

Proceedings

37th International Conference on Parallel Processing

*Portland, Oregon
September 9-12, 2008*

Sponsored by

The International Association for Computers and Communications (IACC)

In cooperation with

Portland State University, USA
The Ohio State University, USA



Los Alamitos, California
Washington • Tokyo



37th International Conference on Parallel Processing

ICPP 2008

Table of Contents

Message from the General Chair
Message from the Program Chair
Organizing Committee
Program Vice-Chairs and Program Committee Members
List of Reviewers

Session 1A: Communication and Data Dissemination in P2P Systems

Towards Minimum Traffic Cost and Minimum Response Latency: A Novel Dynamic Query Protocol in Unstructured P2P Networks	1
<i>Chen Tian, Hongbo Jiang, Xue Liu, Wenyu Liu, and Yi Wang</i>	
Flash Data Dissemination in Unstructured Peer-to-Peer Networks	9
<i>Antonis Papadimitriou and Alex Delis</i>	
Fast Source Switching for Gossip-Based Peer-to-Peer Streaming	17
<i>Zhenhua Li, Jiannong Cao, Guihai Chen, and Yan Liu</i>	

Session 1B: Multi-Core Systems

TFlux: A Portable Platform for Data-Driven Multithreading on Commodity Multicore Systems	25
<i>Kyriakos Stavrou, Marios Nikolaidis, Demos Pavlou, Samer Arandi, Paraskevas Evripidou, and Pedro Trancoso</i>	
Enabling Streaming Remoting on Embedded Dual-Core Processors	35
<i>Kun-Yuan Hsieh, Yen-Chih Liu, Po-Wen Wu, Shou-Wei Chang, and Jenq Kuen Lee</i>	
Scalability Evaluation and Optimization of Multi-Core SIP Proxy Server	43
<i>Zou Jia, Zhiyong Liang, and Yiqi Dai</i>	

Session 1C: Software Transactional Memory

DiSTM: A Software Transactional Memory Framework for Clusters	51
<i>Christos Kotselidis, Mohammad Ansari, Kim Jarvis, Mikel Luján, Chris Kirkham, and Ian Watson</i>	
Implementing and Exploiting Inevitability in Software Transactional Memory	59
<i>Michael F. Spear, Michael Silverman, Luke Dalessandro, Maged M. Michael, and Michael L. Scott</i>	
Scalable Techniques for Transparent Privatization in Software Transactional Memory	67
<i>Virendra J. Marathe, Michael F. Spear, and Michael L. Scott</i>	

Session 2A: Parallel Applications

Parallel Inferencing for OWL Knowledge Bases	75
<i>Ramakrishna Soma and V.K. Prasanna</i>	
Optimizing JPEG2000 Still Image Encoding on the Cell Broadband Engine	83
<i>Seunghwa Kang and David A. Bader</i>	
Bandwidth-Efficient Continuous Query Processing over DHTs	91
<i>Yingwu ZHU</i>	

Session 2B: Job Scheduling

Improving Priority Enforcement via Non-Work-Conserving Scheduling	99
<i>Juan Carlos Sáez, José Ignacio Gomez, and Manuel Prieto</i>	
An Incentive-Compatible Mechanism for Scheduling Non-Malleable Parallel Jobs with Individual Deadlines	107
<i>Thomas E. Carroll and Daniel Grosu</i>	
Thermal Management for 3D Processors via Task Scheduling	115
<i>Xiuyi Zhou, Yi Xu, Yu Du, Youtao Zhang, and Jun Yang</i>	

Session 2C: Parallel Programming

Scalable Dynamic Load Balancing Using UPC	123
<i>Stephen Olivier and Jan Prins</i>	
Thread-Sensitive Modulo Scheduling for Multicore Processors	132
<i>Lin Gao, Quan Hoang Nguyen, Lian Li, Jingling Xue, and Tin-Fook Ngai</i>	
Mapping Algorithms for Multiprocessor Tasks on Multi-Core Clusters	141
<i>Jörg Dümmler, Thomas Rauber, and Gudula Rünger</i>	

