

2009 International Conference on High Performance Computing

(HiPC 2009)

**Kochi, India
16-19 December 2009**



**IEEE Catalog Number: CFP09176-PRT
ISBN: 978-1-4244-4922-4**

Table of Contents

Session I - Scheduling and Resource Management

ERfair Scheduler with Processor Shutdown	4
<i>Arnab Sarkar, Sarthak Swaroop, Sujoy Ghose, and Partha Chakrabarti</i>	
Service Oriented Architecture for Job Submission and Management on Grid Computing Resources	13
<i>Archit Kulshrestha and Gabrielle Allen</i>	
Improved Opportunistic Scheduling Algorithms for WiMAX Mobile Multihop Relay Networks	20
<i>Srinath Narasimha and Krishna Sivalingam</i>	
Automatic Data Placement and Replication in Grids	30
<i>Ying Ding and Ying Lu</i>	
An Effective Scheduling Algorithm for Linear Makespan Minimization on Unrelated Parallel Machines	40
<i>Liya Fan, Fa Zhang, Gongming Wang, and Zhiyong Liu</i>	
Spanning Tree Routing Strategies for Divisible Load Scheduling on Arbitrary Graphs - a Comparative Performance Analysis	50
<i>Sivakumar Viswanathan, Bharadwaj Veeravalli, and Jia Jingxi</i>	

Session II - Best Paper Session

Non Uniform Power Access in Large Caches with Low-Swing Wires	59
<i>Aniruddha Udipi, Naveen Muralimanohar, and Rajeev Balasubramonian</i>	
Group File Operations for Scalable Tools and Middleware	69
<i>Michael Brim and Barton P. Miller</i>	

Session III - Architecture

Distance-Aware Round-Robin Mapping for Large NUCA Caches	79
<i>Alberto Ros, Manuel Acacio, Marcelo Cintra, and José M. García</i>	
Three Scalable Approaches to Improving Many-core Throughput for a Given Peak Power Budget	89
<i>John Sartori and Rakesh Kumar</i>	
Fast Checkpointing by Write Aggregation with Dynamic Buffer and Interleaving on Multicore Architecture	99
<i>Xiangyong Ouyang, Karthik Gopalakrishnan, Tejus Gangadharappa, and Dhabaleswar Panda</i>	
HiPPAI: High Performance Portable Accelerator Interface for SoCs	109
<i>Paul Stillwell, Vineet Chadha, Omesh Tickoo, Steven Zhang, Ramesh Illikkal, Ravishankar Iyer, and Don Newell</i>	

P-Slice Based Efficient Speculative Multithreading	119
<i>Rakesh Ranjan, Pedro Marcuello, Fernando Latorre, and Antonio Gonzalez</i>	
A Framework for Routing and Resource Allocation in Network Virtualization	129
<i>Åshild Solheim, Olav Lysne, Tor Skeie, Thomas Sødning, and Sven-Arne Reinemo</i>	
Cache Streamization for High Performance Stream Processor	140
<i>Nan Wu, Mei Wen, Ju Ren, Yi He, ChangQuing Xun, Wei Wu, and Chunyuan Zhang</i>	
Terascale Chip Multiprocessor Memory Hierarchy and Programming Model	150
<i>Shoumeng Yan, Xiaocheng Zhou, Ying Gao, Hu Chen, Sai Luo, Peinan Zhang, Naveen Cherukuri, Ronny Ronen, and Bratin Saha</i>	

Session IV - Mobile Computing and Sensor Networks

On Providing Event Reliability and Maximizing Network Lifetime Using Mobile Data- Collectors in Wireless Sensor Networks	160
<i>Saamaja Vupputuri, Kiran Rachuri, and C. Siva Ram Murthy</i>	
Taming the Exponential State Space of the Maximum Lifetime Sensor Cover problem	170
<i>Akshaye Dhawan and Sushil K. Prasad</i>	
A Comparison of Secure Data Aggregation Schemes for Wireless Sensor Networks	179
<i>Rajendra Boppana and Pengjun Pan</i>	
An Energy Efficient Deterministic Key Establishment Scheme for Clustered Wireless Sensor Networks	189
<i>A. S. Poornima, B. B. Amberker, and Harihar Baburao Jadhav</i>	
Enhancing Bandwidth Reservation Guarantee for QoS Routing Protocol in Mobile Ad Hoc Networks	195
<i>Khaled Al Mohd. Soufy and Ash Mohammed Abbas</i>	
Integrating Traffic Estimation and Dynamic Channel Reconfiguration in Wireless Mesh Networks	205
<i>Athula Balachandran, A. Antony Franklin, and C. Siva Ram Murthy</i>	

