



International Parallel & Distributed Processing Symposium
April 23-27, 2001 – Hyatt Regency San Francisco Airport
Sponsored by IEEE Computer Society
Technical Committee on Parallel Processing
www.ipdps.org

ADVANCE PROGRAM

International Parallel & Distributed Processing Symposium
April 23-27, 2001
Hyatt Regency - San Francisco Airport
San Francisco, California - USA

This downloadable version of the IPDPS 2001 program was prepared in advance of final submission of papers for publication in the proceedings. Therefore, paper titles and author listings may have been changed from how they appear in this document. . It is also possible that accents or special characters were not accurately translated. Authors of contributed papers listed here are welcome to contact info@ipdps.org regarding changes or errors and every effort will be made to correct this document.

IPDPS 2001 WORKSHOPS All Day Monday

1 Heterogeneous Computing Workshop

2 Workshop on Parallel and Distributed Real-Time Systems

3 Workshop on High-Level Parallel Programming Models & Supportive Environments

4 Workshop on Java for Parallel and Distributed Computing

5 Workshop on Parallel and Distributed Computing in Image Processing, Video Processing, and Multimedia

6 Workshop on Solving Irregularly Structured Problems in Parallel

7 Workshop on Fault-Tolerant Parallel and Distributed Systems

8 Workshop on Biologically Inspired Solutions to Parallel Processing Problems

9 Workshop on Advances in Parallel and Distributed Computational Models

Note: Workshops are open to all IPDPS 2001 registrants. Detailed program schedules will be available on-site. See individual advance program schedules for each workshop on their web sites which may be accessed by link from www.ipdps.org.

MONDAY TUTORIALS TUTORIAL 1

9 AM – 4 PM

High Performance Computing in Java: Compiler, Language, and Application Solutions

WHO SHOULD ATTEND

This tutorial is of interest to those developing computationally intensive applications in Java, with particular emphasis on engineering and scientific applications. The tutorial is also relevant to those developing high performance Java environments (compilers and/or virtual machines) and high performance Java libraries.

COURSE DESCRIPTION

As high performance numerical computing embraces the object-oriented paradigm, the use of Java for such applications is gaining interest. Although Java has many attractive features – reliability, portability, well defined floating point semantics and a growing programmer base – the performance of current commercial implementations in numerical applications is still an impediment to a wider adoption of Java. In this tutorial we describe how (1) standard libraries and currently proposed Java extensions can help in both achieving high performance and writing more maintainable code; and (2) compilers can be improved to provide near-Fortran performance. The proposals of the Java Grande Forum Numerics Working Group, which include a true multidimensional array package, complex arithmetic, and new floating point semantics, will be discussed. Compiler technologies to be addressed include array bounds and null pointer check optimizations, alias analysis techniques, semantic expansion of standard classes, and the interplay of static and dynamic models of computation. Throughout the tutorial the impact of the technology under discussion on compiler writers, language designers, and application developers will be described.

LECTURERS

Manish Gupta is a Research Staff Member and Manager, High Performance Programming Environments, at the IBM Thomas J. Watson Research Center. He received a B.Tech. degree in Computer Science from the Indian Institute of Technology, Delhi, in 1987, an M.S. from Ohio State University in 1988, and a Ph.D. in Computer Science from the University of Illinois in 1992.

Samuel P. Midkiff received a B.S. degree in computer science in 1983 from the University of Kentucky, and M.S. and Ph.D. degrees in computer science from the University of Illinois at Urbana-Champaign in 1986 and 1992, respectively. Dr. Midkiff

is a Research Staff Member in the High Performance Programming Environments Department at the IBM Thomas J. Watson Research Center and an Adjunct Assistant Professor at the University of Illinois, Urbana-Champaign.

Jose E. Moreira received B.S. degrees in physics and electrical engineering in 1987 and an M.S. degree in electrical engineering in 1990, all from the University of Sao Paulo, Brazil. He received his Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign in 1995. Dr. Moreira is a Research Staff Member in the High Performance Programming Environments Department at the IBM Thomas J. Watson Research Center.

TUTORIAL 2

8 AM – 12 Noon

Introduction to Effective Parallel Computing

WHO SHOULD ATTEND

New parallel computing users, as well as students, potential users, managers, and anyone who wants an overview of parallel programming.

COURSE DESCRIPTION

This tutorial provides a comprehensive overview of parallel computing, emphasizing those aspects most relevant to the user. It discusses both hardware and software, with an emphasis on standards, portability, and systems that are now (or soon will be) commercially or freely available. Systems examined range from low-cost clusters to highly integrated supercomputers. The tutorial surveys basic concepts and terminology, and gives parallelization examples selected from engineering, scientific, and data intensive applications. These real-world examples are targeted at distributed memory systems, using MPI, and shared memory systems using OpenMP. The tutorial shows basic parallelization approaches, step-by-step performance improvement, and discusses some of the software engineering aspects of the parallelization process. The tutorial also provides pointers to the literature and web-based resources.

