

# **2008 IEEE International Conference on Cluster Computing**

**Tsukuba, Japan  
29 September - 1 October 2008**



**IEEE Catalog Number:  
ISBN 13:**

**CFP08235-PRT  
978-1-4244-2639-3**

# Table of Contents

<b>Clouds, Clusters and Manycore: the Revolution Ahead .....</b>	<b>1</b>
<i>Daniel A. Reed</i>	
<b>Reliable Adaptable Network Ram.....</b>	<b>2</b>
<i>Tia Newhall, Daniel Amato and Alexandr Pshenichkin</i>	
<b>Magnet: A Novel Scheduling Policy for Power Reduction in Cluster With Virtual Machines.....</b>	<b>13</b>
<i>Litong Hu, Hai Jin, Xiaofei Liao, Xianjie Xiong and Haikun Liu</i>	
<b>Active Coordination (ACT) – Toward Effectively Managing Virtualized Multicore Clouds .....</b>	<b>23</b>
<i>Mukil Kesavan, Adit Ranadive, Ada Gavrilovska, Karsten Schwan</i>	
<b>Variable-Grain and Dynamic Work Generation for Minimal Unique Itemset Mining .....</b>	<b>33</b>
<i>Paraskevas Yiapanis, David J. Haglin, Anna M. Manning, Ken Mayes and John Keane</i>	
<b>A Large-Grained Parallel Algorithm for Nonlinear Eigenvalue Problems and Its Implementation Using Omnirpc .....</b>	<b>42</b>
<i>Takeshi Amako, Yusaku Yamamoto, Shao-Liang Zhang</i>	
<b>Redistribution Aware Two-Step Scheduling for Mixed-Parallel Applications .....</b>	<b>50</b>
<i>Sascha Hunold, Thomas Rauber and Frédéric Suter</i>	
<b>Implications of Non-Constant Clock Drifts for the Timestamps of Concurrent Events .....</b>	<b>59</b>
<i>Daniel Becker, Rolf Rabenseifner and Felix Wolf</i>	
<b>In Search of Sweet-Spots in Parallel Performance Monitoring .....</b>	<b>69</b>
<i>Aroon Nataraj, Allen Malony, Allen Morris, Dorian C. Arnold, Barton P. Miller</i>	
<b>Workflows for Performance Evaluation and Tuning .....</b>	<b>79</b>
<i>Jeffrey L. Tilson, Mark S.C. Reed and Robert J. Fowler</i>	
<b>Combining Virtual Machine Migration With Process Migration for HPC on Multi-Clusters and Grids.....</b>	<b>89</b>
<i>Tal Maoz, Amnon Barak and Lior Amar</i>	
<b>Live and Incremental Whole-System Migration of Virtual Machines Using Block-Bitmap.....</b>	<b>99</b>
<i>Yingwei Luo, Binbin Zhang, Xiaolin Wang, Zhenlin Wang, Yifeng Sun and Haogang Chen</i>	
<b>Efficient One-Copy MPI Shared Memory Communication in Virtual Machines.....</b>	<b>107</b>
<i>Wei Huang, Matthew J. Koop, Dhableswar K. Panda</i>	
<b>Multistage Switches Are Not Crossbars: Effects of Static Routing in High-Performance Networks.....</b>	<b>116</b>
<i>Torsten Hoefler, Timo Schneider and Andrew Lumsdaine</i>	
<b>A Hypertransport-Based Personal Parallel Computer.....</b>	<b>126</b>
<i>Xiaojun Yang, Fei Chen, Hailiang Cheng and Ninghui Sun</i>	
<b>High Message Rate, NIC-Based Atomics: Design and Performance Considerations .....</b>	<b>133</b>
<i>Keith D. Underwood, Michael Levenhagen, K. Scott Hemmert, Ron Brightwell</i>	
<b>A Dependency-Aware Task-Based Programming Environment for Multi-Core Architectures .....</b>	<b>142</b>
<i>Josep M. Perez, Rosa M. Badia and Jesus Labarta</i>	
<b>Are Nonblocking Networks Really Needed for High-End-Computing Workloads? .....</b>	<b>152</b>
<i>N. Desai, P. Balaji, P. Sadayappan, M. Islam</i>	
<b>OpenMP-Centric Performance Analysis of Hybrid Applications .....</b>	<b>160</b>
<i>Karl Furlinger and Shirley Moore</i>	
<b>DWC<sup>2</sup>: A Dynamic Weight-Based Cooperative Caching Scheme for Object-Based Storage Cluster.....</b>	<b>167</b>
<i>Qingsong Wei, Bharadwaj Veeravalli and Lingfang Zeng</i>	