Session V - Software Systems

Integrating and Optimizing Transactional Memory In a Data Mining Middleware	215
<i>Vignesh Ravi and Gagan Agrawal</i>	
Impact of Early Abort Mechanisms on Lock-Based Software Transactional Memory	225
<i>Zhengyu He and Bo Hong</i>	
Supporting Load Balancing For Distributed Data-Intensive Applications	235
<i>Leonid Glimcher, Vignesh Ravi, and Gagan Agrawal</i>	
CellMT: A Cooperative Multithreading Library for the Cell/B.E.	245
<i>Vicenç Beltran, David Carrera, Jordi Torres, and Eduard Ayguadé</i>	
Extracting the Textual and Temporal Structure of Supercomputing Logs	254
<i>Sourabh Jain, Inderpreet Singh, Abhishek Chandra, Zhi-Li Zhang, and Greg Bronevetsky</i>	
Compile-time Disambiguation of MATLAB Types through Concrete Interpretation with Automatic Run-time Fallback	264
<i>Chun-Yu Shei, Arun Chauhan, and Sidney Shaw</i>	
Statistical Workload Shaping for Storage Systems	274
<i>Hui Wang and Peter Varman</i>	
Demand-driven Execution of Static Directed Acyclic Graphs	284
<i>Prabhanjan Kambadur, Anshul Gupta, Torsten Hoefler, and Andrew Lumsdaine</i>	

Session VI - Communication networks

A Throughput Fairness Injection Protocol for Mesh and Torus networks	294
<i>Cruz Izu</i>	
Multiple Virtual Lanes-aware MPI Collective Communication in Multi-core Clusters	304
<i>Bo Li, Zhigang Huo, PanYong Zhang, and Dan Meng</i>	
High Search Performance, Small Document Index: P2P Search Can Have Both	312
<i>Yingwu Zhu and Haiying Shen</i>	
CORP: A Cooperative File Replication Protocol for Structured P2P Networks	322
<i>Haiying Shen</i>	

Session VII - Algorithms

Highly Scalable Algorithm For Distributed Real-Time Text Indexing	332
<i>Ankur Narang, Vikas Agarwal, Monu Kedia, and Vijay Garg</i>	
A Parallel Algorithm for Exact Bayesian Network Inference	342
<i>Olga Nikolova, Jaroslaw Zola, and Srinivas Aluru</i>	
An Efficient Parallel Algorithm for Evaluating Join Queries on Heterogeneous Distributed Systems	350
<i>M. Al Hajj Hassan and M. Bamha</i>	
A Fast Algorithm for Energy-Aware Mapping of Cores Onto WK-Recursive NoC under Performance Constraints	359
<i>Wei Hu, Chen Du, Like Yan, and Chen Tianzhou</i>	

Session VIII - Applications

Towards a Robust, Real-time Face Processing System using CUDA-enabled GPUs	368
<i>Bharatkumar Sharma, Rahul Thota, Naga Vydyanathan, and Amit Kale</i>	
Comparing the performance of Clusters, Hadoop, and Active Disks on Microarray Correlation Computations	378
<i>Jeffrey Delmerico, Nathaniel Byrnes, Andrew E. Bruno, Matthew D. Jones, Steven M. Gallo, and Vipin Chaudhary</i>	
Speculative p-DFAs for Parallel XML Parsing	388
<i>Ying Zhang, Yinfei Pan, and Kenneth Chiu</i>	
Performance Optimizations for Distributed Real-time Text Indexing	398
<i>Ankur Narang, Karthik Swaminathan, and Prashant Agrawal</i>	
Optimizing the use of GPU Memory in Applications with Large data sets	408
<i>Nadathur Satish, Narayanan Sundaram, and Kurt Keutzer</i>	
Detailed analysis of I/O traces of large scale applications	419
<i>Nithin Nakka, Alok Choudhary, Wei-Keng Liao, Lee Ward, Ruth Klundt, and Marlow I. Weston</i>	
Designing Systems for Large-Scale, Discrete-Event Simulations: Experiences with the FastTrans Parallel Microsimulator	428
<i>Sunil Thulasidasan, Shiva Kasiviswanathan, Stephan Eidenbenz, Emanuele Galli, Susan Mniszewski, and Philip Romero</i>	
Acceleration of Conjugate Gradient Method for Circuit Simulation Using CUDA	438
<i>Anirudh Maringanti, Viraj Athavale, and Sachin B. Patkar</i>	

Session IX - Performance Evaluations

Continuous Performance Monitoring for Large-Scale Parallel Applications	445
<i>Isaac Dooley, Chee Wai Lee, and Laxmikant Kale</i>	

Evaluating Implications of Virtual Worlds on Server Architecture using Second Life	453
<i>Srihari Makineni, Omesh Tickoo, Aaron Terrell, Jessica Young, and Don Newell</i>	
A Performance Prediction Model for the CUDA GPGPU Platform	463
<i>Kishore Kothapalli, Rishabh Mukherjee, M. Suhail Rehman, Suryakant Patidar, P. J. Narayanan, and Kannan Srinathan</i>	
Author Index	473