

2006 IEEE International Conference on Cluster Computing

**Barcelona, Spain
25-28 September 2006**

Volume 1 of 2



**IEEE Catalog Number:
ISBN:**

**06TH8880
1-4244-0327-8**

2006 IEEE International Conference on Cluster Computing

Copyright and Reprint Permission: 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 that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. Other copying, reprint, or reproduction requests should be addressed to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331. All rights reserved. Copyright © 2006 by the Institute of Electrical and Electronics Engineers, Inc.

© 2006 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

IEEE Catalog Number:	06TH8880
ISBN:	1-4244-0327-8
ISSN:	1552-5244

Ordering information:

IEEE Operations Center
445 Hoes Lane
Piscataway, NJ 08854-4150
USA

+1 800 678 IEEE (+1 800 678 4333)
+1 732 981 1393
+1 732 981 9667 (FAX)
email: customer-service@ieee.org

Table of Contents

The petacomp machine: A MIMD cluster for parallel pattern-mining	1
<i>O.Birkeland, O. Snøve Jr., A. Halaas, M. Nedland, P. Sætrom</i>	
Empirical Study on Reducing Energy of Parallel Programs using Slack Reclamation by DVFS in a Power-scalable High Performance Cluster.....	11
<i>Hideaki Kimura, Mitsuhsa Sato, Yoshihiko Hotta, Taisuke Boku, Daisuke Takahashi</i>	
Cluster-based IP Router: Implementation and Evaluation	21
<i>Qinghua Ye, Mike H. MacGregor</i>	
Protocol Offload Evaluation Using SIMICS	31
<i>Andres Ortiz, Julio Ortega, Antonio F. Diaz, Alberto Prieto</i>	
Autonomic Management of Clustered Applications.....	40
<i>Sara Bouchenak, Noel De Palma, Daniel Hagimont, Christophe Taton</i>	
XCPU: a new, 9p-based, process management system for clusters and grids.....	51
<i>Ronald Mimich, Andrey Mirtchovski</i>	
Autonomic Resource Management for a Cluster that Executes Batch Jobs	61
<i>L.G. Alex Sung, Johnny W. Wong</i>	
Rapid Node Reallocation between Virtual Clusters for Data Intensive Utility Computing.....	71
<i>Alistair Coles, Aled Edwards</i>	
The Influence of Operating Systems on the Performance of Collective Operations at Extreme Scale	81
<i>Pete Beckman, Kamil Iskra, Kazutomo Yoshii, Susan Coghlan</i>	
Kernel-Level Measurement for Integrated Parallel Performance Views: the KTAU Project.....	93
<i>Aroon Nataraj, Allen Malony, Sameer Shende, Alan Morris</i>	
Efficient MPI Collective Operations for Clusters in Long-and-Fast Networks	105
<i>Motohiko Matsuda, Yutaka Ishikawa, Tomohiro Kudoh, Yuetsu Kodama, Ryousei Takano</i>	
Fine-Grained Message Pipelining for Improved MPI Performance	114
<i>Arun Rodrigues, Kyle Wheeler, Peter Kogge, Keith Underwood</i>	
Designing High Performance and Scalable MPI Intra-node Communication Support for Clusters.....	123
<i>Lei Chai, Albert Hartono, Dhabaleswar K. Panda</i>	
FAIL-MPI: How fault-tolerant is fault-tolerant MPI?.....	133
<i>William Hoarau, Pierre Lemarinier, Thomas Herault, Eric Rodriguez, Sebastien Tixeuil, Franck Cappello</i>	
A Hybrid MPI Simulator	143
<i>Rolf Riesen</i>	
Lightweight I/O for Scientific Applications.....	152
<i>Ron Oldfield, Arthur B. Maccabe, Lee Ward, Todd Kordenbrock, Rolf Riesen, Patrick Widener</i>	
I/O Scheduling Service for Multi-Application Clusters.....	163
<i>Adrien Lebre, Guillaume Huard, Yves Denneulin, Przemyslaw Sowa</i>	
JOSHUA: Symmetric Active/Active Replication for Highly Available HPC Job and Resource Management.....	173
<i>Kai Uhlemann, Christian Engelmann, Stephen L. Scott</i>	
NVisionPA: A Process Accounting Analysis Tool with a Security Focus on Masquerade Detection in HPC Clusters.....	183
<i>Charis Ermopoulos, William Yurcik</i>	
MSSG: A Framework for Massive-Scale Semantic Graphs	193
<i>Timothy D. R. Hartley, Umit Catalyurek, Füsün Özgüner, Andy Yoo, Scott Kohn, Keith Henderson</i>	
Z-align: An Exact and Parallel Strategy for Local Biological Sequence Alignment in User-Restricted Memory Space.....	203
<i>Rodolfo Bezerra Batista, Alba Cristina Magalhaes Alves de Melo</i>	

