



12th International Conference on Parallel and Distributed Systems

**12-15 July 2006
Minneapolis, Minnesota**

Sponsored by

**IEEE Computer Society Technical Committee on Parallel Processing
IEEE Computer Society Technical Committee on Distributed Processing**



Los Alamitos, California
Washington • Tokyo

Copyright © 2006 by The Institute of Electrical and Electronics Engineers, Inc.

All rights reserved

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US 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 republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 133, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society, or the Institute of Electrical and Electronics Engineers, Inc.

IEEE Computer Society Order Number P2612

ISBN 0-7695-2612-8

ISBN-13 978-0-7695-2612-6

ISSN 1521-9097

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1 800 272 6657
Fax: + 1 714 821 4641
<http://computer.org/cspress>
csbooks@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1 732 981 0060
Fax: + 1 732 981 9667
[http://shop.ieee.org/store/
customer-service@ieee.org](http://shop.ieee.org/store/customer-service@ieee.org)


IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81 3 3408 3118
Fax: + 81 3 3408 3553
tokyo.ofc@computer.org

Individual paper REPRINTS may be ordered at: reprints@computer.org

Editorial production by Silvia Ceballos

Cover art production by Joseph Daigle/Studio Productions

Printed in the United States of America by Commercial Communications, Inc.



IEEE Computer Society

Conference Publishing Services

<http://www.computer.org/proceedings/>



Table of Contents

12th International Conference on Parallel and Distributed Systems

Message from the General Chair	xi
Message from the Program Chair	xii
Organizing Committee	xiv
Reviewers	xviii

Keynote #1

Blue Gene/L, a Massively Parallel Supercomputer: An Applications Point-of-View

Carlos P. Sosa, IBM and Univ. of Minnesota Supercomputing Institute

Session I-A: Heterogeneous Scheduling

Critical Path and Area Based Scheduling of Parallel Task Graphs on Heterogeneous Platforms	3
<i>Tchimou N'takpé and Frédéric Suter</i>	

Efficient Compile-Time Task Scheduling for Heterogeneous Distributed Computing Systems	11
<i>Mohammad I. Daoud and Nawwaf Kharma</i>	

Session I-B: Stabilization Algorithms

Verification of Orbitally Self-Stabilizing Distributed Algorithms Using Lyapunov Functions and Poincaré Maps	23
<i>Abhishek Dhama, Jens Oehlerking, and Oliver Theel</i>	

Fast Convergence in Self-Stabilizing Wireless Networks	31
<i>N. Mitton, É. Fleury, I. Guérin Lassous, B. Sericola, and S. Tixeuil</i>	

Snap-Stabilizing PIF and Useless Computations.....	39
<i>Alain Cournier, Stéphane Devismes, and Vincent Villain</i>	

Session I-C: Availability and Security

Efficient Updates in Highly Available Distributed Random Access Memory	49
<i>Damian Cieslicki, Stefan Schäeckeler, and Thomas Schwarz</i>	

iDIBS: An Improved Distributed Backup System	58
<i>Faruck Morcos, Thidapat Chantem, Philip Little, Tiago Gasiba, and Douglas Thain</i>	

Research on Object-Storage-Based Intrusion Detection	68
<i>Youhui Zhang and Dongsheng Wang</i>	

Session I-D: MPI & Shared Memory

LMPI: MPI for Heterogeneous Embedded Distributed Systems	79
<i>Adnan Agbaria, Dong-In Kang, and Karandeep Singh</i>	
Scalable Hybrid Designs for Linear Algebra on Reconfigurable Computing Systems	87
<i>Ling Zhuo and Viktor K. Prasanna</i>	
A Dynamic Lock Protocol for Scope-Consistency sDSM Systems	96
<i>Artemis A. Christopoulou and Eleftherios D. Polychronopoulos</i>	

Session II-A: Scheduling and Resource Management

Dynamic Resource Management Heuristics for Minimizing Makespan while Maintaining an Acceptable Level of Robustness in an Uncertain Environment	107
<i>Ashish M. Mehta, Jay Smith, H. J. Siegel, Anthony A. Maciejewski, Arun Jayaseelan, and Bin Ye</i>	
Resource Allocation Framework for Distributed Real-Time End-to-End Tasks	115
<i>Chin-Fu Kuo, Chi-Sheng Shih, and Tei-Wei Kuo</i>	
The Impact of Information Availability and Workload Characteristics on the Performance of Job Co-allocation in Multi-clusters	123
<i>William M. Jones, Walter B. Ligon III, and Nishant Shrivastava</i>	

Session II-B: Parallel and Distributed Algorithms I

Parallel Leap: Large-Scale Maximal Pattern Mining in a Distributed Environment	135
<i>Mohammad El-Hajj and Osmar R. Zaïane</i>	
Scalable Parallel Implementation of Exact Inference in Bayesian Networks	143
<i>Vasanth Krishna Namasivayam and Viktor K. Prasanna</i>	
A Parallel Independent Component Analysis Algorithm	151
<i>Hongtao Du, Hairong Qi, and Xiaoling Wang</i>	

