2010 39th International Conference on Parallel Processing

(ICPP 2010)

San Diego, California, USA
13-16 September 2010

IEEE Catalog Number: CFP10127-PRT
ISBN: 978-1-4244-7913-9

2010 39th International Conference on Parallel Processing

ICPP 2010

Table of Contents

Message from the Program Chair.....xii

Organizing Committee.....xiv

Reviewers.....xvii

Panel Abstract.....xxii

Session 1A: Architecture I

Achieving Fair or Differentiated Cache Sharing in Power-Constrained Chip Multiprocessors.....1
Xiaorui Wang, Kai Ma, and Yefu Wang

A Theoretical Framework for Value Prediction in Parallel Systems.....11
Shaoshan Liu, Christine Eisenbeis, and Jean-Luc Gaudiot

Heterogeneous Mini-rank: Adaptive, Power-Efficient Memory Architecture.....21
Kun Fang, Hongzhong Zheng, and Zhichun Zhu

Session 1B: Compilers, Programming Models and Languages I

Gossamer: A Lightweight Approach to Using Multicore Machines.....30
Joseph A. Roback and Gregory R. Andrews

Toward Harnessing DOACROSS Parallelism for Multi-GPGPUs.....40
Peng Di, Qing Wan, Xuemeng Zhang, Hui Wu, and Jingling Xue

Exploitation of Dynamic Communication Patterns through Static Analysis.....51
Robert Preissl, Bronis R. de Supinski, Martin Schulz, Daniel J. Quinlan, Dieter Kranzlmüller, and Thomas Panas

Session 1C: Algorithms and Applications I

Parallel Exact Inference on a CPU-GPGPU Heterogenous System.....61
Hyeran Jeon, Yinglong Xia, and Viktor K. Prasanna

Minimizing Stretch and Makespan of Multiple Parallel Task Graphs via Malleable Allocations.....71
Henri Casanova, Frédéric Desprez, and Frédéric Suter

Efficient PageRank and SpMV Computation on AMD GPUs.....81
Tianji Wu, Bo Wang, Yi Shan, Feng Yan, Yu Wang, and Ningyi Xu

Session 2A: Performance Evaluation and Simulation I

Identifying the Root Causes of Wait States in Large-Scale Parallel Applications 90
David Böhme, Markus Geimer, Felix Wolf, and Lukas Arnold

Energy Modeling of Wireless Sensor Nodes Based on Petri Nets 101
Ali Shareef and Yifeng Zhu

Optimal Task Reallocation in Heterogeneous Distributed Computing Systems with Age-Dependent Delay Statistics 111
Jorge E. Pezoa, Majeed M. Hayat, Zhuoyao Wang, and Sagar Dhakal

Session 2B: OS/Resource Management I

Rearchitecting MapReduce for Heterogeneous Multicore Processors with Explicitly Managed Memories 121
Anastasios Papagiannis and Dimitrios S. Nikolopoulos

System-Level, Unified In-band and Out-of-band Dynamic Thermal Control 131
Dong Li, Rong Ge, and Kirk Cameron

Microwiper: Efficient Memory Propagation in Live Migration of Virtual Machines 141
Yuyang Du, Hongliang Yu, Guangyu Shi, Jian Chen, and Weimin Zheng

Session 2C: Algorithms and Applications II

Reliability and Performance Optimization of Pipelined Real-Time Systems 150
Anne Benoit, Fanny Dufossé, Alain Girault, and Yves Robert

An Efficient Randomized Routing Protocol for Single-Hop Radio Networks 160
Sanguthevar Rajasekaran, Dolly Sharma, Reda Ammar, and Nicholas Lownes

Checkpointing vs. Migration for Post-Petascale Supercomputers 168
Franck Cappello, Henri Casanova, and Yves Robert

Session 3A: Wireless Networks and Pervasive Computing

Revisting Tag Collision Problem in RFID Systems 178
Lei Yang, Jinsong Han, Yong Qi, Cheng Wang, Yunhao Liu, Ying Cheng, and Xiao Zhong

Distributed Minimum Transmission Multicast Routing Protocol for Wireless Sensor Networks 188
Long Cheng, Sajal K. Das, Jiannong Cao, Canfeng Chen, and Jian Ma

A Quantitative Study of Accountability in Wireless Multi-hop Networks 198
Zhifeng Xiao, Yang Xiao, and Jie Wu

Session 3B: Cluster, Grid and Cloud Computing I

A Stack-on-Demand Execution Model for Elastic Computing 208
Ricky K. K. Ma, King Tin Lam, Cho-Li Wang, and Chenggang Zhang