LECTURERS

Quentin F. Stout is Professor of Electrical Engineering and Computer Science, and director of the Center for Parallel Computing, at the University of Michigan. In over 15 years of research in parallel computing, he has published over 100 articles, book chapters, and books on the topic and has utilized a range of systems for a wide spectrum of applications. He has won Best Presentation and Best Paper

awards, and two of his doctoral students have won Best Thesis awards for their work in parallel computing. He and his students have had research and consulting funded by DARPA, NSF, NASA, NRL, Digital, IBM, ATT, Unisys, Xerox, Ford Motor, and Dow Research.

Christiane Jablonowski is a meteorologist currently pursuing her Ph.D. in Atmospheric Science and Scientific Computing at the University of Michigan. She has over 10 years of experience in large-scale applications on parallel and parallel vector machines. She has worked at Siemens/Fujitsu; the "Vector and Parallel Computing" section at Aachen University of Technology, Germany; and the European Center for Medium-Range Weather Forecasts in Reading, England. Her research has been focused on climate and weather prediction modeling, with strong emphasis on parallel computing aspects. At the German National Research Center for Information Technology she won the Best Thesis award.

TUTORIAL 3 1 PM – 5 PM

Parallel and Distributed Data Mining

WHO SHOULD ATTEND

This tutorial is meant for researchers, professionals and advanced students who are interested in learning how parallelism is exploited in mining very large volumes of data. People who are dealing with algorithms and applications for the efficient extraction of knowledge using parallel and distributed computing will be interested in this topic.

COURSE DESCRIPTION

Data mining is the automated analysis of large volumes of data looking for relationships and knowledge that are implicit in data and are 'interesting' in the sense of impacting an organization's practice. Data mining and knowledge discovery on large amounts of data can benefit from the use of parallel computers to improve both performance and quality of data selection. The goal of this tutorial is to provide researchers and practitioners with an introduction to mining large data sets by exploiting techniques from parallel and distributed computing and high performance data management.

This tutorial analyzes different forms of parallelism that can be exploited in data mining techniques and algorithms. For each data mining technique, different ways for parallel implementation are presented and discussed. Furthermore, parallel and distributed data mining systems and algorithms are discussed and compared. Finally, current research issues and

perspectives in high-performance data mining are outlined.

LECTURERS

DOMENICO TALIA is a senior researcher at the ISI-CNR – Institute of Systems Analysis and Information Technology of the Italian National Research Council – and a lecturer at the University of Calabria. Talia's main research interests are on parallel computation, parallel programming languages, parallel data mining and cellular automata. Talia is a member of the Editorial Board of the IEEE Computer Society Press, a member of the Editorial Board of the Parallel and Distributed Practices journal, a member of the Advisory Board of Euro-Par, and a Distinguished Speaker in the IEEE Computer Society Tutorial Program. He published three books and more than 90 papers in international journals and conference proceedings. He is member of the ACM and the IEEE Computer Society.

MOHAMMED J. ZAKI is an Assistant Professor of Computer Science at Rensselaer Polytechnic Institute. His research interests include the design of efficient, scalable, and parallel algorithms and systems for various data mining tasks. He has published over 45 papers in this area, and he recently co-edited the book, "Large-scale Parallel Data Mining," Springer-Verlag, 2000. He was co-chair for ACM SIGKDD workshop on Large-scale Parallel KDD Systems (1999), and is a co-chair for IEEE IPDPS International Workshop on High Parallel and Distributed Data Mining (2001,2000). He has given a similar tutorial at the ACM SIGKDD Conference, 2000. He is on many program committees, including the International Conference on Machine Learning (2000), ACM SIGMOD Workshop on Data Mining and Knowledge Discovery (2000), ACM SIGKDD Workshop on Distributed and Parallel Knowledge Discovery (2000), etc. He is a member of ACM and IEEE.

**8:30 AM – 9:30 AM
KEYNOTE SPEAKER**

**Ambuj Goyal, IBM
Next Generation Web:
A goldmine for
distributed computing**

(Break 9:30 – 10:00)

**10:00 AM – 12:00 NOON
SESSION 1
Architectures
Chair: Marc Snir**

*A Dynamic Periodicity Detector:
Application to Speedup
Computation*
Felix Freitag, Julita Corbalan,
Jesus Labarta, Universitat
Politecnica de Catalunya.

*Boosting SMT Performance by
Speculation Control*
Kun Luo, Manoj Franklin,
University of Maryland; Shubhendu
S. Mukherjee, Compaq Computer
Corporation; Andre Sez nec,
IRISA/INRIA.

*Influence of Array Allocation
Mechanisms on Memory System
Energy*
R. Athavale, N. Vijaykrishnan, M.
Kandemir and M.J. Irwin, Pond
Laboratory, Pennsylvania State
University.

A Microserver View of HTMT
Lilia V. Yerosheva, Shannon K.
Kuntz, Peter K. Kogge, Jay B.
Brockman, CSE Department,
University of Notre Dame.

*A PIM-based Multiprocessor
System*
Jinwo Suh, Stephen P. Crago,
Changping Li, and Robert Parker,
University of Southern California /
Information Sciences Institute,
Arlington, VA.