Session II-C: Web Performance

Secure and High-Performance Web Server System for Shared Hosting Service	161
<i>Daisuke Hara and Yasuichi Nakayama</i>	
A Pure Nash Equilibrium Guaranteeing Game Theoretical Replica Allocation Method for Reducing Web Access Time	169
<i>Samee Ullah Khan and Ishfaq Ahmad</i>	
Experiences with Simulations—A Light and Fast Model for Secure Web Applications	177
<i>Ramon Nou, Jordi Guitart, David Carrera, and Jordi Torres</i>	

Session II-D: Data Dissemination in Sensor Networks

Multicommodity Flow Based Maximum Lifetime Routing in Wireless Sensor Network	187
<i>Kumar Padmanabh and Rajarshi Roy</i>	
On the Effects of Consistency in Data Operations in Wireless Sensor Networks	195
<i>Kewei Sha and Weisong Shi</i>	
Efficient Cache Replacement in Mobile Environment Using Data Profit	203
<i>Narottam Chand, R.C. Joshi, and Manoj Misra</i>	

Session III-A: QoS in Interconnection Networks

Destination-Based HoL Blocking Elimination 213
T. Natchiondo and J. Flich J. Duato

Providing Quality of Service over Advanced Switching 223
Raul Martinez, Francisco Jose Alfaro, and Jose Luis Sanchez

Session III-B: Streaming

OBN: Peering for Finding Suppliers in P2P On-Demand Streaming Systems 235
Chi-Shiang Liao, Wen-Hung Sun, Chung-Ta King, and Hung-Chang Hsiao

Distributed Coordination for Scalable Multi-source Streaming Model 243
Satoshi Itaya, Naohiro Hayashibara, Tomoya Enokido, and Makoto Takizawa

Session III-C: Sensor Network Deployment

A Delaunay Triangulation Based Method for Wireless Sensor Network Deployment 253
Chun-Hsien Wu, Kuo-Chuan Lee, and Yeh-Ching Chung

Optimal Movement-Assisted Sensor Deployment and Its Extensions
in Wireless Sensor Networks 261
Jie Wu and Shuhui Yang

Panel:

How will we Develop and Program Emerging Robust, Low-Power, Adaptive Multicore Computing Systems?

Panelists:

- David Lilja, U. of Minnesota, *Panel Moderator*
- Dennis Abts, Cray, Inc.
- Mark Hill, U. of Wisconsin, Madison
- Margaret Martonosi, Princeton University
- Jose Moreira, IBM
- D.K. Panda, Ohio State University

Keynote #2

Mobile Sensor Networks & the Princeton ZebraNet Project: Experiences and Challenges
Margaret Martonosi, Princeton University

Session IV-A: Power-aware Protocols

Energy-Efficient and Fault-tolerant Resolution of Topographic Queries in Networked
Sensor Systems 271
Mitali Singh and Viktor K. Prasanna

A Novel Software Framework for Power-Aware Reconfiguration in Distributed Embedded
Smart Cameras 281
Andreas Doblander, Arnold Maier, Bernhard Rinner, and Helmut Schwabach

Session IV-B: Parallel & Distributed Algorithms II

Acceleration of Maximum Likelihood Estimation for Tomosynthesis Mammography 291
Juemin Zhang, Waleed Meleis, David Kaeli, and Tao Wu

A Performance Study of Group Editing Algorithms 300
Du Li and Rui Li

Distributed Algorithms for Building Hamiltonian Cycles in k -ary n -cubes and Hypercubes with Faulty Links	308
<i>Iain A. Stewart</i>	
Session IV-C: Grids	
Policy-Directed Data Movement in Grids	319
<i>Jun Feng, Lingling Cui, Glenn Wasson, and Marty Humphrey</i>	
Flexible Resource Reservation Using Slack Time for Service Grid	327
<i>Chunning Hu, Jinpeng Huai, and Tianyu Wo</i>	
A Pre-evaluating Set based Bias-tuned Reputation Evaluation Method for Trust Establishment in Grid	335
<i>Xiangli Qu and Xuejun Yang</i>	
Session IV-D: Multiprocessor Scheduling	
Memory and Network Bandwidth Aware Scheduling of Multiprogrammed Workloads on Clusters of SMPs	345
<i>Evangelos Koukis and Nectarios Koziris</i>	
Accuracy versus Migration Overhead in Real-Time Multiprocessor Reweighting Algorithms	355
<i>Aaron Block and James H. Anderson</i>	
A Dynamic Critical Path Duplication Task Scheduling Algorithm for Distributed Heterogeneous Computing Systems	365
<i>Chun-Hsien Liu, Chia Feng Li, Kuan-Chou Lai, and Chao-Chin Wu</i>	
Session V-A: Multicore Scheduling	
Loop Scheduling with Complete Memory Latency Hiding on Multi-core Architecture	375
<i>Chun Xue, Zili Shao, Meilin Liu, Meikang Qiu, and Edwin H.-M. Sha</i>	
Multi-Level On-Chip Memory Hierarchy Design for Embedded Chip Multiprocessors	383
<i>Ozcan Ozturk, Mahmut Kandemir, Mary Jane Irwin, and Suleyman Tosun</i>	
SPM Conscious Loop Scheduling for Embedded Chip Multiprocessors	391
<i>Liping Xue, Mahmut Kandemir, Guangyu Chen, and Talyan Yemliha</i>	
Session V-B: P2P Networks I	
Reconciliation in the APPA P2P System	401
<i>Vidal Martins, Reza Akbarinia, Esther Pacitti, and Patrick Valduriez</i>	
Load Balancing with Multiple Hash Functions in Peer-to-Peer Networks	411
<i>Ye Xia, Shigang Chen, and Vivekanand Korgaonkar</i>	
XYZ: A Scalable, Partially Centralized Lookup Service for Large-Scale Peer-to-Peer Systems	421
<i>Jianying Zhang and Jie Wu</i>	
Session V-C: Switch Design	
Hardware Efficient Two Step Iterative Matching Algorithms for VOQ Switches	431
<i>Deng Pan and Yuanyuan Yang</i>	