Session 3A: Fault Tolerance and Reliability

On Modeling Fault Tolerance of Gossip-Based Reliable Multicast Protocols	149
<i>Xiaopeng Fan, Jiannong Cao, Weigang Wu, and Michel Raynal</i>	
Dynamic Meta-Learning for Failure Prediction in Large-Scale Systems: A Case Study	157
<i>Jiexing Gu, Ziming Zheng, Zhiling Lan, John White, Eva Hocks, and Byung-Hoon Park</i>	

On the Reliability of Large-Scale Distributed Systems — A Topological View	165
<i>Yuan He, Hao Ren, Yunhao Liu, and Baijian Yang</i>	

Session 3B: Simultaneous Multithreading Architectures

MFLUSH: Handling Long-Latency Loads in SMT On-Chip Multiprocessors	173
<i>Carmelo Acosta, Francisco J. Cazorla, Alex Ramirez, and Mateo Valero</i>	
Two-Level Reorder Buffers: Accelerating Memory-Bound Applications on SMT Architectures	182
<i>Jason Loew and Dmitry Ponomarev</i>	
Optimizing Issue Queue Reliability to Soft Errors on Simultaneous Multithreaded Architectures	190
<i>Xin Fu, Wangyuan Zhang, Tao Li, and José Fortes</i>	

Session 3C: Grid Computing

A List-Based Strategy for Optimal Replica Placement in Data Grid Systems	198
<i>Yi-Fang Lin, Jan-Jan Wu, and Pangfeng Liu</i>	
Ocean-Atmosphere Modelization over the Grid	206
<i>Y. Caniou, E. Caron, G. Charrier, A. Chis, F. Desprez, and É. Maisonnave</i>	
Optimized Workflow Orchestration of Database Aggregate Operations on Heterogenous Grids	214
<i>Werner Mach and Erich Schikuta</i>	

Session 4A: Message Passing

Designing an Efficient Kernel-Level and User-Level Hybrid Approach for MPI Intra-Node Communication on Multi-Core Systems	222
<i>Lei Chai, Ping Lai, Hyun-Wook Jin, and Dhabaleswar K. Panda</i>	
Detecting Patterns in MPI Communication Traces	230
<i>Robert Preissl, Thomas Köckerbauer, Martin Schulz, Dieter Kranzlmüller, Bronis R. de Supinski, and Daniel J. Quinlan</i>	
VELO: A Novel Communication Engine for Ultra-Low Latency Message Transfers	238
<i>Heiner Litz, Holger Froening, Mondrian Nuessle, and Ulrich Bruening</i>	

Session 4B: Scheduling Algorithms

Realistic Models and Efficient Algorithms for Fault Tolerant Scheduling on Heterogeneous Platforms	246
<i>Anne Benoit, Mourad Hakem, and Yves Robert</i>	
A Duplication Based Algorithm for Optimizing Latency Under Throughput Constraints for Streaming Workflows	254
<i>Nagavijayalakshmi Vydyanathan, Umit Catalyurek, Tahsin Kurc, Ponnuswamy Sadayappan, and Joel Saltz</i>	

Scheduling CPU-Intensive Grid Applications Using Partial Information	262
<i>Nelson Nóbrega-Júnior, Leonardo Assis, and Francisco Brasileiro</i>	

Session 4C: Optimization for Parallelization

Adaptive Software Speculation for Enhancing the Cost-Efficiency of Behavior-Oriented Parallelization	270
---	-----

Yunlian Jiang and Xipeng Shen

Application of Automatic Parallelization to Modern Challenges of Scientific Computing Industries	279
---	-----

Brian Armstrong and Rudolf Eigenmann

Achieving Multi-Level Parallelism in the Filter-Labeled Stream Programming Model	287
---	-----

*George Teodoro, Daniel Fireman, Dorgival Guedes, Wagner Meira Jr.,
and Renato Ferreira*

Session 5A: Interconnects

Address Compression and Heterogeneous Interconnects for Energy-Efficient High-Performance in Tiled CMPs	295
--	-----

