

# **2010 IEEE International Symposium on Parallel & Distributed Processing**

**(IPDPS 2010)**

**Atlanta, Georgia, USA  
19-23 April 2010**

**Pages 3/8: ;**



**IEEE Catalog Number: CFP10023-PRT  
ISBN: 978-1-4244-6442-5**

# Table of Contents

## 2010 IEEE International Symposium on Parallel & Distributed Processing (IPDPS)

### Algorithms for Network Management

**Distributed Advance Network Reservation with Delay Guarantees**  
Nilloofar Fazlollahi and David Starobinski

**A General Algorithm for Detecting Faults under the Comparison Diagnosis Model**  
Iain A. Stewart

**Broadcasting on Large Scale Heterogeneous Platforms under the Bounded Multi-Port Model**  
Olivier Beaumont, Lionel Eyraud-Dubois and Shailesh Kumar Agrawal

**On the Importance of Bandwidth Control Mechanisms for Scheduling on Large Scale Heterogeneous Platforms**  
Olivier Beaumont and Hejer Rejeb

### Scientific Computing with GPUs

**Improving Numerical Reproducibility and Stability in Large-Scale Numerical Simulations on GPUs**  
Michela Taufer, Omar Padron, Philip Saponaro and Sandeep Patel

**Implementing the Himeno Benchmark with CUDA on GPU Clusters**  
Everett H. Phillips and Massimiliano Fatica

**Direct Self-Consistent Field Computations on GPU Clusters**  
Guochun Shi, Volodymyr Kindratenko, Ivan Ufimtsev and Todd Martinez

**Parallelization of Tau-Leap Coarse-Grained Monte Carlo Simulations on GPUs**  
Lifan Xu, Michela Taufer, Stuart Collins and Dionisios G. Vlachos

### Data Storage and Memory Systems

**DEBAR: A Scalable High-Performance De-duplication Storage System for Backup and Archiving**  
Tianming Yang, Hong Jiang, Dan Feng, Zhongying Niu, Ke Zhou and Yaping Wan

**HPDA: A Hybrid Parity-based Disk Array for Enhanced Performance and Reliability**  
Bo Mao, Hong Jiang, Dan Feng, Suzhen Wu, Jianxi Chen, Lingfang Zeng and Lei Tian

**Fine-Grained QoS Scheduling for PCM-based Main Memory Systems**  
Ping Zhou, Yu Du, Youtao Zhang and Jun Yang

**Performance Impact of Resource Contention in Multicore Systems**  
Robert Hood, Haoqiang Jin, Piyush Mehrotra, Johnny Chang, Jahed Djomehri, Sharad Gaval, Dennis Jespersen, Kenichi Taylor and Rupak Biswas