Scalable Low-Cost QoS Support for Single-Chip Switches	439
<i>Alejandro Martinez, Francisco Jose Alfaro, Jose Luis Sánchez Garcia, and Jose Duato</i>	
A Local Coefficient Based Load Sensitive Routing Protocol for Providing QoS	447
<i>Anunay Tiwari and Anirudha Sahoo</i>	
Session VI-A: Load Balancing	
A Distributed Algorithm for Sharing Web Cache Disk Capacity	457
<i>George Alyfantis, Stathes Hadjiefthymiades, Lazaros Merakos, and Panagiotis Kostopoulos</i>	
Optimal Placement of Replicas in Data Grid Environments with Locality Assurance	465
<i>Yi-Fang Lin, Pangfeng Liu, and Jan-Jan Wu</i>	
Session VI-B: Parallel Programming Tools	
PARSE: A Tool for Parallel Application Run Time Sensitivity Evaluation	475
<i>Jeffrey J. Evans and Cynthia S. Hood</i>	
Productivity and Performance in Parallel Programming Environments Using a Novel Virtual Machine Framework	485
<i>Raja Neogi</i>	
Session VI-C (short papers): Communication Networks	
Incentive Compatible Cost- and Stability-Based Routing in Ad Hoc Networks	495
<i>Mingming Lu, Feng Li, and Jie Wu</i>	
SF-RED: A novel server-based AQM to provide inter-server fairness service	501
<i>Hsien-Ming Wu, Chin-Chi Wu, and Woei Lin</i>	
Fast and Robust TCP Session Lookup by Digest Hashing	507
<i>Fong Pong</i>	
Keynote #3	
Emerging Networking Technologies and Protocols for Next Generation Clusters	
<i>Dhabaleswar K. Panda, The Ohio State University</i>	
Banquet Address	
HPC: Been There, Done That?	
<i>Craig Stunkel, Deep Computing Software & Applications, IBM T.J. Watson Research Center</i>	
Keynote #4	
An Overview of Middleware Architectures for Scheduling Grid Applications	
<i>Frédéric Desprez, INRIA-ENS Lyon, France</i>	
Session VII-A: Fault Tolerance	
Reachability-Based Fault-Tolerant Routing	515
<i>J. M. Montanana, J. Flich, A. Robles, and J. Duato</i>	
Software-Based Adaptive and Concurrent Self-Testing in Programmable Network Interfaces	525
<i>Yizheng Zhou, Vijay Lakamraju, Israel Koren, and C.M. Krishna</i>	
Design, Analysis and Performance Evaluation of a New Algorithm for Developing a Fault Tolerant Distributed System	533
<i>Umasankar Malladi</i>	

Session VII-B: P2P Networks II

Oriented Overlays for Clustering Client Requests to Data-Centric Network Services..... 545
Congchun He and Vijay Karamcheti

RASTER: A Light-Weight Routing Protocol to Discover Shortest Overlay Routes
in Randomized DHT Systems 553
Chih-Chiang Wang and Khaled Harfoush

Preferential and Strata Based P2P Model: Selfishness to Altruism and Fairness 561
Darshan Purandare and Ratan Guha

Session VII-C: Resource Monitoring

Job Centric Cluster Monitoring 571
Roger Curry and Rob Simmonds

The Node Monitoring Component of a Scalable Systems Software Environment..... 579
Sam Miller and Brett Bode

Coloring the Internet: IP Traceback 589
M. Muthuprasanna, G. Manimaran, Mansoor Alicherry, and Vijay Kumar

Session VII-D: Ad Hoc and Wireless Networks

Admission Control and Bandwidth Allocation above Packet Level
for IEEE 802.16 Wireless MAN 599
Haitang Wang, Bing He, and Dharma P. Agrawal

Utilizing Call Admission Control for Deriving Optimal Pricing of Multiple
Service Classes in Mobile Wireless Cellular Networks 605
Okan Yilmaz and Ing-Ray Chen

Agent-Based Cooperative Anomaly Detection for Wireless Ad hoc Networks 613
Hongmei Deng, Roger Xu, Jason Li, Frank Zhang, Renato Levy, and Wenke Lee

Author Index 621