Table of Contents

A Simple Synchronous Distributed-Memory Algorithm for the HPCC RandomAccess Benchmark	213
<i>Steven J. Plimpton, Ron Brightwell, Courtenay Vaughan, Keith Underwood, Mike Davis</i>	
Energy-Aware Duplication Strategies for Scheduling Precedence Constrained Parallel Tasks on Clusters	220
<i>Ziliang Zong, Adam Manzanares, Brian Stinar, Xiao Qin</i>	
Locality Conscious Processor Allocation and Scheduling for Mixed Parallel Applications.....	228
<i>Nagavijayalakshmi Vydyanathan, S. Krishnamoorthy, G. Sabin, U. Catalyurek, T. Kurc, P., Sadayappan, J. Saltz</i>	
Using Simulation, Historical and Hybrid Estimation Systems for Enhancing Job Scheduling on NOWs	238
<i>M. Hanzich, P. Hernández, E. Luque, Francesc Giné, F. Solsona, J.L. Lérida</i>	
An Iteration Aware Multidimensional Data Distribution Prototype for Computing Clusters.....	250
<i>Baoqiang Yan, and Philip J. Rhodes</i>	
Improving Communication Performance on InfiniBand by Using Efficient Data Placement Strategies	260
<i>Robert Rex, Frank Mietke, Wolfgang Rehm, Christoph Raisch, and Hoang-Nam Nguyen</i>	
A New Flexible MPI Collective I/O Implementation	267
<i>Kenin Coloma, Avery Ching, Alok Choudhary, Wei-keng Liao, Rob Ross, Rajeev Thakur, Lee Ward</i>	
Stochastic Scheduling with Availability Constraints in Heterogeneous Clusters.....	276
<i>Tao Xie, Xiao Qin</i>	
Parallel Morphological/Neural Classification of Remote Sensing Images Using Fully Heterogeneous and Homogeneous Commodity Clusters	287
<i>Javier Plaza, Rosa Pérez, Antonio Plaza, Pablo Martínez, David Valencia</i>	
Robust task scheduling in non-deterministic heterogeneous computing systems	297
<i>Zhiao Shi, Emmanuel Jeannot, Jack Dongarra</i>	
Parallelizing Lattice Gauge Theory Models on Commodity Clusters.....	307
<i>Cyril Banino-Rokkones, Jørn Amundsen, Eivind Smørgrav</i>	
HPC Cluster Readiness of Xen and User Mode Linux.....	315
<i>Wesley Emenecker, Dan Stanzione</i>	
Scheduling Workflow-based Parameter-Sweep Applications with Best- Intermediate-Result-First Heuristic	323
<i>Kunaporn Srimanatham, Veera Muangsin</i>	
Increasing the cluster availability using RADIC	329
<i>Angelo Duarte, Dolores Rexachs, Emilio Luque</i>	
Multi-Objective Models for Scheduling Jobs on Parallel Computer Systems.....	337
<i>Sangsuree Vasupongayya, Su-Hui Chiang</i>	
Can a Helmholtz solver run on a cluster?.....	343
<i>Ralf Gruber, Vincent Keller, Emmanuel Leriche, Marc-Antoine Habisreutinger</i>	
SLA-Based Coordinated Superscheduling Scheme for Computational Grids.....	351
<i>Rajiv Ranjan, Aaron Harwood, Rajkumar Buyya</i>	
Performance Analysis, Modeling and Prediction of a Parallel Multiblock Lattice Boltzmann Application Using Prophecy System	359
<i>Xingfu Wu, Valerie Taylor, Shane Garrick, Dazhi Yu, Jacques Richard</i>	
Implementation Tradeoffs of the Array Files Library for Out-of-Core Computations	367
<i>Yueyue Zhang, Amy Apon</i>	
Implementation and Evaluation of Parallel Sparse Matrix-Vector Products on Distributed Memory Parallel Computers.....	375
<i>Rukhsana Shahnaz, Anila Usman, Imran R. Chughtai</i>	
A Performance Prediction Methodology for Data-dependent Parallel Applications.....	381
<i>Paula Cecilia Fritzsche Frugoni, Concepció Roig, Ana Ripoll, Aura Hernández, Emilio Luque</i>	