## Fault Tolerance

### **Improving the Performance of Hypervisor-Based Fault Tolerance**

Jun Zhu, Wei Dong, Zhefu Jiang, Xiaogang Shi, Zhen Xiao and Xiaoming Li

### **Supporting Fault Tolerance in a Data-Intensive Computing Middleware**

Tekin Bicer, Wei Jiang and Gagan Agrawal

### **A High-Performance Fault-Tolerant Software Framework for Memory on Commodity GPUs**

Naoya Maruyama, Akira Nukada and Satoshi Matsuoka

### **Scalable Failure Recovery for High-performance Data Aggregation**

Dorian C. Arnold and Barton P. Miller

## Sorting

### **High Performance Comparison-Based Sorting Algorithm on Many-Core GPUs**

Xiaochun Ye, Dongrui Fan, Wei Lin, Nan Yuan and Paolo Ienne

### **GPU Sample Sort**

Nikolaj Leischner, Vitaly Osipov and Peter Sanders

### **Highly Scalable Parallel Sorting**

Edgar Solomonik and Laxmikant V. Kalé

## Scheduling

### **A Scheduling Framework for Large-Scale, Parallel, and Topology-Aware Applications**

Valentin Kravtsov, Pavel Bar, David Carmeli, Assaf Schuster and Martin Swain

### **Load Regulating Algorithm for Static-Priority Task Scheduling on Multiprocessors**

Risat Mahmud Pathan and Jan Jonsson

### **Scheduling Algorithms for Linear Workflow Optimization**

K. Agrawal, A. Benoit, L. Magnan and Y. Robert

### **Hypergraph-based Task-Bundle Scheduling Towards Efficiency and Fairness in Heterogeneous Distributed Systems**

Han Zhao, Xinxin Liu and Xiaolin Li

## Performance/Scalability Improvement for Scientific Applications

### **Improving the Performance of Uintah: A Large-Scale Adaptive Meshing Computational Framework**

Justin Luitjens and Martin Berzins

### **Optimizing and Tuning the Fast Multipole Method for State-of-the-Art Multicore Architectures**

Aparna Chandramowliswaran, Samuel Williams, Leonid Oliker, Ilya Lashuk, George Biros and Richard Vuduc

## **Parallelization of DQMC Simulation for Strongly Correlated Electron Systems**

Che-Rung Lee, I-Hsin Chung and Zhaojun Bai

## **Parallel I/O Performance: From Events to Ensembles**

Andrew Uselton, Mark Howison, Nicholas J. Wright, David Skinner, Noel Keen, John Shalf, Karen L. Karavanic and Leonid Oliker

## **Network Architecture and Algorithms**

### **Achieve Constant Performance Guarantees using Asynchronous Crossbar Scheduling without Speedup**

Deng Pan, Kia Makki and Niki Pissinou

### **Distributive Waveband Assignment in Multi-granular Optical Networks**

Yang Wang and Xiaojun Cao

### **QoS Aware BiNoC Architecture**

Shih-Hsin Lo, Ying-Cherng Lan, Hsin-Hsien Yeh, Wen-Chung Tsai, Yu-Hen Hu and Sao-Jie Chen

### **First Experiences with Congestion Control in InfiniBand Hardware**

Ernst Gunnar Gran, Magne Eimot, Sven-Arne Reinemo, Tor Skeie, Olav Lysne and Lars Paul Huse

## **Software Support for Using GPUs**

### **Object-Oriented Stream Programming using Aspects**

Mingliang Wang and Manish Parashar

### **Optimal Loop Unrolling For GPGPU Programs**

Giridhar Sreenivasa Murthy, Mahesh Ravishankar, Muthu Manikandan Baskaran and P. Sadayappan

### **Speculative Execution on Multi-GPU Systems**

Gregory Damos and Sudhakar Yalamanchili

### **Dynamic Load Balancing on Single- and Multi-GPU Systems**

Long Chen, Oreste Villa, Sriram Krishnamoorthy and Guang R. Gao

## **Performance Prediction and Benchmarking Tools**

### **Servet: A Benchmark Suite for Autotuning on Multicore Clusters**

Jorge González-Domínguez, Guillermo L. Taboada, Basilio B. Fraguera, María J. Martín and Juan Touriño

### **KRASH: Reproducible CPU Load Generation on Many-Core Machines**

Swann Perarnau and Guillaume Huard

### **Power-aware MPI Task Aggregation Prediction for High-End Computing Systems**

Dong Li, Dimitrios S. Nikolopoulos, Kirk Cameron, Bronis R. de Supinski and Martin Schulz

## Resource Allocation

### **Varying Bandwidth Resource Allocation Problem with Bag Constraints**

Venkatesan T. Chakaravarthy, Vinayaka Pandit, Yogish Sabharwal and Deva P. Seetharam