Designing Power-Aware Collective Communication Algorithms for InfiniBand Clusters 218
Krishna Kandalla, Emilio P. Mancini, Sayantan Sur, and Dhabaleswar K. Panda
Starling: Minimizing Communication Overhead in Virtualized Computing Platforms Using Decentralized Affinity-Aware Migration 228
Jason Sonnek, James Greensky, Robert Reutiman, and Abhishek Chandra

Session 3C: P2P and Service-Oriented Architectures I

Dual-Phase Just-in-Time Workflow Scheduling in P2P Grid Systems 238
Sheng Di and Cho-Li Wang
A DHT-Aided Chunk-Driven Overlay for Scalable and Efficient Peer-to-Peer Live Streaming 248
Haiying Shen, Lianyu Zhao, Ze Li, and Jin Li
Towards Building Efficient Content-Based Publish/Subscribe Systems over Structured P2P Overlays 258
Shengdong Zhang, Ji Wang, Rui Shen, and Jie Xu

Session 4A: Architecture II

A G-Line-Based Network for Fast and Efficient Barrier Synchronization in Many-Core CMPs 267
José L. Abellán, Juan Fernández, and Manuel E. Acacio
Hyperscalar: A Novel Dynamically Reconfigurable Multi-core Architecture 277
Jih-Ching Chiu, Yu-Liang Chou, and Po-Kai Chen
Dynamic Switching-Frequency Scaling: Scheduling Overcommitted Domains in Xen VMM 287
Huacai Chen, Hai Jin, Kan Hu, and Jian Huang

Session 4B: Compilers, Programming Models and Languages II

Block-Parallel Programming for Real-Time Embedded Applications 297
David Black-Schaffer and William J. Dally
Automatic Generation of Stream Descriptors for Streaming Architectures 307
Lei Gao, David Zaretsky, Gaurav Mittal, Dan Schonfeld, and Prith Banerjee
Efficient Work Stealing for Fine Grained Parallelism 313
Karl-Filip Faxén

Session 4C: Algorithms and Applications III

Parameterized Schedulability Analysis on Uniform Multiprocessors 323
Risat Mahmud Pathan and Jan Jonsson
A Scalable Parallel Algorithm for Large-Scale Protein Sequence Homology Detection 333
Changjun Wu, Ananth Kalyanaraman, and William R. Cannon
A MapReduce Style Framework for Computations on Trees 343
Abhinav Sarje and Srinivas Aluru

Session 5A: Performance Evaluation and Simulation II

Characterizing the Relation Between Apex-Map Synthetic Probes and Reuse Distance Distributions 353
Khaled Z. Ibrahim and Erich Strohmaier
Extending the Monte Carlo Processor Modeling Technique: Statistical Performance Models of the Niagara 2 Processor 363
Waleed Alkohlani, Jeanine Cook, and Ram Srinivasan
Modelization and Performance Evaluation of the DIET Middleware 375
Eddy Caron, Benjamin Depardon, and Frédéric Desprez

Session 5B: Cluster, Grid and Cloud Computing II

Power Management in Heterogeneous Multi-tier Web Clusters 385
Peijian Wang, Yong Qi, Xue Liu, Ying Chen, and Xiao Zhong
Cyberaide onServe: Software as a Service on Production Grids 395
Tobias Kurze, Lizhe Wang, Gregor von Laszewski, Jie Tao, Marcel Kunze, David Kramer, and Wolfgang Karl
Energy Efficient Prefetching with Buffer Disks for Cluster File Systems 404
Adam Manzanares, Xiaojun Ruan, Shu Yin, Jiong Xie, Zhiyang Ding, Yun Tian, James Majors, and Xiao Qin

Session 5C: Information Retrieval and Knowledge Discovery

DISQO: A Distributed Framework for Spatial Queries over Moving Objects 414
Baihua Zheng, Wang-Chien Lee, Ken C. K. Lee, Julian Winter, and Meng-Chang Chen
Model-Driven Traffic Data Acquisition in Vehicular Sensor Networks 424
Chih-Chieh Hung and Wen-Chih Peng
Distributing a Metric-Space Search Index onto Processors 433
Mauricio Marin, Flavio Ferrarotti, and Veronica Gil-Costa

Session 6A: Architecture III

Cubic Ring Networks: A Polymorphic Topology for Network-on-Chip 443
Bilal Zafar, Jeff Draper, and Timothy M. Pinkston
Speculative Execution on GPU: An Exploratory Study 453
Shaoshan Liu, Christine Eisenbeis, and Jean-Luc Gaudiot
Improving Application Performance and Predictability Using Multiple Virtual Lanes in Modern Multi-core InfiniBand Clusters 462
Hari Subramoni, Ping Lai, Sayantan Sur, and Dhabaleswar K. (DK) Panda