**10:00 AM – 12:00 NOON
SESSION 2
Computational Science &
Numerical Algorithms
Chair: Pavlos Spirakis**

*High Performance Computing in
Coastal and Hydraulic Applications*
Shahrouz Aliabadi, Clark Atlanta
University; Andrew Johnson,
Network Computing Services;
Charlie Berger, Jane Smith,
ERDC-CHL; Bruce Zellars, Adetola
Abata, Clark Atlanta University.

*Large Scale Parallel Distributed
Simulations and Visualizations of
the Olami-Feder-Christiensen
Earthquake Model*
Gonzalo Hernandez, School of
Engineering, Andres Bello
University.

*Benchmark of Parallelization
Methods for Unstructured Shock
Capturing Code*
Tsutomu Saito, Atsushi Abe,
Kazuyoshi Takayama, Tohoku
University, Japan.

*Parallel Simulation of Radio-Base
Antennas on Massively Parallel
Systems*
L. Catarinucci, P. Palazzari, ENEA
- HPCN Project - Via Anguillarese,
Rome, Italy; L. Tarricone,
University of Perugia, Perugia,
Italy.

*Fast and Scalable Parallel
Algorithms for Matrix Chain
Product and Matrix Powers on
Distributed Memory Systems*
Keqin Li, State University of New
York.

*Mixed Parallel Implementations of
Strassen and Winograd Matrix
Multiplication Algorithms*
Frédéric Desprez and Frédéric
Suter, LIP/Ecole Normale
Supérieure de Lyon.

**10:00 AM – 12:00 NOON
SESSION 3
Visualization, Graphics &
Multimedia
Chair: Burkhard Monien**

*A Rotate-Tiling Image Composition
Method for Parallel Volume
Rendering on Distributed Memory
Multicomputers*
Chin-Feng Lin, Don-Lin Yang,
Yeh-Ching Chung, Feng Chia
University, Taiwan.

*A Parallel Real Time
Implementation of Stereo Matching*
Hong Jeong, Yuns Oh, Dept. of
E.E., POSTECH, Pohang,
Republic of Korea.

*DDDDRRaW: A Prototype Toolkit
for Distributed Real-Time
Rendering on Commodity Clusters*
Thu D. Nguyen, Christopher
Peery, Department of Computer
Science, Rutgers University; John
Zahorjan, Department of Computer
Science & Engineering, University
of Washington.

*The Layering Scaler of MPEG
Video to Transmit over Non-
Guaranteed QoS Network*
Woo Jong Yoo, Sung In Lee,
Taejon Health Sciences College,
Korea; Ho Shin Son, Tae Gun
Kang, Kwan Jong Yoo, Chungnam
National University, Korea; Doo
Hyun Kim, ETRI, Korea.

*Directory Based Composite
Routing and Scheduling for
Dynamic Multimedia Environments*
Zhenghua Fu, Nalini
Venkatasubramanian, UC Irvine.

(Lunch 12:00 - 1:30)

1:30 PM – 3:30 PM

SESSION 4

Mobile Computing, WANs, and LANs

Chair: Sajal Das

Checkpointing and Rollback of Wide-area Distributed Applications using Mobile Agents

J. Cao, G.H. Chan, T. Dillon, Internet Computing and E-Commerce Lab, Department of Computing, Hong Kong

Polytechnic University; W. Jia, Department of Computer Science, City University of Hong Kong.

Competitive Analysis of On-line Randomized Call Control in Cellular Networks

Ioannis Caragiannis, Christos Kaklamanis, Evi Papaioannou, Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, Greece.

A Queueing Model for Space-Division Packets Switches and Its Application to the Performance Evaluation of Computer Networks

V. Shurbanov, D. Avresky, Boston University; P. Mehra, Compaq Tandem Labs, CA.

Robust Routing in Wide-Area WDM Networks

Weifa Liang, Dept. of Computer Science, Australian National University, Canberra, Australia.

Flow Generation for IP/ATM Label-Switched Routing over Random Networks

Aaron Harwood and Hong Shen, Griffith University, Australia.

Bandwidth Tracking in Distributed Heterogeneous Networking Environments

Craig Sullivan and Michael Jurczyk, Department of Computer Engineering and Computer Science, University of Missouri.

1:30 PM – 3:30 PM

SESSION 5

Financial Applications & Databases

Chair: Ananth Grama

Virtual Card Payment Protocol and Risk Analysis Using Performance Scoring

Sung Samyuan, Xu Xianhua, Ge Ling, Tan Chewlim, National University of Singapore.

Multithreaded Algorithms for Pricing a Class of Complex Options

Ruppa K. Thulasiram, Department of Computer Science, University of Manitoba; Lubomir Litov, Stern School of Business, New York University; Hassan Nojumi, Dept. of Mathematics, University of Delaware; Christopher T. Downing, Federal Reserve Board, Washington DC; Guang R. Gao, Dept. of Elec. & Comp. Engg., University of Delaware.

Permutation-Based Range-Join Algorithms on N-Dimensional Meshes

Shao Dong Chen, Hong Shen, Rodney Topor, Griffith University, Australia.