### **Decentralized Resource Management for Multi-core Desktop Grids**

Jaehwan Lee, Pete Keleher and Alan Sussman

### **Dynamic Fractional Resource Scheduling for HPC Workloads**

Mark Stillwell, Frédéric Vivien and Henri Casanova

### **ADEPT Scalability Predictor in Support of Adaptive Resource Allocation**

Arash Deshmeh, Jacob Machina and Angela Sodan

## Image Processing and Data Mining

### **Exploiting the Forgiving Nature of Applications for Scalable Parallel**

#### **Execution**

Jiayuan Meng, Anand Raghunathan, Srimat Chakradhar and Surendra Byna

### **Fisheye Lens Distortion Correction on Multicore and Hardware Accelerator Platforms**

Konstantis Daloukas, Christos D. Antonopoulos, Nikolaos Bellas and Sek M. Chai

### **Large-Scale Multi-Dimensional Document Clustering on GPU Clusters**

Yongpeng Zhang, Frank Mueller, Xiaohui Cui and Thomas Potok

### **eScience in the Cloud: A MODIS Satellite Data Reprojection and Reduction Pipeline in the Windows Azure Platform**

Jie Li, Marty Humphrey, Deb Agarwal, Keith Jackson, Catharine van Ingen and Youngryel Ryu

## Transactional Memory

### **Locality-Aware Adaptive Grain Signatures for Transactional Memories**

Woojin Choi and Jeff Draper

### **Dynamic Analysis of the Relay Cache-Coherence Protocol for Distributed Transactional Memory**

Bo Zhang and Binoy Ravindran

### **Runtime Checking of Serializability in Software Transactional Memory**

Arnab Sinha and Sharad Malik

### **Consistency in Hindsight: A Fully Decentralized STM Algorithm**

Annette Bieniusa and Thomas Fuhrmann

## Tools for Performance and Correctness Analysis

### **Identifying Ad-hoc Synchronization for Enhanced Race Detection**

Ali Jannesari and Walter F. Tichy

### **Improving the Performance of Program Monitors with Compiler Support in Multi-Core Environment**

Guojin He and Antonia Zhai

### **On-line detection of large-scale parallel application's structure**

German Llort, Juan Gonzalez, Harald Servat, Judit Gimenez and Jesus Labarta

## **Adaptive Sampling-Based Profiling Techniques for Optimizing the Distributed JVM Runtime**

King Tin Lam, Yang Luo and Cho-Li Wang

## **Parallel Linear Algebra I**

### **Algorithmic Cholesky Factorization Fault Recovery**

Doug Hakkarinen and Zizhong Chen

### **Analyzing the Soft Error Resilience of Linear Solvers on Multicore Multiprocessors**

Konrad Malkowski, Padma Raghavan and Mahmut Kandemir

### **Parallel Architecture for Meaning Comparison**

Suneil Mohan, Amitava Biswas, Aalap Tripathy, Jagannath Pannigrahy and Rabi Mahapatra

## **P2P Algorithms**

### **A Hybrid Interest Management Mechanism for Peer-to-Peer Networked Virtual Environments**

Ke Pan, Wentong Cai, Xueyan Tang, Suiping Zhou and Stephen John Turner

### **Attack-Resistant Frequency Counting**

Bo Wu, Jared Saia and Valerie King

### **Overlays with Preferences: Approximation Algorithms for Matching with Preference Lists**

Giorgos Georgiadis and Marina Papatriantafylou

### **Analysis of Durability in Replicated Distributed Storage Systems**

Sriram Ramabhadran and Joseph Pasquale

## **Parallel Solutions for String and Sequence Problems**

### **Scalable Multi-Pipeline Architecture for High Performance Multi-Pattern String Matching**

Weirong Jiang, Yi-Hua E. Yang and Viktor K. Prasanna

### **Head-Body Partitioned String Matching for Deep Packet Inspection with Scalable and Attack-Resilient Performance**

Yi-Hua E. Yang, Viktor K. Prasanna and Chenqian Jiang

### **Parallel de novo Assembly of Large Genomes from High-Throughput Short Reads**

B.G. Jackson, M. Regennitter, X. Yang, P.S. Schnable and S. Aluru

### **Efficient Parallel Algorithms for Maximum-Density Segment Problem**

Xue Wang, Fasheng Qiu, Sushil K. Prasad and Guantao Chen

## **Energy-aware Task Management**

### **Hybrid MPI/OpenMP Power-Aware Computing**

Dong Li, Bronis R. de Supinski, Martin Schulz, Kirk Cameron and Dimitrios S. Nikolopoulos

