log in

HIGH PERFORMANCE COMPUTATION CONFERENCE 2007
UNIVERSITY *of* HOUSTON

Home > **Conference Program**

Home

Meeting Location

Program

Registration

Organizing Committee

Program Committee

Conference Papers

Hotel Information

Parking

Contact Us

Conference proceedings

View conference proceedings

Conference: 26-28 Sep 07

It is advisable to reserve your hotel room today. All events will take place at Hotel ZaZa unless otherwise indicated.

Please read these instructions for speakers and session chairs.

The proceedings is available online now!

26 Sep | Wednesday

07.00	Registration
08.15	Welcome by John Bear, Dean of the College of Natural Sciences and Mathematics, University of Houston
08.30-9.30	Vivek Sarkar | Rice University Programming Challenges for Petascale and Multicore Parallel Systems abstract
9.30-10.00	Break

10.00-12.00	Session 1:
	
	A: Peer-to-Peer Computing (Session Chair: Rong Zheng)
	
	   Proactive Method for Content Distribution in a Data Indexed DHT Overlay
	   Bassam A. Alqaralleh, Chen Wang, Bing Bing Zhou, Albert Y. Zomaya
	
	   CDACAN: A Scalable Structured P2P Network Based on Continuous Discrete Approach and CAN
	   Lingwei Li, Qunwei Xue, Deke Guo
	
	   Multi-Domain Topology-Aware Grouping for Application-Layer Multicast
	   Jianqun Cui, Yanxiang He, Libing Wu, Hui Jin
	
	   Parallel Database Sort and Join Operations Revisited on Grids
	   Werner Mach, Erich Schikuta
	
	B: Reliability and Fault Tolerance (Session Chair: Edgar Garbriel)
	
	   Dynamic Preemptive Multi-Class Routing Scheme under Dynamic Traffic in Survivable WDM Mesh Networks
	   Xuetao Wei, Lemin Li, Hongfang Yu, Du Xu
	
	   Quantification of Cut Sequence Set for Fault Tree Analysis
	   Dong Liu, Chunyuan Zhang, Weiyan Xing, Rui Li, Haiyan Li
	
	   Improving a Fault-Tolerant Routing Algorithm Using Detailed Traffic Analysis
	   Abbas Nayebi, Arash Shamaei, Hamid Sarbazi-Azad
	
	C: Data Mining, Management and Optimization (Session Chair: Carlos Ordonez)
	
	   Hybrid Line Search for Multiobjective Optimization
	   Crina Grosan, Ajith Abraham

	A Data Imputation Model in Sensor Databases Nan Jiang An Adaptive Parallel Hierarchical Clustering Algorithm Zhaopeng Li, Kenli Li, Degui Xiao, Lei Yang
12.00-13.15	Lunch
13.15-13.50	Mitsuhisa Sato | University of Tsukuba Bandwidth-aware Design of Large-scale Clusters for Scientific Computations abstract
14.00-15:30	Session 2: **A: Distributed, Mobile and Pervasive Computing** (Session Chair: Laurence Yang) Resource Aggregation and Workflow with Webcom Oisín Curran, Paddy Downes, John Cunniffe, Andy Shearer, John P.Morrison Maximum-Objective-Trust Clustering Solution and Analysis in Mobile Ad Hoc Networks Qiang Zhang, Guangming Hu, Zhenghu Gong A New Method for Multi-Objective TDMA Scheduling in Wireless Sensor Networks Using Pareto-Based PSO and Fuzzy Comprehensive Judgement Tao Wang, Zhiming Wu, Jianlin Mao **B: Networking: Protocols, Routing, Algorithms** (Session Chair: Jaspal Subhlok) Power-Aware Fat-Tree Networks Using On/Off Links Marina Alonso, Salvador Coll, Vicente Santonja, Juan-Miguel Martínez,Pedro López, José Duato Efficient Broadcasting in Multi-Radio Multi-Channel and Multi-Hop Wireless Networks Based on Self-Pruning Li Li, Bin Qin, Chunyuan Zhang, Haiyan Li Open Box Protocol (OBP) Paulo Loureiro, Saverio Mascolo, Edmundo Monteiro **C: High Performance Scientific and Engineering Computing** (Session Chair: Kengo Nakajama) A Block JRS Algorithm for Highly Parallel Computation of SVDs Mostafa I. Soliman, Sanguthevar Rajasekaran, Reda Ammar Concurrent Number Cruncher: An Efficient Sparse Linear Solver on the GPU Luc Buatois, Guillaume Caumon, Bruno L´evy Adaptive Computation of Self Sorting In-Place FFTs on Hierarchical Memory Architectures Ayaz Ali, Lennart Johnsson, Jaspal Subhlok
15.30-16.00	Break
16.00-17.30	Session 3: **A: Grid Computing** (Session Chair: Mitsuhisa Sato) A Complex Network-Based Approach for Job Scheduling in Grid Environments Renato P. Ishii, Rodrigo F. de Mello, Laurence T. Yang Performance Prediction Based Resource Selection in Grid Environments Peggy Lindner, Edgar Gabriel, Michael M. Resch Online Algorithms for Single Machine Schedulers to Support Advance Reservations from Grid Jobs Bo Li, Dongfeng Zhao **B: Networking: Protocols, Routing, Algorithms** (Session Chair: Jane Yuan) Stability Aware Routing: Exploiting Transient Route Availability in MANETs Pramita Mitra, Christian Poellabauer, Shivajit Mohapatra