Performance Analysis of a Distributed Question/Answering System

Mihai Surdeanu, Dan I. Moldovan, and Sanda M. Harabagiu, Language Computer Corporation and Department of Computer Science and Engineering, Southern Methodist University.

Optimizations Enabled by Relational Data Model View to Querying Data Streams

Beth Plale and Karsten Schwan, Georgia Institute of Technology.

Experiments in Parallel Execution of Answer Set Programs

Enrico Pontelli, New Mexico State University.

1:30 PM – 3:30 PM

SESSION 6

System Software & Tools

Chair: Nalini

Venkatasubramanian

MPX: Software for Multiplexing Hardware Performance Counters in Multithreaded Programs

John M. May, Lawrence Livermore National Laboratory.

Efficient Comparison-Based Fault Diagnosis of Multiprocessor Systems Using Genetic Algorithms

Mourad Elhadef and Bechir Ayeb, Department of Mathematics and Computer Science, University of Sherbrooke, Canada.

A Perturbation-Free Replay Platform for Cross-Optimized Multithreaded Applications

Bowen Alpern, Jong-Deok Choi, Ton Ngo, IBM T.J. Watson Research Center; Manu Sridharan, MIT; John Vlissides, IBM T.J. Watson Research Center.

Asynchronous Resource Management

Suvas Vajracharya, Los Alamos National Laboratory; Daniel G. Chavarria-Miranda, Rice University.

On-Line Debugging and Performance Monitoring with Barriers

Ernesto Novillo, Paul Lu, Dept. of Computing Science, University of Alberta, Canada.

VIBe: A Micro-benchmark Suite for Evaluating Virtual Interface Architecture (VIA) Implementations

M. Banikazemi, J. Liu, S. Kutlug, A. Ramakrishnan, P. Sadayappan, Ohio State University; H. Shah, Intel Corporation; D.K. Panda, Ohio State University.

(Break 3:30 – 4:00)

4:00 PM – 6:00 PM

SESSION 7

Compilers

Chair: Subhash Saini

Data Locality Exploitation in Algorithms including Sparse Communications

Gerardo Bandera, Emilio L. Zapata, Computer Architecture Dept, Univ. Malaga.

Minimum Register Instruction Sequence Problem: Revisiting Optimal Code Generation for DAGs

R. Govindarajan, H. Yang, C. Zhang, J.N. Amaral, G.R. Gao, University of Delaware.

A Comparison of Parallelization Techniques for Irregular Reductions

Hwansoo Han, Chau-Wen Tseng, Department of Computer Science, University of Maryland.

4:00 PM – 6:00 PM

SESSION 8

Real-time Systems & Fault Tolerance

Chair: Karsten Schwan

A Predictive Algorithm for Adaptive Resource Management of Periodic Tasks in Asynchronous Real-Time Distributed Systems

Binoy Ravindran and Tamir Hegazy, The Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University Blacksburg, VA.

An Automatic Scheduler for Real-Time Vision Applications

Mau-Tsuen Yang, Rangachar Kasturi, Anand Sivasubramaniam, Pennsylvania State University.

A Context Switch Technique for Real-time Task Synchronization

Fan Jiang and Albert M.K. Cheng, Real-Time Systems Laboratory, Department of Computer Science, University of Houston and Rice University Houston, TX.

Comparing Fail-Sailence Provided by Process Duplication versus Internal Error Detection for DHCP Server

David T. Stott, Zbigniew Kalbarczyk, Josh Scheid, Jun Xu, Ravishankar K. Iyer, Center for Reliable and High-Performance Computing, University of Illinois at Urbana-Champaign, IL; Neil A. Speirs, Department of Computing Science, University of Newcastle upon Tyne, United Kingdom.

A New Fault-Tolerant Technique for Improving the Schedulability in Multiprocessor Real-time Systems

R. Al-Omari, Arun K. Somani, G. Manimaran, Iowa State University.

RACE: A Software-Based Fault Tolerance Scheme for Systematically Transforming Ordinary Algorithms to Robust Algorithms

Chi-Hsiang Yeh, Dept. of Electrical & Computer Engineering, Queen's University, Canada; Behrooz Parhami and Emmanouel A. Varvarigos, Dept. of Electrical & Computer Engineering, University of California, Santa Barbara; Theodora A. Varvarigou, Dept. of Electrical and Computer Engineering, National Technical University of Athens, Greece.

4:00 PM – 6:00 PM

SESSION 9

Graph & Non-numerical Algorithms

Chair: David Bader

Parallel 2D Delaunay Triangulations in HPF and MPI

Min-Bin Chen, Tyng-Ruey Chuang, and Jan-Jan Wu, Institute of Information Science, Academia Sinica Nankang, Taipei, Taiwan.

Efficient Graph Algorithms on a Linear Array with a Reconfigurable Pipelined Bus System

Amitava Datta, Department of Computer Science and Software Engineering, University of Western Australia.

Parallel Algorithms for Hamiltonian 2-Separator Chordal Graphs

B.S. Panda, Vijay Natarajan, Birla Institute of Technology and Science, Pilani, India; Sajal K. Das, The University of Texas at Arlington.