### **Performance and Energy Optimization of Concurrent Pipelined**

## **Applications**

Anne Benoit, Paul Renaud-Goud and Yves Robert

## **Robust Control-theoretic Thermal Balancing for Server Clusters**

Yong Fu, Chenyang Lu and Hongan Wang

## **A Simple Thermal Model for Multi-core Processors and Its Application to Slack Allocation**

Zhe Wang and Sanjay Ranka

## **Parallel Operating Systems and System Software**

### **GenerOS: An Asymmetric Operating System Kernel for Multi-core Systems**

Qingbo Yuan, Jianbo Zhao, Mingyu Chen and Ninghui Sun

### **Palacios and Kitten: New High Performance Operating Systems For Scalable Virtualized and Native Supercomputing**

John Lange, Kevin Pedretti, Trammell Hudson, Peter Dinda, Zheng Cui, Lei Xia, Patrick Bridges, Andy Gocke, Steven Jaconette, Mike Levenhagen and Ron Brightwell

### **MMT: Exploiting Fine-Grained Parallelism in Dynamic Memory Management**

Devesh Tiwari, Sanghoon Lee, James Tuck and Yan Solihin

### **Optimization of Applications with Non-blocking Neighborhood Collectives via Multisends on the Blue Gene/P Supercomputer**

Sameer Kumar, Philip Heidelberger, Dong Chen and Michael Hines

## **Parallel Graph Algorithms I**

### **A Multi-Source Label-Correcting Algorithm for the All-Pairs Shortest Paths Problem**

Hiroki Yanagisawa

### **Parallel Computation of Best Connections in Public Transportation Networks**

Daniel Delling, Bastian Katz and Thomas Pajor

### **Dynamically Tuned Push-Relabel Algorithm for the Maximum Flow Problem on CPU-GPU-Hybrid Platforms**

Zhengyu He and Bo Hong

### **A Novel Application of Parallel Betweenness Centrality to Power Grid Contingency Analysis**

Shuangshuang Jin, Zhenyu Huang, Yousu Chen, Daniel Chavarría-Miranda, John Feo and Pak Chung Wong

## **Parallel Linear Algebra II**

### **Adapting Communication-Avoiding LU and QR Factorizations to Multicore Architectures**

Simplice Donfack, Laura Grigori and Alok Kumar Gupta

### **QR Factorization of Tall and Skinny Matrices in a Grid Computing Environment**

Emmanuel Agullo, Camille Coti, Jack Dongarra, Thomas Herault and Julien Langou

## **Tile QR Factorization with Parallel Panel Processing for Multicore**

**Architectures**, \* \$

Bilel Hadri, Hatem Ltaief, Emmanuel Agullo and Jack Dongarra

## **Linpack Evaluation on a Supercomputer with Heterogeneous Accelerators**, + \$

Toshio Endo, Akira Nukada, Satoshi Matsuoka and Naoya Maruyama

## **Caches and Caching**

### **Adapting Cache Partitioning Algorithms to Pseudo-LRU Replacement**

**Policies**, +,

Kamil Kędzierski, Miquel Moreto, Francisco J. Cazorla and Mateo Valero

### **Exploiting Set-Level Non-Uniformity of Capacity Demand to Enhance CMP**

**Cooperative Caching**, - \$

Dongyuan Zhan, Hong Jiang and Sharad C. Seth

### **Masking I/O Latency using Application Level I/O Caching and Prefetching on Blue Gene Systems**, - \$\$

Seetharami Seelam, I-Hsin Chung, John Bauer and Hui-Fang Wen

### **Intra-Application Cache Partitioning**, - %&

Sai Prashanth Muralidhara, Mahmut Kandemir and Padma Raghavan

## **Thread Scheduling**

### **SLAW: a Scalable Locality-aware Adaptive Work-stealing Scheduler**, - &(

Yi Guo, Jisheng Zhao, Vincent Cave and Vivek Sarkar

### **Executing Task Graphs Using Work-Stealing**, - ' \*

Kunal Agrawal, Charles E. Leiserson and Jim Sukha

### **Structuring Execution of OpenMP Applications for Multicore Architectures**, - (,

François Broquedis, Olivier Aumage, Brice Goglin, Samuel Thibault, Pierre-Andr Wacrenier and Raymond Namyst