Reliable Event Detection and Congestion Avoidance in
Wireless Sensor Networks
Md. Mamun-Or-Rashid, Muhammad Mahbub Alam,
Md. Abdur Razzaque, Choong Seon Hong

Systolic Routing in an Optical Ring with Logarithmic Shortcuts
Risto T. Honkanen, Juha-Pekka Liimatainen

C: High Performance Scientific and Engineering Computing
(Session Chair: Henry Jin)

Parallel Multistage Preconditioners Based on a Hierarchical Graph
Decomposition for SMP Cluster Architectures with a Hybrid
Parallel Programming Model
Kengo Nakajima

High Performance FFT on SGI Altix 3700
Akira Nukada, Daisuke Takahashi, Reiji Suda, Akira Nishida

Highly Scalable Distributed Component Framework
for Scientific Computing
Kostadin Damevski, Ashwin Deepak Swaminathan, Steven Parker

27 Sep | Thursday

09.00-10.00	Ray Simar | Texas Instruments The Changing Impact of Semiconductor Technology on Processor Architecture abstract
10.00-10.15	Break
10.15-12.15	Session 4: **A: Grid Computing** (Session Chair: Matthias Mueller) CROWN FlowEngine: A GPEL-Based Grid Workflow Engine Jin Zeng, Zongxia Du, Chunming Hu, Jinpeng Huai Dynamic System-Wide Reconfiguration of Grid Deployments in Response to Intrusion Detections Jonathan Rowanhill, Glenn Wasson, Zach Hill, Jim Basney, Yuliyan Kiryakov, John Knight, Anh Nguyen-Tuong, Andrew Grimshaw, Marty Humphrey File and Memory Security Analysis for Grid Systems Unnati Thakore, Lorie M. Liebrock Business Model and the Policy of Mapping Light Communication Grid-Based Workflow within the SLA Context Dang Minh Quan, Jörn Altmann **B: Embedded Systems** (Session Chair: Eric Stotzer) Energy-Aware Online Algorithm to Satisfy Sampling Rates with Guaranteed Probability for Sensor Applications Meikang Qiu, Edwin H.-M. Sha A Low-Power Globally Synchronous Locally Asynchronous FFT Processor Yong Li, Zhiying Wang, Jian Ruan, Kui Dai Parallel Genetic Algorithms for DVS Scheduling of Distributed Embedded Systems Man Lin, Chen Ding Journal Remap-Based FTL for Journaling File System with Flash Memory Seung-Ho Lim, Hyun Jin Choi, Kyu Ho Park **C: Languages and Compilers** (Session Chair: Larry Meadows) Strategies and Implementation for Translating OpenMP Code for Clusters Deepak Eachempati, Lei Huang, Barbara Chapman Optimizing Array Accesses in High Productivity Languages Mackale Joyner, Zoran Budimlíc, Vivek Sarkar Software Pipelining for Packet Filters Yoshiyuki Yamashita, Masato Tsuru Speculative Parallelization – Eliminating the Overhead of Failure Mikel Luján, Phyllis Gustafson, Michael Paleczny, Christopher A. Vick

12.15-13.45	Lunch
13.45-14.30	Larry Meadows | Intel OpenMP 3.0: A preview of the upcoming standard abstract
14.30-15.00	Break
15.00-17.00	**Vendor Session:** (Session Chair: Ronald Perrott) Ruud van der Pas | Sun Microsystems Evaluating the throughput performance of the heavily threaded UltraSPARC T2 multicore processor abstract Conrad Geiger | Sun Microsystems Current and Emerging Multicore HPC solutions abstract David G. Solt | Hewlett-Packard Co. A profile-based approach for topology aware MPI rank placement abstract Slides Alan Gatherer | Texas Instruments Designing Highly Integrated Basestation Transceiver SoCs abstract

28 Sep | Friday

09.00-10.00	Rob Schreiber | Hewlett Packard Labs Manycores in the Future abstract
10.00-10.30	Break
10.30-12.30	Session 5: **A: Grid Computing** (Session Chair: Angela Sodan) The One-Click Grid-Resource Model Martin Rehr, Brian Vinter Optimizing Performance of Automatic Training Phase for Application Performance Prediction in the Grid Farrukh Nadeem, Radu Prodan, Thomas Fahringer Latency in Grid over Optical Burst Switching with Heterogeneous Traffic Yuhua Chen, Wenjing Tang, Pramode K. Verma **B: Parallel/Distributed Architectures** (Session Chair: Lei Huang) MC2DR: Multi-Cycle Deadlock Detection and Recovery Algorithm for Distributed Systems Md. Abdur Razzaque, Md. Mamun-Or-Rashid, Choong Seon Hong FROCM: A Fair and Low-Overhead Method in SMT Processor Shuming Chen, Pengyong Ma A Highly Efficient Parallel Algorithm for H.264 Encoder Based on Macro-Block Region Partition Shuwei Sun, Dong Wang, Shuming Chen Towards Scalable and High Performance I/O Virtualization – A Case Study Jinpeng Wei, Jeffrey R. Jackson, John A. Wiegert **C: Performance, Evaluation and Measurements Tools** (Session Chair: Oscar Hernandez) Parallel Performance Prediction for Multigrid Codes on Distributed Memory Architectures Giuseppe Romanazzi, Peter K. Jimack Netgauge: A Network Performance Measurement Framework