Optimal Tree Access by Elementary and Composite Templates in Parallel Memory Systems

Vincenzo Auletta, Universita di Salerno, Italy; Sajal K. Das, The University of Texas at Arlington; Amelia De Vivo, Universita di Salerno, Italy; M. Cristina Pinotti, Universita' degli Studi di Trento, Italy; Vittorio Scarano, Universita di Salerno, Italy.

Linear-time Matrix Transpose Algorithms Using Vector Register File With Diagonal Registers

Bedros Hanounik, Xiaobo (Sharon) Hu, Department of Computer Science and Engineering, University of Notre Dame.

**8:30 AM – 9:30 AM
KEYNOTE SPEAKER**

Tom Leighton, MIT
***The challenges of delivering
content on the internet***

(Break 9:30 – 10:00)

10:00 AM – 12:00 NOON



BEST PAPERS SESSION

Chair: Vipin Kumar

*Transparent Adaptation of Sharing
Granularity in MultiView-Based
DSM Systems*

Nitzann Niv and Assaf Schuster,
Technion -- Israel Institute of
Technology.

*Improving Effective Bandwidth
through Compiler Enhancement of
Global Cache Reuse*

Chen Ding, Department of
Computer Science, University of
Rochester, NY; Ken Kennedy,
Department of Computer Science,
Rice University, Houston, TX.

*Minimizing Completion Time for
Loop Tiling with Computation and
Communication Overlapping*

Georgios Goumas, Aristidis
Sotiropoulos, and Nectarios
Koziris, National Technical
University of Athens, Dept. of
Electrical and Computer
Engineering, Computing Systems
Laboratory, Zografou, Greece.

*CPR: Mixed Task and Data
Parallel Scheduling for Distributed
Systems*

Andrei Radulescu, Cristina
Nicolescu, Arjan J.C. van Gemund,
Pieter P. Jonker, Delft University of
Technology.

(Lunch 12:00 - 1:30)

1:30 PM – 3:30 PM

SESSION 10
***Performance Modeling &
Latency Management***

Chair: Amitava Datta

*HiHCoHP---Toward a Realistic
Communication Model for
Hierarchical HyperClusters of
Heterogeneous Processors*

Franck Cappello, Pierre
Fraigniaud, Lab. de Recherche en
Informatique, Univ. of Paris-Sud;
Bernard Mans, Dept. of
Computing, Univ. Macquarie;
Arnold L. Rosenberg, Dept. of
Computer Science, Univ.
Massachusetts.

*Effect of Speculative Prefetching
on Network Load in Distributed
Systems*

N.J. Tuah, Universiti Brunei
Darussalam; M. Kumar, University
of Texas at Arlington; S.
Venkatesh, Curtin University of
Technology, Australia.

*Predicting Scalability of Parallel
Garbage Collectors on Shared
Memory Multiprocessors*

Toshio Endo, Kenjiro Taura,
Akinori Yonezawa, The University
of Tokyo.

*On the Exploitation of Value
Predication and Producer
Identification to Reduce Barrier
Synchronization Time*

Khaled Z. Ibrahim, Gregory T.
Byrd, North Carolina State
University.

*Stochastic Prediction of Execution
Time for Dynamic Bulk*

Synchronous Computations
Cheng-Zhong Xu, Le Yi Wang, and
Ngo-Tai Fong, Wayne State
University.

*Performance Analysis of a CC-
NUMA Operating System*

Moon-Seok Chang and Hae-Jin
Kim, ETRI, Korea.

1:30 PM – 3:30 PM

SESSION 11
Scheduling I

Chair: Denis Trystram

*On the Design of Clustering-based
Scheduling Algorithms for Realistic
Machine Models*

Cristina Boeres and Vinod E.F.
Rebello, Instituto de Computacao,
Universidade Federal Fluminense,
Niteroi, Brazil.

*G-commerce: Market Formulations
Controlling Resource Allocation on
the Computational Grid*

Rich Wolski, James Plank,
Computer Science Department,
University of Tennessee; John
Brevik, Mathematics and
Computer Science Department,
College of the Holy Cross; Todd
Bryan, Computer Science
Department, University of
Tennessee.

*Simulation Based HPC Workload
Analysis*

David B. Jackson, Heather L.
Jackson, Quinn O. Snell, Brigham
Young University, UT.

*Production Job Scheduling for
Parallel Shared Memory Systems*

Su-Hui Chiang, Mary K. Vernon,
Computer Sciences Department,
University of Wisconsin - Madison.

*CMC: A Coscheduling Model for
non-Dedicated Cluster Computing*

Francesc Solsona, Francesc Gine,
Departamento de Informatica e
Ingenieria Industrial, Universitat de
Lleida, Spain; Porfidio Hernandez,
Emilio Luque, Departamento de
Informatica, Universitat Autonoma
de Barcelona, Spain.

1:30 PM – 3:30 PM

SESSION 12

System Software

Chair: Josep Torrellas

Asynchronous MPI messaging on Myrinet

Chamath Keppitiyagama, Alan Wagner, Department of Computer Science, University of British Columbia, Canada.