Table of Contents

Combining Clusters for Business Continuity	389
<i>Steve McKinty</i>	
Replay: A Model-Based Service for Supporting Transparent Cluster Analysis Tools.....	395
<i>Diwakar Krishnamurthy, Cameron Kiddle, Jerry Rolia, Rob Simmonds</i>	
STAS: A Scalability Testing and Analysis System.....	406
<i>Yong Chen, Xian-He Sun</i>	
A Model-Based Framework for the Integration of Parallel Tools.....	416
<i>Greg Watson, Nathan DeBardeleben</i>	
An Integrated Adaptive Management System for Cluster-based Web Services	427
<i>Ying Jiang, Dan Meng, Chao Ren, Jianfeng Zhan</i>	
On the Effectiveness of Content-aware Load Distribution for Web Clusters	437
<i>Mon-Yen Luo</i>	
Application-aware interface for SOAP communication in Web Services.....	447
<i>Hao Wang, Yizhu Tong, Hong Liu, Taoying Liu</i>	
A Performance Instrumentation Framework to Characterize Computation-Communication Overlap in Message-Passing Systems	455
<i>Aniruddha G. Shet, P. Sadayappan, David E. Bernholdt, Jarek Nieplocha, Vinod Tipparaju</i>	
Predictive Performance Analysis of a Parallel Pipelined Synchronous Wavefront Application for Commodity Processor Cluster Systems	467
<i>Gihan R. Mudalige, Stephen A. Jarvis, Daniel P. Spooner, Graham R. Nudd</i>	
Modeling Network Contention Effects on All-to-All Operations	479
<i>Luiz Angelo Steffene</i>	
MPJ Express: Towards Thread Safe Java HPC	489
<i>Mark Baker, Bryan Carpenter, Aamir Shafi</i>	
Automatic Clustering for Self-Organizing Grids.....	499
<i>Weishuai Yang, Nael Abu-Ghazaleh, Michael Lewis</i>	
Resource Management for Interactive Jobs in a Grid Environment	508
<i>Enol Fernández, Elisa Heymann, Miquel A. Senar</i>	
A Virtual Registry For Wide-Area Messaging.....	518
<i>Mark Baker, Matthew Grove</i>	
Parallel Basic Matrix Algebra on the Grid'5000 Large Scale Distributed Platform	528
<i>Lamine Aouad, Serge Petiton</i>	
Marching Towards Nirvana: Configurations for Very High Performance Parallel File Systems	536
<i>Phil Andrews, Chris Jordan, Wayne Pfeiffer</i>	
Using Lightweight Transactions and Snapshots for Fault-Tolerant Services Based on Shared Storage Bricks	548
<i>Michail D. Flouris, Renaud Lachaize, and Angelos Bilas</i>	
Efficient Data-Movement for Lightweight I/O.....	558
<i>Ron A. Oldfield, Patrick Widener, Arthur B. Maccabe, Lee Ward and Todd Kordenbrock</i>	
Improving the Performance of Cluster Applications through I/O Proxy Architecture	567
<i>L. M. Sanchez, F. Isaila, A. Calderón, D. E. Singh and J. D. García</i>	
Fairness and Performance Isolation: an Analysis of Disk Scheduling Algorithms.....	574
<i>Seetharami R. Seelam and Patricia J. Teller</i>	
Positioning Dynamic Storage Caches for Transient Data	584
<i>Sudharshan S. Vazhkudai, Douglas Thain, Xiaosong Ma, Vincent W. Freeh</i>	

