

2009 38th International Conference on Parallel Processing (ICPP 2009)

**Vienna, Austria
22 – 25 September 2009**



**IEEE Catalog Number: CFP09127-PRT
ISBN: 978-1-4244-4961-3**

2009 International Conference on Parallel Processing

ICPP 2009

Table of Contents

Message from the General Chairs
Message from the Program Co-Chairs
Organizing and Program Committee
Reviewers

Session 1A: Architecture I

Using Coherence Information and Decay Techniques to Optimize L2 Cache Leakage in CMPs	1
<i>Matteo Monchiero, Ramon Canal, and Antonio González</i>	
Fine-grain Parallelism Using Multi-core, Cell/BE, and GPU Systems: Accelerating the Phylogenetic Likelihood Function	9
<i>Frederico Pratas, Pedro Trancoso, Alexandros Stamatakis, and Leonel Sousa</i>	
Bank-aware Dynamic Cache Partitioning for Multicore Architectures	18
<i>Dimitris Kaseridis, Jeffrey Stuecheli, and Lizy K. John</i>	

Session 1B: Grid / Cloud / Cluster I

Hardware Implementation Study of the Deficit Table Egress Link Scheduling Algorithm	26
<i>Raúl Martínez, José M. Claver, Francisco J. Alfaro, and José L. Sánchez</i>	
Accelerating Checkpoint Operation by Node-Level Write Aggregation on Multicore Systems	34
<i>Xiangyong Ouyang, Karthik Gopalakrishnan, and Dhabaleswar K. Panda</i>	
Speeding Up Distributed MapReduce Applications Using Hardware Accelerators	42
<i>Yolanda Becerra, Vicenç Beltran, David Carrera, Marc González, Jordi Torres, and Eduard Ayguadé</i>	

Session 1C: Wireless / Sensor I

Run to Potential: Sweep Coverage in Wireless Sensor Networks	50
<i>Min Xi, Kui Wu, Yong Qi, Jizhong Zhao, Yunhao Liu, and Mo Li</i>	
CED: A Community-Based Event Delivery Protocol in Publish/Subscribe Systems for Delay Tolerant Sensor Network (DTSN)	58
<i>Jinqi Zhu, Ming Liu, Jiannong Cao, Guihai Chen, Haigang Gong, and Fulong Xu</i>	
End-User Diagnosis of Communication Paths in Sensor Network Systems	66
<i>Qing Cao, Dong Wang, and Tarek Abdelzaher</i>	

Session 2A: Programming Model I

Efficient Scheduling of Nested Parallel Loops on Multi-Core Systems	74
<i>Arun Kejariwal, Alexandru Nicolau, Alexander V. Veidenbaum, Utpal Banerjee, and Constantine D. Polychronopoulos</i>	
A Parallel Skeleton Library for Multi-core Clusters	84
<i>Yuki Karasawa and Hideya Iwasaki</i>	
Parallel Phase Model: A Programming Model for High-end Parallel Machines with Manycores	92
<i>Ron Brightwell, Mike Heroux, Zhaofang Wen, and Junfeng Wu</i>	

Session 2B: Algorithms and Applications I

Fast Isosurface Extraction for Medical Volume Dataset on Cell BE	100
<i>Hai Jin, Bo Li, Ran Zheng, and Qin Zhang</i>	
Constructing Gene Regulatory Networks on Clusters of Cell Processors	108
<i>Jaroslav Zola, Abhinav Sarje, and Srinivas Aluru</i>	
Computing Equilibria in Bimatrix Games by Parallel Vertex Enumeration	116
<i>Jonathan Widger and Daniel Grosu</i>	