JECho - Supporting Distributed High Performance Applications with Java Event Channels

Dong Zhou, Karsten Schwan, Greg Eisenhauer, Yuan Chen, College of Computing, Georgia Institute of Technology.

MPICH/Madeleine: a True Multi-Protocol MPI for High Performance Networks

Olivier Aumage, Guillaume Mercier, and Raymond Namyst, LIP, Ecole Normale Supérieure de Lyon.

Data Collection and Restoration for Heterogeneous Process Migration
Kasidit Chanchio and Xian-He Sun, Illinois Institute of Technology.

Fast NIC-Based Barrier over Myrinet/GM

Darius Buntinas, Dhableswar K. Panda, P. Sadayappan, Network-Based Computing Laboratory, The Ohio State University.

Protocol and Performance Analysis of the MPC Parallel Computer

J.L. Desbarbieux, O. Gluck, A. Zerrouki, A. Fenyó, A. Greiner, F. Wajsburt, C. Spasevski, F. Silva, E. Dreyfus, University P. & M. Curie, France.

(Break 3:30 – 4:00)

4:00 PM – 6:00 PM

PANEL DISCUSSION

Microprocessor design beyond the PC era: Is there room for innovation?

MODERATOR

Marc Snir, IBM Research

PARTICIPANTS

Norman P Jouppi, Compaq WRL
Peter M Kogge, University of Notre Dame

David A. Patterson, UC Berkeley
Larry Rudolph, MIT

Josep Torrellas, University of Illinois at Urbana-Champaign

INTRODUCTION TO DISCUSSION

A pessimistic view of microprocessor architecture is that, as the design cost of microprocessors keeps climbing, the number of players in this domain keeps shrinking, and players become more risk adverse. Some practitioners have bemoaned the end of computer architecture as an interesting research field, as industry becomes increasingly unlikely to accept revolutionary changes in microprocessor design.

An alternative view is that pessimists ignore both the pressure of new applications (mobile, games, communication ...), and the pressures of technological shifts (embedded DRAM, billion transistors on a chip, higher error rates ...). An optimist view is that we are at the threshold of a renaissance where we shall see a flourishing of many different architectures and drastically changed implementation techniques.

The panel will discuss these contrasting views of the state of the art in Computer Architecture.

6:30 PM – 9:00 PM

BANQUET

INVITED SPEAKER

David A. Patterson, UC Berkeley
A Server for the Post PC Era: ISTORE-1



Note

Tickets for the banquet are included in non-student registration. Students and guests may purchase tickets as part of Advance Registration but on-site availability is not guaranteed.

A brief reception will precede seating for the banquet, & the buffet menu will accommodate vegetarian attendees.

**8:30 AM – 9:30 AM
KEYNOTE SPEAKER**

Jack Dongarra
University of Tennessee
High Performance Computing
and Trends: Connecting
Computational Requirements
with Computing Resources

(Break 9:30 – 10:00)

**10:00 AM – 12:00 NOON
SESSION 13**

Broadcasts & Wormhole
Routing

Chair: Dhableswar Panda

Log-Time Multicast to Local Vertices in the Star Graph
Satoshi Fujita, Hiroshima University.

Near-Optimal All-to-All Broadcast in Multidimensional All-Port Meshes and Tori
Yuanyuan Yang, State University of New York at Stony Brook;
Jianchao Wang, GTE Laboratories, MA.

Gossiping in the Multicasting Communication Environment
Teofilo F. Gonzalez, Department of Computer Science, University of California, Santa Barbara.

A Simple Incremental Network Topology for Wormhole Switch-Based Networks
Pangfeng Liu, National Chung Cheng University; Jan-Jan Wu, Academia Sinica; Yi-Fang Lin, National Chung Cheng University; Shih-Hsien Yeh, Academia Sinica.

Performance Analysis of Wormhole-Switched k-Ary n-Cubes with Bursty Traffic
Geyong Min, Department of Computer Science, University of Strathclyde, UK; Mohamed Ould-Khaoua, Dept. of Computing Science, University of Glasgow, UK

Adaptive Fault-tolerant Wormhole Routing in 2D Meshes
Jipeng Zhou and Francis C.M. Lau, University of Hong Kong.

**10:00 AM – 12:00 NOON
SESSION 14
Scheduling II**
Chair: P. Sadayappan

User-Level Communication in a System with Gang Scheduling
Yoav Etsion and Dror G. Feitelson, The Hebrew University, Israel.

Gang Scheduling with a Queue for Large Jobs
B.B. Zhou, Deakin University; R.P. Brent, Oxford University
Computing Laboratory, UK.

Scheduling Task In-Trees on Distributed Memory Systems
Sanjeev Baskiyar, Auburn University, AL.

A Model for Moldable Supercomputer Jobs
Walfredo Cirne, Francine Berman, Computer Science and Engineering, University of California, San Diego.

Trading Execution Time for Reliability in Scheduling Precedence-Constrained Tasks in Heterogeneous Computing
Atakan Dogan, Anadolu University; Fusun Ozguner, The Ohio State University.

**10:00 AM – 12:00 NOON
SESSION 15
Software Distributed Shared Memory**
Chair: Guang Gao

A Comparison of Two Strategies of Dynamic Data Prefetching in Software DSM
Liu Haiming, Hu Weiwu, Chinese Academy of Sciences, Beijing.