Antonio Flores, Manuel E. Acacio, and Juan L. Aragón

Performance of HPC Middleware over InfiniBand WAN	304
---	-----

S. Narravula, H. Subramoni, P. Lai, R. Noronha, and D.K. Panda

Network Reconfiguration Suitability for Scientific Applications	312
---	-----

Héctor Montaner, Federico Silla, Vicente Santonja, and José Duato

Session 5B: Cluster Computing

Optimization of All-to-All Communication on the Blue Gene/L Supercomputer	320
--	-----

Sameer Kumar, Yogish Sabharwal, Rahul Garg, and Philip Heidelberger

A Performance Counter Based Workload Characterization on Blue Gene/P	330
--	-----

Karthik Ganesan, Lizy John, Valentina Salapura, and James Sexton

Tracking Nanostructural Evolution in Alloys: Large-Scale Analysis of Atom Probe Tomography Data on Blue Gene/L	338
---	-----

*Sudip Seal, Michael Moody, Anna Ceguerra, Simon Ringer, Krishna Rajan,
and Srinivas Aluru*

Session 5C: Parallel Algorithms

Parallel Construction of Bidirected String Graphs for Genome Assembly	346
---	-----

Benjamin G. Jackson and Srinivas Aluru

An Efficient Parallel Algorithm for the Multiple Longest Common Subsequence (MLCS) Problem	354
---	-----

Dmitry Korkin, Qingguo Wang, and Yi Shang

XMT-GPU: A PRAM Architecture for Graphics Computation	364
<i>Thomas M. DuBois, Bryant Lee, Yi Wang, Marc Olano, and Uzi Vishkin</i>	

Session 6A: Network Algorithms

A Distributed Context-Free Language Constrained Shortest Path Algorithm	373
<i>Charles B. Ward, Nathan M. Wiegand, and Phillip G. Bradford</i>	
On Clustering Tasks in IC-Optimal Dags	381
<i>Mark Sims, Gennaro Cordasco, and Arnold L. Rosenberg</i>	
Prefix Computation and Sorting in Dual-Cube	389
<i>Yamin Li, Shietung Peng, and Wanming Chu</i>	

Session 6B: OS/Resource Management

Non-Blocking Concurrent FIFO Queues with Single Word Synchronization Primitives	397
<i>Claude Evequoz</i>	
Memory Access Scheduling Schemes for Systems with Multi-Core Processors	406
<i>Hongzhong Zheng, Jiang Lin, Zhao Zhang, and Zhichun Zhu</i>	
Resource Allocation for Distributed Streaming Applications	414
<i>Qian Zhu and Gagan Agrawal</i>	
Exploring Parallel I/O Concurrency with Speculative Prefetching	422
<i>Yong Chen, Surendra Byna, Xian-He Sun, Rajeev Thakur, and William Gropp</i>	

Session 6C: Parallel Simulations

A Scalable Architecture for Crowd Simulation: Implementing a Parallel Action Server	430
<i>G. Viguera, M. Lozano, C. Pérez, and J.M. Orduna</i>	
A Multiway Partitioning Algorithm for Parallel Gate Level Verilog Simulation	438
<i>Lijun Li and Carl Tropper</i>	
TPTS: A Novel Framework for Very Fast Manycore Processor Architecture Simulation	446
<i>Sangyeun Cho, Socrates Demetriades, Shayne Evans, Lei Jin, Hyunjin Lee, Kiyeon Lee, and Michael Moeng</i>	

Session 7A: Network Resource Management

A Replication Overlay Assisted Resource Discovery Service for Federated Systems	454
<i>Hao Yang, Fan Ye, and Zhen Liu</i>	
IMCa: A High Performance Caching Front-End for GlusterFS on InfiniBand	462
<i>Ranjit Noronha and Dhabaleswar K. Panda</i>	
Location Dependent Cooperative Caching in MANET	470
<i>Yilin Wang, Edward Chan, Wenzhong Li, and Sanglu Lu</i>	

Session 7B: Architecture

Accommodation of the Bandwidth of Large Cache Blocks Using Cache/Memory Link Compression	478
<i>Martin Thuresson and Per Stenstrom</i>	
Taming Single-Thread Program Performance on Many Distributed On-Chip L2 Caches	487
<i>Lei Jin and Sangyeun Cho</i>	
Machine Learning Models to Predict Performance of Computer System Design Alternatives	495
<i>Berkin Ozisikyilmaz, Gokhan Memik, and Alok Choudhary</i>	