# Table of Contents

<b>An OSD-Based Approach to Managing Directory Operations in Parallel File Systems</b> .....	175
<i>Nawab Ali, Ananth Devulapalli, Dennis Dalessandro Pete Wyckoff and P. Sadayappan</i>	
<b>Differstore: A Differentiated Storage Service in Object-Based Storage System</b> .....	185
<i>Qingsong Wei</i>	
<b>A Novel Hint-Based I/O Mechanism for Centralized File Server of Cluster</b> .....	194
<i>Huan Chen, Jin Xiong, Ninghui Sun</i>	
<b>Designing Next Generation Clusters With Infiniband and 10GE/IWARP: Opportunities and Challenges</b> .....	202
<i>Dhabaleswar K. Panda</i>	
<b>Scalable MPI Design Over Infiniband Using Extended Reliable Connection</b> .....	203
<i>Matthew J. Koop, Jaidev K. Sridhar, Dhabaleswar K. Panda</i>	
<b>Message Progression in Parallel Computing – to Thread Or Not to Thread?</b> .....	213
<i>Torsten Hoefler, Andrew Lumsdaine</i>	
<b>Improving Message Passing Over Ethernet With I/OAT Copy Offload in Open-MX</b> .....	223
<i>Brice Goglin</i>	
<b>A Trace-Driven Emulation Framework to Predict Scalability of Large Clusters in Presence of OS Jitter</b> .....	232
<i>Pradipta De, Ravi Kothari, Vijay Mann</i>	
<b>Prediction of Behavior of MPI Applications</b> .....	242
<i>Marc Casas, Rosa M. Badia and Jesús Labarta</i>	
<b>A Novel Model for Synthesizing Parallel I/O Workloads in Scientific Applications</b> .....	252
<i>Dan Feng, Qiang Zou, Hong Jiang and Yifeng Zhu</i>	
<b>Divisible Load Scheduling With Improved Asymptotic Optimality</b> .....	262
<i>Reiji Suda</i>	
<b>DLM: A Distributed Large Memory System Using Remote Memory Swapping Over Cluster Nodes</b> .....	268
<i>Hiroko Midorikawa, Motoyoshi Kurokawa, Ryutaro Himeno, Mitsuhsisa Sato</i>	
<b>RI2N: High-Bandwidth and Fault-Tolerant Network With Multi-Link Ethernet for PC Clusters</b> .....	274
<i>Shin'ichi Miura, Takayuki Okamoto, Taisuke Boku, Toshihiro Hanawa, Mitsuhsisa Sato</i>	
<b>Impact of Topology and Link Aggregation on a PC Cluster With Ethernet</b> .....	280
<i>Takafumi Watanabe, Masahiro Nakao, Tomoyuki Hiroyasu, Tomohiro Otsuka, Michihiro Koibuchi</i>	
<b>Predictive Models for Bandwidth Sharing in High Performance Clusters</b> .....	286
<i>Vienne Jerome, Martinasso Maxime, Vincent Jean-Marc, Mehaut Jean-Francois</i>	
<b>Context-Aware Address Translation for High Performance SMP Cluster System</b> .....	292
<i>Moon-Sang Lee, Joonwon Lee, Seungryoul Maeng</i>	
<b>Parallel Multistage Preconditioners by Hierarchical Interface Decomposition on “T2K Open Super Computer (Todai Combined Cluster)” With Hybrid Parallel Programming Models</b> .....	298
<i>Kengo Nakajima</i>	
<b>Load-Balancing Methods for Parallel and Distributed Constraint Solving</b> .....	304
<i>Carl Christian Rolf, Krzysztof Kuchcinski</i>	
<b>Enabling Lock-Free Concurrent Fine-Grain Access to Massive Distributed Data: Application to Supernovae Detection</b> .....	310
<i>Bogdan Nicolae, Gabriel Antoniu, Luc Bouge</i>	
<b>A Multicore-Enabled Multirail Communication Engine</b> .....	316
<i>Elisabeth Brunet, Francois Trahay, Alexandre Denis</i>	