Torsten Hoefler, Torsten Mehlan, Andrew Lumsdaine,
Wolfgang Rehm

Toward a Complexity Model for Design and Analysis of
PGAS-Based Algorithms
Mohamed Bakhouya, Jaafar Gaber, Tarek El-Ghazawi

An Exploration of Performance Attributes for Symbolic Modeling of
Emerging Processing Devices
Sadaf R. Alam, Nikhil Bhatia, Jeffrey S. Vetter

12.30-14.00	Lunch
14.00-14.45	Haoqiang Jin	NASA Ames Research Center Towards Enhancing OpenMP Expressiveness and Performance abstract
14.45-15.00	Break

15.00-17.00	Session 6: **A: Web Services and Internet Computing** (Session Chair: Babu Sundaram) An Ontology for Semantic Web Services Qizhi Qiu, Qianxing Xiong DISH – Dynamic Information-Based Scalable Hashing on a Cluster of Web Cache Servers Andrew Sohn, Hukeun Kwak, Kyusik Chung CIVIC: A Hypervisor Based Virtual Computing Environment Jinpeng Huai, Qin Li, Chunming Hu On Pancyclicity Properties of OTIS Networks Mohamadreza R. Hoseinyfarahabady, Hamid Sarbazi-Azad **B: Cluster Computing** (Session Chair: Edgar Gabriel) PARMI: A Publish/Subscribe Based Asynchronous RMI Framework for Cluster Computing Heejin Son, Xiaolin Li Coarse-Grain Time Slicing with Resource-Share Control in Parallel-Job Scheduling Bryan Esbaugh, Angela C. Sodan Quality Assurance for Clusters: Acceptance-, Stress-, and Burn-In Tests for General Purpose Clusters Matthias S. Müller, Guido Juckeland, Matthias Jurenz, Michael Kluge Performance Evaluation of Distributed Computing over Heterogeneous Networks Ouissem Ben Fredj, Éric Renault A Windows-Based Parallel File System Lungpin Yeh, Juei-Ting Sun, Sheng-Kai Hung, Yarsun Hsu **C: Performance, Evaluation and Measurements Tools and Experiences** (Session Chair: Torsten Hoefler) Towards Scalable Event Tracing for High End Systems Kathryn Mohror, Karen L. Karavanic Checkpointing Aided Parallel Execution Model and Analysis Laura Mereuta, Éric Renault Throttling I/O Streams to Accelerate File-IO Performance Seetharami Seelam, Andre Kerstens, Patricia J. Teller Performance Evaluation of View-Oriented Parallel Programming on Cluster of Computers Haifeng Shang, Jiaqi Zhang, Wenguang Chen, Weimin Zheng, Zhiyi Huang Security Enhancement and Performance Evaluation of an Object-Based Storage System Po-Chun Liu, Sheng-Kai Hong, Yarsun Hsu

Tutorials: 29 Sep 07

29 Sep | Saturday

OpenMP: 8:00 am to 3:00 pm, Ballroom

Alternatively entitled "How do I take advantage of all those cores?", this tutorial targets developers who are new to parallel computing and want to learn more about parallelizing applications using the OpenMP shared memory parallel programming model. OpenMP was originally designed for SMP systems, but is equally suitable for multicore systems. Taught by Ruud van der Pas from Sun Microsystems. More info

Schedule
08:00 - 08:45 Concepts in parallelization
08:45 - 09:15 Multicore Processor Architectures
09:15 - 10:00 OpenMP Overview, part 1
10:00 - 10:30 Break
10:30 - 11:15 OpenMP Overview, part 2
11:15 - 12:00 Data Races
12:00 - 13:00 Lunch
13:00 - 13:30 OpenMP under the hood (Lei Huang, UH)
13:30 - 14:00 Cluster OpenMP (Larry Meadows, Intel)
14:00 - 14:30 OpenMP and Performance
14:30 - 15:00 Q&A

Threading Building Blocks: 8:00 am to 12:00 noon, Ballroom

This tutorial is an introduction to Intel Threading Building Blocks (Intel TBB), a commercially supported C++ template library for shared-memory parallel programming, notably for multi-core processors. Though threads are a popular means of shared-memory parallel programming, they are a low-level unstructured construct whose undisciplined use can cause both correctness and performance problems. Taught by Arch D. Robison from Intel. More info

Quick Links: Online Registration | Conference venue | Book a hotel room

State of Texas | Compact with Texans | Statewide Search | Homeland Security | UH System | University of Houston | Search | Hot Page | Links | Site Map