Session 7C: Numerical Parallel Algorithms

Challenges and Advances in Parallel Sparse Matrix-Matrix Multiplication	503
<i>Aydin Buluc and John R. Gilbert</i>	
Improving the Performance of Multithreaded Sparse Matrix-Vector Multiplication Using Index and Value Compression	511
<i>Kornilios Kourtis, Georgios Goumas, and Nectarios Koziris</i>	
On the Design of Fast Pseudo-Random Number Generators for the Cell Broadband Engine and an Application to Risk Analysis	520
<i>David A. Bader, Aparna Chandramowlishwaran, and Virat Agarwal</i>	

Session 8A: Algorithms

Cellular ANTomata: Food-Finding and Maze-Threading	528
<i>Arnold L. Rosenberg</i>	
Solving Large, Irregular Graph Problems Using Adaptive Work-Stealing	536
<i>Guojing Cong, Sreedhar Kodali, Sriram Krishnamoorthy, Doug Lea, Vijay Saraswat, and Tong Wen</i>	

Session 8B: File Systems and I/O

Impacts of Indirect Blocks on Buffer Cache Energy Efficiency	546
<i>Jianhui Yue, Yifeng Zhu, and Zhao Cai</i>	
Bridging the Gap Between Parallel File Systems and Local File Systems: A Case Study with PVFS	554
<i>Peng Gu, Jun Wang, and Robert Ross</i>	
ParColl: Partitioned Collective I/O on the Cray XT	562
<i>Weikuan Yu and Jeffrey Vetter</i>	

Session 8C: Software Tools

The MAP3S Static-and-Regular Mesh Simulation and Wavefront Parallel-Programming Patterns	570
<i>Robert Niewiadomski, José Nelson Amaral, and Duane Szafron</i>	
Overcoming Scalability Challenges for Tool Daemon Launching	578
<i>Dong H. Ahn, Dorian C. Arnold, Bronis R. de Supinski, Gregory L. Lee, Barton P. Miller, and Martin Schulz</i>	
Scioto: A Framework for Global-View Task Parallelism	586
<i>James Dinan, Sriram Krishnamoorthy, D. Brian Larkins, Jarek Nieplocha, and P. Sadayappan</i>	

Session 9A: Routing

On the Potentials of Segment-Based Routing for NoCs	594
<i>Andres Mejia, Jose Flich, and Jose Duato</i>	
Utility-Based Distributed Routing in Intermittently Connected Networks	604
<i>Ze Li and Haiying Shen</i>	
Deadlock-Free Fully Adaptive Routing in Tori Based on a New Virtual Network Partitioning Scheme	612
<i>Dong Xiang, Qi Wang, and Yi Pan</i>	

Session 9B: Parallel Computing

On-the-Fly Recovery of Job Input Data in Supercomputers	620
<i>Chao Wang, Zhe Zhang, Sudharshan S. Vazhkudai, Xiaosong Ma, and Frank Mueller</i>	
Parallelization and Characterization of Probabilistic Latent Semantic Analysis	628
<i>Chuntao Hong, Wenguang Chen, Weimin Zheng, Jiulong Shan, Yurong Chen, and Yimin Zhang</i>	
Maotai: View-Oriented Parallel Programming on CMT Processors	636
<i>Jiaqi Zhang, Zhiyi Huang, Wenguang Chen, Qihang Huang, and Weimin Zheng</i>	

Session 9C: Search and Security in P2P Systems

Bounded LSH for Similarity Search in Peer-to-Peer File Systems	644
<i>Yu Hua, Bin Xiao, Dan Feng, and Bo Yu</i>	
The Content Pollution in Peer-to-Peer Live Streaming Systems: Analysis and Implications	652
<i>Sirui Yang, Hai Jin, Bo Li, Xiaofei Liao, Hong Yao, and Xuping Tu</i>	
GeWave: Geographically-Aware Wave for File Consistency Maintenance in P2P Systems	660
<i>Haiying Shen</i>	

Author Index