Session 2C: Performance Evaluation I

Barcelona OpenMP Tasks Suite: A Set of Benchmarks Targeting the Exploitation of Task Parallelism in OpenMP	124
<i>Alejandro Duran, Xavier Teruel, Roger Ferrer, Xavier Martorell, and Eduard Ayguadé</i>	
Performance Characterization of a Hierarchical MPI Implementation on Large-scale Distributed-memory Platforms	132
<i>Sadaf R. Alam, Richard Barrett, Jeffery Kuehn, and Steve Poole</i>	
Integrated Performance Views in Charm++: Projections Meets TAU	140
<i>Scott Biersdorff, Chee Wai Lee, Allen D. Malony, and Laxmikant V. Kale</i>	

Session 3A: Data Center I

DORA: A Dynamic File Assignment Strategy with Replication	148
<i>Jonathan Tjioe, Renata Widjaja, Abraham Lee, and Tao Xie</i>	
Designing Efficient FTP Mechanisms for High Performance Data-Transfer over InfiniBand	156
<i>Ping Lai, Hari Subramoni, Sundeep Narravula, Amit Mamidala, and Dhableswar K. Panda</i>	
Performance Evaluation of Energy-Efficient Parallel I/O Systems with Write Buffer Disks	164
<i>Xiaojun Ruan, Adam Manzanares, Shu Yin, Ziliang Zong, and Xiao Qin</i>	

Session 3B: Resource Management I

A Strategy-proof Pricing Scheme for Multiple Resource Type Allocations	172
<i>Yong Meng Teo and Marian Mihailescu</i>	
Broker Selection Strategies in Interoperable Grid Systems	180
<i>Ivan Rodero, Francesc Guim, Julita Corbalan, Liana Fong, and S. Masoud Sadjadi</i>	
Stochastic-Based Robust Dynamic Resource Allocation in a Heterogeneous Computing System	188
<i>Jay Smith, Edwin K.P. Chong, Anthony A. Maciejewski, and H.J. Siegel</i>	

Session 3C: Algorithms and Application II

Complexity Analysis and Performance Evaluation of Matrix Product on Multicore Architectures	196
<i>Mathias Jacquelin, Loris Marchal, and Yves Robert</i>	
Computing the Throughput of Replicated Workflows on Heterogeneous Platforms	204
<i>Anne Benoit, Matthieu Gallet, Bruno Gaujal, and Yves Robert</i>	
Analysis of Parallel Algorithms for Energy Conservation in Scalable Multicore Architectures	212
<i>Vijay Anand Korthikanti and Gul Agha</i>	

Session 4A: Architecture II

A Resource Optimized Remote-Memory-Access Architecture for Low-latency Communication	220
<i>Mondrian Nüssle, Martin Scherer, and Ulrich Brüning</i>	
DTM: Decoupled Hardware Transactional Memory to Support Unbounded Transaction and Operating System	228
<i>Shaogang Wang, Weixia Xu, Zhengbin Pang, Dan Wu, Qiang Dou, and Xiaodong Yang</i>	
CIFTS: A Coordinated Infrastructure for Fault-Tolerant Systems	237
<i>Rinku Gupta, Pete Beckman, Byung-Hoon Park, Ewing Lusk, Paul Hargrove, Al Geist, Dhableswar Panda, Andrew Lumsdaine, and Jack Dongarra</i>	

Session 4B: Grid / Cloud / Cluster II

Performance Analysis of DHT Algorithms for Range-Query and Multi-Attribute Resource Discovery in Grids	246
<i>Haiying Shen and Cheng-Zhong Xu</i>	
Using Standards-Based Interfaces to Share Data across Grid Infrastructures	254
<i>Karolina Sarnowska, Andrew Grimshaw, and Erwin Laure</i>	
GePSeA: A General-Purpose Software Acceleration Framework for Lightweight Task Offloading	261
<i>Ajeet Singh, Pavan Balaji, and Wu-chun Feng</i>	