### **Oversubscription on Multicore Processors**, - ) ,

Costin Iancu, Steven Hofmeyr, Filip Blagojević and Yili Zheng

## **Distributed Algorithms**

### **A Scalable Algorithm for Maintaining Perpetual System Connectivity in**

**Dynamic Distributed Systems**, - \* -

Tarun Bansal and Neeraj Mittal

### **Algorithmic Mechanisms for Internet-based Master-Worker Computing with Untrusted and Selfish Workers**, - , %

Antonio Fernández Anta, Chryssis Georgiou and Miguel A. Mosteiro

### **Stabilizing Pipelines for Streaming Applications**, - - &

Andrew Berns, Anurag Dasgupta and Sukumar Ghosh

### **A Dynamic Approach for Characterizing Collusion in Desktop Grids**, - %\$\$\$%

Louis-Claude Canon, Emmanuel Jeannot and Jon Weissman

## **Automatic Tuning and Automatic Parallelization**



## **Offline Library Adaptation Using Automatically Generated Heuristics**

Frédéric de Mesmay, Yevgen Voronenko and Markus Püschel

## **An Auto-Tuning Framework for Parallel Multicore Stencil Computations**

Shoaib Kamil, Cy Chan, Leonid Oliker, John Shalf and Samuel Williams

## ***DynTile*: Parametric Tiled Loop Generation for Parallel Execution on Multicore Processors**

Albert Hartono, Muthu Manikandan Baskaran, J. Ramanujam and P. Sadayappan

## **Using Focused Regression for Accurate Time-Constrained Scaling of Scientific Applications**

Brad Barnes, Jeonifer Garren, David K. Lowenthal, Jaxk Reeves, Bronis R. de Supinski, Martin Schulz and Barry Rountree

## **Architectural Support for Runtime Systems**

### **A Low Cost Split-Issue Technique to Improve Performance of SMT Clustered VLIW Processors**

Manoj Gupta, Fermín Sánchez and Josep Llosa

### **Exploiting Inter-thread Temporal Locality for Chip Multithreading**

Jiayuan Meng, Jeremy W. Sheaffer and Kevin Skadron

### **Profitability-Based Power Allocation for Speculative Multithreaded Systems**

Polychronis Xekalakis, Nikolas Ioannou, Salman Khan and Marcelo Cintra

### **Evaluating Standard-Based Self-Virtualizing Devices: A Performance Study on 10 GbE NICs with SR-IOV Support**

Jiuxing Liu

## **Client-Server System Management and Analysis**

### **QoS Assessment of WS-BPEL Processes through non-Markovian Stochastic Petri Nets**

Dario Bruneo, Salvatore Distefano, Francesco Longo and Marco Scarpa

### **Power-aware Resource Provisioning in Cluster Computing**

Kaiqi Xiong

### **Using the Middle Tier to Understand Cross-Tier Delay in a Multi-tier Application**

Haichuan Wang, Qiming Teng, Xiao Zhong and Peter F. Sweeney

### **Service and Resource Discovery in Cycle-Sharing Environments with a Utility Algebra**

João Nuno Silva, Paulo Ferreira and Luís Veiga

## **Parallel Graph Algorithms II**

### **Optimization of Linked List Prefix Computations on Multithreaded GPUs Using CUDA**

Zheng Wei and Joseph JaJa

### **Parallel External Memory Graph Algorithms**

Lars Arge, Michael T. Goodrich and Nodari Sitchinava

**Engineering a Scalable High Quality Graph Partitioner** ,  
Manuel Holtgrewe, Peter Sanders and Christian Schulz

## Algorithms for Wireless Networks

**Sparse Power-Efficient Topologies for Wireless Ad Hoc Sensor Networks** \$  
Amitabha Bagchi

**Contention-based Georouting with Guaranteed Delivery, Minimal Communication Overhead, and Shorter Paths in Wireless Sensor Networks** \$  
Stefan Rührup and Ivan Stojmenović