Implementing Shared Memory on Clustered Machines
Carlo Fantozzi, Andrea Pietracaprina, Geppino Pucci, Dipartimento di Elettronica e Informatica, Universita' di Padova, Italy.


Efficient Categorization of Sharing Patterns in Software DSM Systems
M.C.S. de Castro, Computer Science Department, Federal University of Juiz de Fora, Brazil; C.L. de Amorim, COPPE - Systems Engineering, Federal University of Rio de Janeiro, Brazil.

Message Passing Vs. Shared Address Space on a Clusters of SMPs
Hongzhang Shan, Jaswinder Pal Singh, Princeton University; Leonid Oliker, NERSC, Lawrence Berkeley National Lab; Rupak Biswas, NASA Ames Research Center, Moffett field, CA.

Efficient Fine-Grain Sharing Support for Software DSMs Through Segmentation
Yangwoo Roh, Korea Advanced Institute of Science and Technology; Jaewoong Chung, TMax Software Inc.; Cheol Ho Park, Daeyeon Park, Korea Advanced Institute of Science and Technology.

(Lunch 12:00 - 1:30)

**1:30 PM – 3:30 PM
INDUSTRIAL TRACK
Chair: Prashanth Bhat**



Technical papers from the industrial sponsors will be presented in this plenary session

(Break 3:30 – 4:00)

4:00 PM – 6:00 PM
SESSION 16

***Routing, Switching, and
Interconnection Networks***
Chair: Christos Kaklamanis

*A Simple Formation of Orthogonal
Convex Fault Regions in Mesh-
Connected Multicomputers*
Jie Wu, Florida Atlantic University.

*Improving Network Performance
by Reducing Network Contention
in Source-Based COWs with a
Low Path-Computation Overhead*
J. Flich, P. López, M.P.
Malumbres, J. Duato Dpto.
Informática de Sistemas y
Computadores, Universidad
Politécnica de Valencia; T. Rokicki,
Instantis, Incorporated.

*A Packet-size Aware Adaptive
Routing Algorithm for Parallel
Transmission Server Systems*
Kazumasa Oida, Jun Shiozaki,
Shigeru Saito, ATR Adaptive
Communications Research
Laboratories.

*Efficient Multicast Algorithms for
Heterogeneous Switch-based
Irregular Networks of Workstations*
Amit Singhal, Mohammad
Banikazemi, P. Sadayappan, and
Dhableswar K. Panda, The Ohio
State University.

*On the VLSI Area and Bisection
Width of Star Graphs and
Hierarchical Cubic Networks*
Chi-Hsiang, Yeh Dept. of Electrical
& Computer Engineering, Queen's
University Kingston, Ontario,
Canada; Behrooz Parhami, Dept.
of Electrical & Computer
Engineering, University of
California, Santa Barbara.

4:00 PM – 6:00 PM
SESSION 17

Load Balancing
Chair: Shahrouz Aliabadi

*A High-Performance Mapping
Algorithm for Heterogeneous
Computing Systems*
Min-You Wu and Wei Shu, The
University of New Mexico.

*Improving Processor Allocation
through Run-Time Measured
Efficiency*
Julita Corbalan, Jesus Labarta,
Universitat Politecnica de
Catalunya.

*An Implementation of Parallel
Dynamic Load Balancing for
Adaptive Computing in VLSI
Device Simulation*
Yiming Li, Cheng-Kai Chen, Shui-
Sheng Lin, Jinn-Liang Liu, and
S.M. Sze, National Chiao Tung
University, Hsinchu, Taiwan.

*A Latency-Tolerant Partitioner for
Distributed Computing on the
Information Power Grid*
Sajal K. Das, Daniel J. Harvey, The
University of Texas at Arlington;
Rupak Biswas, NASA Ames
Research Center, CA.

*Dynamic Load-balancing Using
Prediction in a Parallel Object-
oriented System*
Wei Jie, Wentong Cai, and
Stephen J. Turner, Nanyang
Technological University,
Singapore.

4:00 PM – 6:00 PM
SESSION 18
Distributed Computing
Chair: Susamma Barua

*Distributed coloring and
communication in rings with local
knowledge*
Anders Dessmark, Lund
University, Sweden; Andrzej Pelc,
Universite du Quebec a Hull,
Canada.

*A Condition for k-Set Agreement in
Asynchronous Distributed Systems*
Achour Mostefaoui, Michel Raynal,
IRISA, Campus de Beaulieu,
France.

*A distributed self-stabilizing
algorithm for argumentation*
P. Baroni, M. Giacomini, Universita'
di Brescia, Dipartimento di
Elettronica per l'Automazione,
Italy.

*A Fast Algorithm for Detecting
Distributed Deadlocks in the OR
Request Model*
Soojung Lee, Dept. of Computer
Education, Incheon National
University of Education, Korea.

*Efficient Handling of Message-
Dependent Deadlock*
Yong Ho Song and Timothy Mark
Pinkston, SMART Interconnects
Group, UCLA.