Session 4C: Wireless / Sensor II

SkipStream: A Clustered Skip Graph Based On-demand Streaming Scheme over Ubiquitous Environments	269
<i>Qifeng Yu, Tianyin Xu, Baoliu Ye, Sanglu Lu, and Daoxu Chen</i>	
A Distributed Three-hop Routing Protocol to Increase the Capacity of Hybrid Networks	277
<i>Ze Li and Haiying Shen</i>	
Performance Limits of Fair-Access in Underwater Sensor Networks	285
<i>Yang Xiao, Miao Peng, John Gibson, Geoffrey G. Xie, and Ding-Zhu Du</i>	

Session 5A: Programming Model II

Investigating High Performance RMA Interfaces for the MPI-3 Standard	293
<i>Vinod Tipparaju, William Gropp, Hubert Ritzdorf, Rajeev Thakur, and Jesper L. Träff</i>	
Optimizing Communication Scheduling Using Dataflow Semantics	301
<i>Adrian Soviani and Jaswinder Pal Singh</i>	
Mapping the FDTD Application to Many-Core Chip Architectures	309
<i>Daniel A. Orozco and Guang R. Gao</i>	

Session 5B: Resource Management II

Slotted Wavelength Scheduling for Bulk Transfers in Research Networks	317
<i>Zhe Wang, Sanjay Ranka, and Ye Xia</i>	
Optimizing the Latency of Streaming Applications under Throughput and Reliability Constraints	325
<i>Anne Benoit, Mourad Hakem, and Yves Robert</i>	
Improving Resource Availability by Relaxing Network Allocation Constraints on Blue Gene/P	333
<i>Narayan Desai, Darius Buntinas, Daniel Buettner, Pavan Balaji, and Anthony Chan</i>	

Session 5C: Algorithms and Applications III

A Parallel Algorithm for Computing Betweenness Centrality	340
<i>Guangming Tan, Dengbiao Tu, and Ninghui Sun</i>	
Load Balance in the Phylogenetic Likelihood Kernel	348
<i>Alexandros Stamatakis and Michael Ott</i>	
Performance Models for Blocked Sparse Matrix-Vector Multiplication Kernels	356
<i>Vasileios Karakasis, Georgios Goumas, and Nectarios Koziris</i>	

Session 6A: Performance Evaluation II

On the Scalability of Parallel Verilog Simulation	365
<i>Sina Meraji, Wei Zhang, and Carl Tropper</i>	
Exploiting Simulation Slack to Improve Parallel Simulation Speed	371
<i>Jianwei Chen, Murali Annavaram, and Michel Dubois</i>	
Direct N-body Kernels for Multicore Platforms	379
<i>Nitin Arora, Aashay Shringarpure, and Richard W. Vuduc</i>	

Session 6B: Algorithms and Applications IV

Employing Transactional Memory and Helper Threads to Speedup Dijkstra's Algorithm	388
<i>Konstantinos Nikas, Nikos Anastopoulos, Georgios Goumas, and Nectarios Koziris</i>	
A Partition-Merge Based Cache-Conscious Parallel Sorting Algorithm for CMP with Shared Cache	396
<i>Song Hao, Zhihui Du, David Bader, and Yin Ye</i>	
Scalability of Time- and Space-Efficient Embedded Runge-Kutta Solvers for Distributed Address Space	404
<i>Matthias Korch and Thomas Rauber</i>	

Session 6C: P2P I

Heterogeneity-Aware Erasure Codes for Peer-to-Peer Storage Systems	412
<i>Lluís Pamies-Juarez, Pedro García-López, and Marc Sánchez-Artigas</i>	
SandStone: A DHT Based Carrier Grade Distributed Storage System	420
<i>Guangyu Shi, Jian Chen, Hao Gong, Lingyuan Fan, Haiqiang Xue, Qingming Lu, and Liang Liang</i>	
Exploring the Cost-Availability Tradeoff in P2P Storage Systems	429
<i>Zhi Yang, Yafei Dai, and Zhen Xiao</i>	

