

# **2010 IEEE International Conference on Cluster Computing (CLUSTER 2010)**

**Heraklion, Greece  
20-24 September 2010**



**IEEE Catalog Number: CFP10235-PRT  
ISBN: 978-1-4244-8373-0**

# 2010 IEEE International Conference on Cluster Computing

## CLUSTER 2010

### Table of Contents

Foreword.....	ix
Conference Organization.....	x
Program Committee.....	xii
External Reviewers.....	xiv
Sponsors and Supporters.....	xvi
Keynotes.....	xvii

---

#### Session 1

Minimizing MPI Resource Contention in Multithreaded Multicore Environments .....	1
<i>David Goodell, Pavan Balaji, Darius Buntinas, Gábor Dózsa, William Gropp, Sameer Kumar, Bronis R. de Supinski, and Rajeev Thakur</i>	
TCCluster: A Cluster Architecture Utilizing the Processor Host Interface as a Network Interconnect .....	9
<i>Heiner Litz, Maximilian Thuermer, and Ulrich Bruening</i>	
Adaptive Optimization for Petascale Heterogeneous CPU/GPU Computing .....	19
<i>Canqun Yang, Feng Wang, Yunfei Du, Juan Chen, Jie Liu, Huizhan Yi, and Kai Lu</i>	

#### Session 2

How to scale Nested OpenMP Applications on the ScaleMP vSMP Architecture .....	29
<i>Dirk Schmidl, Christian Terboven, Andreas Wolf, Dieter an Mey, and Christian Bischof</i>	
Synchronizing the Timestamps of Concurrent Events in Traces of Hybrid MPI/OpenMP Applications .....	38
<i>Daniel Becker, Markus Geimer, Rolf Rabenseifner, and Felix Wolf</i>	
Getting Rid of Coherency Overhead for Memory-Hungry Applications .....	48
<i>Héctor Montaner, Federico Silla, Holger Fröning, and Jose Duato</i>	

### Session 3

Energy-Aware Scheduling in Virtualized Datacenters .....	58
<i>Íñigo Goiri, Ferran Julià, Ramón Nou, Josep Ll. Berral, Jordi Guitart, and Jordi Torres</i>	
TRACER: A Trace Replay Tool to Evaluate Energy-Efficiency of Mass Storage Systems .....	68
<i>Zhuo Liu, Fei Wu, Xiao Qin, Changsheng Xie, Jian Zhou, and Jianzong Wang</i>	
Designing OS for HPC Applications: Scheduling .....	78
<i>Roberto Gioiosa, Sally A. McKee, and Mateo Valero</i>	

### Session 4

Exploiting Data Deduplication to Accelerate Live Virtual Machine Migration .....	88
<i>Xiang Zhang, Zhigang Huo, Jie Ma, and Dan Meng</i>	
SHelp: Automatic Self-Healing for Multiple Application Instances in a Virtual Machine Environment .....	97
<i>Gang Chen, Hai Jin, Deqing Zou, Bing Bing Zhou, Weizhong Qiang, and Gang Hu</i>	
Virtualizing Modern High-Speed Interconnection Networks with Performance and Scalability .....	107
<i>Bo Li, Zhigang Huo, Panyong Zhang, and Dan Meng</i>	

### Session 5

RDMA-Based Job Migration Framework for MPI over InfiniBand .....	116
<i>Xiangyong Ouyang, Sonya Marcarelli, Raghunath Rajachandrasekar, and Dhabaleswar K. Panda</i>	
Host Side Dynamic Reconfiguration with InfiniBand .....	126
<i>Wei Lin Guay, Sven-Arne Reinemo, Olav Lysne, Tor Skeie, Bjørn Dag Johnsen, and Line Holen</i>	
Multiplexing Endpoints of HCA for Scaling MPI Applications: Design and Performance Evaluation with uDAPL .....	136
<i>Jasjit Singh and Yogeshwar Sonawane</i>	

### Session 6

The Impact of System Design Parameters on Application Noise Sensitivity .....	146
<i>Kurt B. Ferreira, Patrick G. Bridges, Ron Brightwell, and Kevin T. Pedretti</i>	
Computing Contingency Statistics in Parallel: Design Trade-Offs and Limiting Cases .....	156
<i>Philippe Pébay, David Thompson, and Janine Bennett</i>	
Integration Experiences and Performance Studies of A COTS Parallel Archive System .....	166
<i>Hsing-bung Chen, Gary Grider, Cody Scott, Milton Turley, Aaron Torres, Kathy Sanchez, and John Bremer</i>	

## Session 7

Enforcing SLAs in Scientific Clouds .....	178
<i>Oliver Niehörster, André Brinkmann, Gregor Fels, Jens Krüger, and Jens Simon</i>	
CDRM: A Cost-Effective Dynamic Replication Management Scheme for Cloud Storage Cluster .....	188
<i>Qingsong Wei, Bharadwaj Veeravalli, Bozhao Gong, Lingfang Zeng, and Dan Feng</i>	
An Efficient Process Live Migration Mechanism for Load Balanced Distributed Virtual Environments .....	197
<i>Balazs Gerofi, Hajime Fujita, and Yutaka Ishikawa</i>	

## Session 8

Acceleration of Streamed Tensor Contraction Expressions on GPGPU-Based Clusters .....	207
<i>Wenjing Ma, Sriram Krishnamoorthy, Oreste Villay, and Karol Kowalski</i>	
Efficient Parallel Subgraph Counting Using G-Tries .....	217
<i>Pedro Ribeiro, Fernando Silva, and Luís Lopes</i>	
Cluster versus GPU Implementation of an Orthogonal Target Detection Algorithm for Remotely Sensed Hyperspectral Images .....	227
<i>Abel Paz and Antonio Plaza</i>	

## Session 9

Breaking the MapReduce Stage Barrier .....	235
<i>Abhishek Verma, Nicolas Zea, Brian Cho, Indranil Gupta, and Roy H. Campbell</i>	
Asynchronous Algorithms in MapReduce .....	245
<i>Karthik Kambatla, Naresh Rapolu, Suresh Jagannathan, and Ananth Grama</i>	
Reducing Communication Overhead in Large Eddy Simulation of Jet Engine Noise .....	255
<i>Yingchong Situ, Lixia Liu, Chandra S. Martha, Matthew E. Louis, Zhiyuan Li, Ahmed H. Sameh, Gregory A. Blaisdell, and Anastasios S. Lyrintzis</i>	

## Session 10

Performance Analysis of Multi-level Time Sharing Task Assignment Policies on Cluster-Based Systems .....	265
<i>Malith Jayasinghe, Zahir Tari, and Panlop Zeephongsekul</i>	
A Simulation Framework to Automatically Analyze the Communication-Computation Overlap in Scientific Applications .....	275
<i>Vladimir Subotic, Jose Carlos Sancho, Jesus Labarta, and Mateo Valero</i>	
Analysis of Tasks Reallocation in a Dedicated Grid Environment .....	284
<i>Yves Caniou, Ghislain Charrier, and Frédéric Desprez</i>	

## Session 11

Replication-Based Highly Available Metadata Management for Cluster File Systems .....	292
<i>Zhuan Chen, Jin Xiong, and Dan Meng</i>	
Improving Parallel I/O Performance with Data Layout Awareness .....	302
<i>Yong Chen, Xian-He Sun, Rajeev Thakur, Huaiming Song, and Hui Jin</i>	
Optimization Techniques at the I/O Forwarding Layer .....	312
<i>Kazuki Ohta, Dries Kimpe, Jason Cope, Kamil Iskra, Robert Ross, and Yutaka Ishikawa</i>	

## Author Index