# Table of Contents

<b>Multi-Core Aware Optimization for MPI Collectives .....</b>	<b>322</b>
<i>Bibo Tu, Ming Zou, Jianfeng Zhan, Xiaofang Zhao, Jianping Fan</i>	
<b>Environmental-Aware Optimization of MPI Checkpointing Intervals .....</b>	<b>326</b>
<i>Hideyuki Jitsumoto, Toshio Endo, Satoshi Matsuoka</i>	
<b>Towards an Understanding of the Performance of MPI-IO in Lustre File Systems .....</b>	<b>330</b>
<i>Jeremy Logan, Phillip Dickens</i>	
<b>Gather-Arrange-Scatter: Node-Level Request Reordering for Parallel File Systems on Multi-Core Clusters.....</b>	<b>336</b>
<i>Kazuki Ohta, Hiroya Matsuba, Yutaka Ishikawa</i>	
<b>Empirical-Based Probabilistic Upper Bounds for Urgent Computing Applications.....</b>	<b>342</b>
<i>Nick Trebon, Pete Beckman</i>	
<b>Supporting Storage Resources in Urgent Computing Environments .....</b>	<b>348</b>
<i>Jason Cope, Henry M. Tufo</i>	
<b>Runtime DVFS Control With Instrumented Code in Power-Scalable Cluster System .....</b>	<b>354</b>
<i>Hideaki Kimura, Mitsuhsa Sato, Takayuki Imada, Yoshihiko Hotta</i>	
<b>Intelligent Compilers .....</b>	<b>360</b>
<i>John Cavazos</i>	
<b>Continuous Adaptation for High Performance Throughput Computing Across Distributed Clusters .....</b>	<b>369</b>
<i>Edward Walker</i>	
<b>Performance Models for Dynamic Tuning of Parallel Applications on Computational Grids.....</b>	<b>376</b>
<i>Genaro Costa, Josep Jorba, Anna Morajko, Tomas Margalef, Emilio Luque</i>	
<b>SPRAT: Runtime Processor Selection for Energy-Aware Computing .....</b>	<b>386</b>
<i>Hiroyuki Takizawa, Katuto Sato, Hiroaki Kobayashi</i>	
<b>Using Cluster Computing to Support Automatic and Dynamic Database Clustering.....</b>	<b>394</b>
<i>Sylvain Guinepain, Le Gruenwald</i>	
<b>A Dynamic Programming Approach to Optimizing the Blocking Strategy for the Householder QR Decomposition .....</b>	<b>402</b>
<i>Takeshi Fukaya, Yusaku Yamamoto, Shao-Liang Zhang</i>	
<b>Reinforcement Learning for Automated Performance Tuning: Initial Evaluation for Sparse Matrix Format Selection .....</b>	<b>411</b>
<i>Warren Armstrong, Alistair P. Rendell</i>	
<b>A Comparison of Search Heuristics for Empirical Code Optimization.....</b>	<b>421</b>
<i>Keith Seymour, Haihang You, Jack Dongarra</i>	
<b>An Optimized Dynamic Load Balancing Method for Parallel 3-D Mesh Refinement for Finite Element Electromagnetics With Tetrahedra .....</b>	<b>430</b>
<i>Da Qi Ren, Dennis D. Giannacopoulos, Reiji Suda</i>	
<b>Design and Implementation of an Effective Hypertransport Core in FPGA .....</b>	<b>437</b>
<i>Fei Chen, Hailiang Cheng, Xiaojun Yang, Rui Liu</i>	
<b>New Techniques for Simulating High Performance MPI Applications on Large Storage Networks.....</b>	<b>444</b>
<i>Alberto Nunez, Javier Fernandez, Jose Daniel Garcia, Jesus Carretero</i>	
<b>A Scalable, High Performance Infiniband-Attached SAN Volume Controller .....</b>	<b>453</b>
<i>D. Scott Guthridge</i>	

# Table of Contents

<b>Jetter: A Multi-Pattern Parallel I/O Benchmark</b> .....	459
<i>Liqiang Cao, Hongbing Luo, Baoyin Zhang</i>	
<b>Enhancing Write Performance of a Shared-Disk Cluster Filesystem Through a Fine-Grained Locking Strategy</b> .....	464
<i>Paulo A. Lopes, Pedro D. Medeiros</i>	
<b>Active Storage Using Object-Based Devices</b> .....	472
<i>Tina Miriam John, Anuradharthi Thiruvenkata Ramani, John A. Chandy</i>	
<b>Exploiting Data Compression in Collective I/O Techniques</b> .....	479
<i>Rosa Filgueira, David E. Singh, Juan C. Pichel, Jesus Carretero</i>	