Session 7A: Architecture III

Code Semantic-Aware Runahead Threads	437
<i>Tanausú Ramírez, Alex Pajuelo, Oliverio J. Santana, and Mateo Valero</i>	
Thread Merging Schemes for Multithreaded Clustered VLIW Processors	445
<i>Manoj Gupta, Fermín Sánchez, and Josep Llosa</i>	

Register Versioning: A Low-Complexity Implementation of Register Renaming in Out-of-Order Microarchitectures	453
<i>Hui Zeng, Kanad Ghose, and Dmitry Ponomarev</i>	
Session 7B: Grid / Cloud / Cluster III	
Cache-Efficient, Intranode, Large-Message MPI Communication with MPICH2-Nemesis	462
<i>Darius Buntinas, Brice Goglin, David Goodell, Guillaume Mercier, and Stéphanie Moreaud</i>	
Using Subfiling to Improve Programming Flexibility and Performance of Parallel Shared-file I/O	470
<i>Kui Gao, Wei-keng Liao, Arifa Nisar, Alok Choudhary, Robert Ross, and Robert Latham</i>	
A Parallel Branch and Bound Algorithm for Workflow QoS Optimization	478
<i>Kevin Kofler, Irfan ul Haq, and Erich Schikuta</i>	
Session 7C: P2P II	
Mediacoop: Hierarchical Lookup for P2P-VoD Services	486
<i>Tieying Zhang, Jianming Lv, and Xueqi Cheng</i>	
SEIF: Search Enhanced by Intelligent Feedback in Unstructured P2P Networks	494
<i>Xiaoyu Yang and Yiming Hu</i>	
On Maximum Stability with Enhanced Scalability in High-Churn DHT Deployment	502
<i>Junfeng Xie, Zhenhua Li, Guihai Chen, and Jie Wu</i>	
Session 8A: Resource Management III	
Green Multicore-SoC Software-Execution Framework with Timely-Power-Gating Scheme	510
<i>Masafumi Onouchi, Keisuke Toyama, Toru Nojiri, Makoto Sato, Masayoshi Mase, Jun Shirako, Mikiko Sato, Masashi Takada, Masayuki Ito, Hiroyuki Mizuno, Mitaro Namiki, Keiji Kimura, and Hironori Kasahara</i>	
A Heuristic for Mapping Virtual Machines and Links in Emulation Testbeds	518
<i>Rodrigo N. Calheiros, Rajkumar Buyya, and César A.F. De Rose</i>	
LeWI: A Runtime Balancing Algorithm for Nested Parallelism	526
<i>Marta Garcia, Julita Corbalan, and Jesus Labarta</i>	
Session 8B: Algorithms and Applications V	
Scalable Parallel Execution of an Event-Based Radio Signal Propagation Model for Cluttered 3D Terrains	534
<i>Sudip K. Seal and Kalyan S. Perumalla</i>	

Generalizing k-Betweenness Centrality Using Short Paths and a Parallel Multithreaded Implementation	542
<i>Karl Jiang, David Ediger, and David A. Bader</i>	
Accelerating Lattice Boltzmann Fluid Flow Simulations Using Graphics Processors	550
<i>Peter Bailey, Joe Myre, Stuart D.C. Walsh, David J. Lilja, and Martin O. Saar</i>	
Session 8C: Information Retrieval I	
An Efficient Collaborative Filtering Approach Using Smoothing and Fusing	558
<i>Daqiang Zhang, Jiannong Cao, Jingyu Zhou, Minyi Guo, and Vaskar Raychoudhury</i>	
End-to-End Study of Parallel Volume Rendering on the IBM Blue Gene/P	566
<i>Tom Peterka, Hongfeng Yu, Robert Ross, Kwan-Liu Ma, and Rob Latham</i>	
Group Operation Assembly Language - A Flexible Way to Express Collective Communication	574
<i>Torsten Hoefler, Christian Siebert, and Andrew Lumsdaine</i>	

Author Index