, E. Schmetter, E. Seidel</i>	
V-MCS: A Configuration System for Virtual Machines	681
<i>X. Sun, C. Du, H. Zou, Y. Chen, P. Shukla</i>	
Author Index	