IPDPS 2001 WORKSHOPS All Day Friday

10 Reconfigurable
Architectures Workshop

11 Workshop on Formal
Methods for Parallel
Programming

12 Workshop on Optical
Communications for Computing
Systems CANCELLED

13 Workshop on Parallel and
Distributed Data Mining

14 Workshop on
Communication Architecture
for Clusters

15 Workshop on Internet
Computing and E-Commerce

16 Workshop on Parallel and
Distributed Scientific and
Engineering Computing with
Applications

17 Workshop on Parallel and
Distributed Computing Issues
in Wireless Networks and
Mobile Computing

18 Workshop on Massively
Parallel Processing

19 Workshop on Scheduling
and Telecommunications

20 Workshop on Exploiting
Instruction Level Parallelism

Note: Workshops are open to all IPDPS 2001 registrants. Detailed program schedules will be available on-site. See individual advance program schedules for each workshop on their web sites which may be accessed by link from www.ipdps.org.

FRIDAY TUTORIALS

TUTORIAL 4

9 AM – 4 PM

Grid Computing, Globus, and Java Interface to the Grid

WHO SHOULD ATTEND

Those wishing to get an introduction into Grid computing and those who want to learn more about Java and Grid computing.

COURSE DESCRIPTION

Overview:

- A. General introduction to the Grid
- B. Grid programming with Globus
- C. Grid programming with Java

Grid computing technologies enable controlled resource sharing in distributed communities and the coordinated use of those shared resources as community members tackle common goals. These technologies include new protocols, services, and APIs for secure resource access, resource management, fault detection, communication, and so forth, that in turn enable new application concepts such as virtual data, smart instruments, collaborative design spaces, and metacomputations. In this tutorial, we teach review applications that are motivating widespread interest in Grid concepts within the scientific and engineering communities. Then, we describe the Globus Grid architecture that has been adopted by many Grid projects, focusing in particular on our security, resource management, and data management technologies. In the second half of the tutorial we introduce a Java library that enables to build sophisticated client applications utilizing the basic Grid services.

LECTURER

Gregor von Laszewski is an assistant research scientist at Argonne National Laboratory. His research interests are the utilization of commodity technologies in distributed and parallel computing, as well as, the development of the novel use of parallel computing in scientific applications. He has been part of the Globus project since the very early stages. He received a Ph.D. from Syracuse University.

TUTORIAL 5
8 AM – 12 Noon
SGI Pro64 Open Source
Compiler Infrastructure

WHO SHOULD ATTEND

Compiler researchers, developers, faculty, and students.

COURSE DESCRIPTION

We present an overview of the Pro64 compiler infrastructure made available as open source software by SGI in May 2000. We will address questions that must be answered by compiler researchers and students in deciding on a compiler infrastructure to be used in their work, or by others who must use compilers, especially for development of high-performance parallel software. The tutorial will provide a general overview of the Pro64 infrastructure, in particular the optimization strategy and methods. We will emphasize loop nest optimization and parallelization, but will also address traditional global optimization and code generation. We illustrate our experience in using the infrastructure for compiler research.

LECTURER(S)

The lecturers will include experts from both academia and industry (SGI). Please see <http://www.capsl.udel.edu/~ggao/Pro64.htm> for an introduction.

TUTORIAL 6
1 PM – 5 PM
Distributed Object
Computing with Java/ORB

WHO SHOULD ATTEND

The tutorial is for all those who are involved in distributed application development with an exposure to OO programming languages like C++/Java.

COURSE DESCRIPTION

With the widespread growth of Internet, the nature of the application program is changing from desktop application to the distributed application. The emerging standard like OMG's CORBA is not only facilitating the application development in such an environment but it has become a preferred platform for integrating the widely distributed objects. This has been manifested by the adoption of CORBA as a communication protocol for Sun's Java.

Common Object Request Broker Architecture (CORBA) is the standard based on Object Oriented (OO) model supporting application software development in distributed and heterogeneous environment. The main objective of CORBA is to provide

transparency at three different levels, i.e. Network, Operating System and Programming language. To reap the benefits of OO and Client Server (C/S) model, CORBA combines both technologies. One of the useful aspects of CORBA is that it applies equally across the entire market segment, such as, Finance, Tele-communication, Education, Transportation and E-Commerce.

The proposed tutorial reviews, at the outset, the conventional methodologies for distributed application development. Then it discusses the architecture of CORBA, in terms of its Components, Services and Facilities, independent of any particular implementation. Further, the CORBA application development life cycle is enumerated with Java/ORB implementation. In the end, the performance of an application developed with Java's ORB is compared with Java's RMI and other ORB product.

LECTURER

Arvind W. Kiwelekar is presently working as a Head of Computer Engineering Department of Dr. Babasaheb Ambedkar Technological University, Lonere, India. He did his M. E. (1998) from Mumbai University (India) and B. E. (19991) from Marathwada University (India). For a short period he has worked in the System Administration group of C-DAC, India's premier institute engaged in building supercomputing facilities. He is the recipient of Research Fellowship from Indian Academy of Sciences. His areas of research include Distributed Systems, Databases and Cluster Computing.