Session 6B: Compilers, Programming Models and Languages III

PacketC Language and Parallel Processing of Masked Databases 472
Ralph Duncan, Peder Jungck, and Kenneth Ross

Handling Conflicts with Compiler's Help in Software Transactional Memory Systems 482
Sandya Mannarswamy and R. Govindarajan

Design and Implementation of a Hybrid Parallel Performance Measurement System 492
Alan Morris, Allen D. Malony, Sameer Shende, and Kevin Huck

Session 6C: Algorithms and Applications IV

Scalability of a Parallel JPEG Encoder on Shared Memory Architectures 502
David Castells-Rufas, Jaume Joven, and Jordi Carrabina

Integer Number Crunching on the Cell Processor 508
Hsieh-Chung Chen, Chen-Mou Cheng, Shih-Hao Hung, and Zong-Cing Lin

Incentive Compatible Online Scheduling of Malleable Parallel Jobs with Individual Deadlines 516
Thomas E. Carroll and Daniel Grosu

Session 7A: Performance Evaluation and Simulation III

Optimizing HPC Fault-Tolerant Environment: An Analytical Approach 525
Hui Jin, Yong Chen, Huaiyu Zhu, and Xian-He Sun

Detailed Load Balance Analysis of Large Scale Parallel Applications 535
Kevin A. Huck and Jesús Labarta

A Machine Learning Approach for Optimizing Parallel Logic Simulation 545
Sina Meraji and Carl Tropper

Session 7B: OS/Resource Management II

A Gray-Box Feedback Control Approach for System-Level Peak Power Management 555
Jiayu Gong and Cheng-Zhong Xu

A Lightweight, GPU-Based Software RAID System 565
Matthew L. Curry, H. Lee Ward, Anthony Skjellum, and Ron Brightwell

Task Assignment with Cache Partitioning and Locking for WCET Minimization on MPSoC 573
Tiantian Liu, Yingchao Zhao, Minming Li, and Chun Jason Xue

Session 7C: Algorithms and Applications V

Massive Social Network Analysis: Mining Twitter for Social Good 583
David Ediger, Karl Jiang, Jason Riedy, David A. Bader, Courtney Corley, Rob Farber, and William N. Reynolds

Subgraph Enumeration in Large Social Contact Networks Using Parallel Color Coding and Streaming 594
Zhao Zhao, Maleq Khan, V. S. Anil Kumar, and Madhav V. Marathe

Lock-Free Multiway Search Trees 604
Michael Spiegel and Paul F. Reynolds Jr.

Session 8A: Data Intensive and I/O Computing

SAM: A Semantic-Aware Multi-tiered Source De-duplication Framework for Cloud Backup 614
Yujuan Tan, Hong Jiang, Dan Feng, Lei Tian, Zhichao Yan, and Guohui Zhou

Hint-K: An Efficient Multi-level Cache Using K-Step Hints 624
Chentao Wu, Xubin He, Qiang Cao, and Changsheng Xie

FlashCoop: A Locality-Aware Cooperative Buffer Management for SSD-Based Storage Cluster 634
Qingsong Wei, Bozhao Gong, Suraj Pathak, and Y. C. Tay

Session 8B: Cluster, Grid and Cloud Computing III

Near-Optimal Rendezvous Protocols for RDMA-Enabled Clusters 644
Matthew Small, Zheng Gu, and Xin Yuan

Performance Management of Accelerated MapReduce Workloads in Heterogeneous Clusters 653
Jordà Polo, David Carrera, Yolanda Becerra, Vicenç Beltran, Jordi Torres, and Eduard Ayguadé

MemX: Virtualization of Cluster-Wide Memory 663
Umesh Deshpande, Beilan Wang, Shafee Haque, Michael Hines, and Kartik Gopalan

Session 8C: P2P and Service-Oriented Architectures II

Using Mobile Mules for Collecting Data from an Isolated Wireless Sensor Network 673
Yu-Chee Tseng, Wan-Ting Lai, Chi-Fu Huang, and Fang-Jing Wu

A Simple Effective Scheme to Enhance the Capability of Web Servers Using P2P Networks 680
Jie Yu, Liming Lu, Zhoujun Li, Xiaofeng Wang, and Jinshu Su

Optimal Overlay Construction on Heterogeneous Live Peer-to-Peer Streaming Systems 690
Min Yang and Yuanyuan Yang

Author Index 699