Table of Contents

Matrix Multiplication on Two Interconnected Processors	593
<i>Brett A. Becker and Alexey Lastovetsky</i>	
Analytical Network Modeling of Heterogeneous Large-Scale Cluster Systems	602
<i>Bahman Javadi, Jemal H. Abawajy, Mohammad K. Akbari, and Saeid Nahavandi</i>	
Heterogeneous Parallel Computing in Remote Sensing Applications: Current Trends and Future Perspectives	611
<i>Antonio J. Plaza</i>	
Open MPI: A High-Performance, Heterogeneous MPI	621
<i>Richard L. Graham, Galen M. Shipman, Brian W. Barrett, Ralph H. Castain, George Bosilca, and Andrew Lumsdaine</i>	
A 2-Approximation Algorithm for Scheduling Independent Tasks onto a Uniform Parallel Machine and its Extension to a Computational Grid	630
<i>Noriyuki Fujimoto and Kenichi Hagihara</i>	
JaceP2P: an Environment for Asynchronous Computations on Peer-to-Peer Networks	637
<i>Jacques M. Bahi, Raphaël Couturier and Philippe Vuillemin</i>	
Virtual Structured P2P Network Topology for Distributed Computing	647
<i>Euloge Edi, Tahar Kechadi and Ronan McNulty</i>	
A Parallel Algorithm for the Solution of the Deconvolution Problem on Heterogeneous Networks	656
<i>Pedro Alonso, Alexey Lastovetsky and Antonio M. Vidal</i>	
Well Balanced Sparse Matrix-vector Multiplication on a Parallel Heterogeneous System	665
<i>Clovis Dongmo Jiogo, Pierre Kuonen and Pierre Manneback</i>	
TGrid - Grid Runtime Support for Hierarchically Structured Task-parallel Programs	673
<i>Sascha Hunold, Thomas Rauber and Gudula Rünger</i>	
A Quadratic Self-Scheduling Algorithm for Heterogeneous Distributed Computing Systems	683
<i>J. Díaz, S. Reyes, A. Niño and C. Muñoz-Caro</i>	
Self-Adapting Scheduling for Tasks with Dependencies in Stochastic Environments	691
<i>I. Riakotakis, F. M. Ciorba, T. Andronikos and G. Papakonstantinou</i>	
A Framework for Adaptive Communication Modeling on Heterogeneous Hierarchical Clusters	699
<i>Wahid Nasri, Hajer Hamad and Hadhemi Fejjari</i>	
Out of User Space Storage and RDMA	709
<i>Michael Ko, Renato Recio, Claudia Salzberg</i>	
Initial Performance Evaluation of the NetEffect 10 Gigabit iWARP Adapter	719
<i>Dennis Dalessandro, Pete Wyckoff, Gary Montry</i>	
Exploiting RDMA operations for Providing Efficient Fine-Grained Resource Monitoring in Cluster-based Servers	726
<i>K. Vaidyanathan, H.-W. Jin, D. K. Panda</i>	
Software iSCSI over iWARP on a General Purpose CPU	736
<i>Abhijeet A. Joglekar, Steve King</i>	
Architecture and Implementation of Sockets Direct Protocol in Windows	745
<i>Dror Goldenberg, Tzachi Dar, Gilad Shainer</i>	
Experiences from Debugging a PCI-X-based RDMA-capable NIC	754
<i>Manolis Marazakis, Vassilis Papafstathiou, Giorgos Kalokairinos, Angelos Bilas</i>	
Performance Study of Winsock Direct with 10 Gb Ethernet RDMA-Enabled NIC	764
<i>Brian Hausauer</i>	