**Midpoint Routing Algorithms for Delaunay Triangulations** -  
Weisheng Si and Albert Y. Zomaya

**A Local, Distributed Constant-Factor Approximation Algorithm for the Dynamic Facility Location Problem** \$\*  
Bastian Degener, Barbara Kempkes and Peter Pietrzyk

## Analysis of Heterogeneity and Future Platforms

**Toward Understanding Heterogeneity in Computing** %\*  
Arnold L. Rosenberg and Ron C. Chiang

**Balls into Non-uniform Bins** %&\*<br>Petra Berenbrink, André Brinkmann, Tom Friedetzky and Lars Nagel

**An Introductory Exascale Feasibility Study for FFTs and Multigrid** %' \*  
Hormozd Gahvari and William Gropp

## Data Management

**A Cost-Effective Strategy for Intermediate Data Storage in Scientific Cloud Workflow Systems** %&()  
Dong Yuan, Yun Yang, Xiao Liu and Jinjun Chen

**BlobSeer: Bringing High Throughput under Heavy Concurrency to Hadoop Map-Reduce Applications** %&) +  
Bogdan Nicolae, Diana Moise, Gabriel Antoniu, Luc Bougé and Matthieu Dorier

**PreDatA - Preparatory Data Analytics on Peta-Scale Machines** %&\* ,  
Fang Zheng, Hasan Abbasi, Ciprian Docan, Jay Lofstead, Qing Liu, Scott Klasky, Manish Parashar, Norbert Podhorszki, Karsten Schwan and Matthew Wolf

**Reconciling Scratch Space Consumption, Exposure, and Volatility to Achieve Timely Staging of Job Input Data** %&, \$  
Henry M. Monti, Ali R. Butt and Sudharshan S. Vazhkudai

## Synchronization

**Hierarchical Phasers for Scalable Synchronization and Reductions in Dynamic Parallelism** %&- &  
Jun Shirako and Vivek Sarkar

**Clustering JVMs with Software Transactional Memory Support** % \$(  
Christos Kotselidis, Mikel Luján, Mohammad Ansari, Konstantinos Malakasis, Behram Kahn, Chris Kirkham and Ian Watson

## **Inter-Block GPU Communication via Fast Barrier Synchronization**

Shucaï Xiao and Wu-chun Feng

## **A Lock-Free, Cache-Efficient Multi-Core Synchronization Mechanism for Line-Rate Network Traffic Monitoring**

Patrick P. C. Lee, Tian Bu and Girish Chandranmenon

## **Extreme Scale Computing: Modeling the Impact of System Noise in Multicore Clustered Systems**

Seetharami Seelam, Liana Fong, Asser Tantawi, John Lewars, John Divirgilio and Kevin Gildea

## **Oblivious Algorithms for Multicores and Network of Processors**

Rezaul Alam Chowdhury, Francesco Silvestri, Brandon Blakeley and Vijaya Ramachandran

## **Analyzing and Adjusting User Runtime Estimates to Improve Job Scheduling on the Blue Gene/P**

Wei Tang, Narayan Desai, Daniel Buettner and Zhiling Lan

## **Performance Evaluation of Concurrent Collections on High-Performance Multicore Computing Systems**

Aparna Chandramowliswaran, Kathleen Knobe and Richard Vuduc

## **Best Papers - Plenary**

### **Extreme Scale Computing: Modeling the Impact of System Noise in Multicore Clustered Systems**

Seetharami Seelam, Liana Fong, Asser Tantawi, John Lewars, John Divirgilio and Kevin Gildea

### **Oblivious Algorithms for Multicores and Network of Processors**

Rezaul Alam Chowdhury, Francesco Silvestri, Brandon Blakeley and Vijaya Ramachandran

### **Analyzing and Adjusting User Runtime Estimates to Improve Job Scheduling on the Blue Gene/P**

Wei Tang, Narayan Desai, Daniel Buettner and Zhiling Lan

### **Performance Evaluation of Concurrent Collections on High-Performance Multicore Computing Systems**

Aparna Chandramowliswaran, Kathleen Knobe and Richard Vuduc