Program & Proceedings

Patrons and Commercial Participants

CRAY THE SUPERCOMPUTER COMPANY

ELSEVIER

Google

IBM

NVIDIA

24th IEEE INTERNATIONAL PARALLEL & DISTRIBUTED PROCESSING SYMPOSIUM (IPDPS) and the IPDPS 2010 WORKSHOPS & PhD FORUM (IPDPSW)

19-23 April 2010
Atlanta (Georgia) USA
www.ipdps.org

Hosted by:

Georgia Institute of Technology

IPDPS
IEEE Catalog Number: CFP10023-CDR
ISBN: 978-1-4244-6441-8
ISSN: 1530-2075

IPDPSW
IEEE Catalog Number: CFP1051J-CDR
ISBN: 978-1-4244-6532-3

Sponsored by IEEE Computer Society Technical Committee on Parallel Processing.
In cooperation with ACM SIGARCH, IEEE Computer Society Technical Committee on Computer Architecture, and IEEE Computer Society Technical Committee on Distributed Processing

IPDPS 2010 ATLANTA

Proceedings of the

24th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2010)

and

IPDPS 2010 Workshops & PhD Forum (IPDPSW)

Program Book and CD-ROM

IPDPS 2010
IEEE Catalog Number: CFP10023-CDR
ISBN: 978-1-4244-6441-8
ISSN: 1530-2075

•

IPDPSW
IEEE Catalog Number: CFP1051J-CDR
ISBN: 978-1-4244-6532-3

Program
24th IEEE International Parallel & Distributed Processing Symposium

April 19-23, 2010
ATLANTA (Georgia) USA

Sponsored by
IEEE Computer Society Technical Committee on Parallel Processing

In cooperation with
IEEE Computer Society Technical Committee on Computer Architecture
IEEE Computer Society Technical Committee on Distributed Processing
ACM SIGARCH


IEEE


IEEE
computer
society

IPDPS 2010 Program Book · Summary of Contents

This is the program book for IEEE IPDPS 2010. It includes the full schedule of events for April 19-23 in Atlanta and lists all of the papers and posters to be presented in thirty-three technical sessions and the eighteen workshops and PhD Forum. This year the conference proceedings will be published by the IEEE Digital Library under two separate ISBN's, one for regular conference papers and the other for the workshop papers and the PhD Forum short papers. The single CD-ROM of the full proceedings, pocketed in this program book, contains the two separate volumes of papers and is labeled **24th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2010) and IPDPS 2010 Workshops & PhD Forum (IPDPSW)**. Material for this program book was compiled from various sources, including workshop Web sites. Errors are possible in names & titles, so please refer to the proceedings on the CD-ROM as the final source. Apologies to authors and speakers not accurately listed. Note that workshops and tutorials are part of the IPDPS program and open to all registrants.

* The RAW workshop starts on Monday and runs through Tuesday. The program for both days is listed in the Monday section of the program.

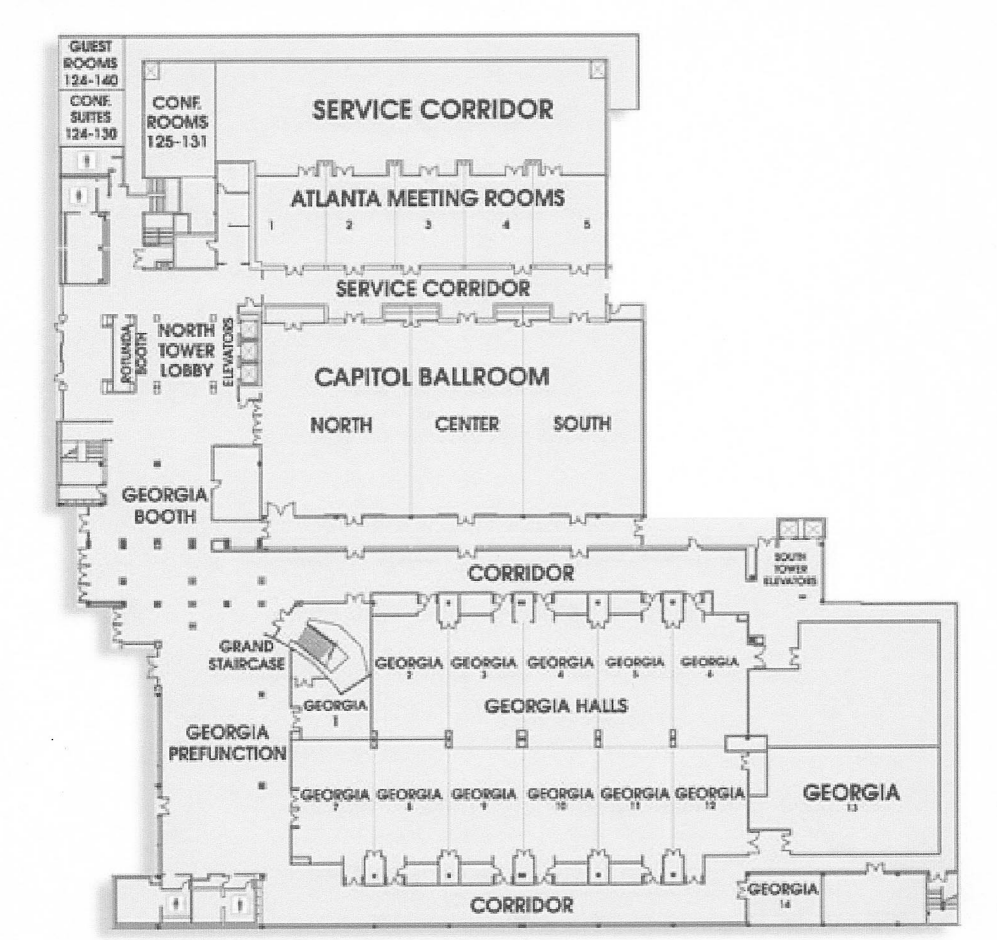

IPDPS 2010
Sheraton Atlanta Hotel
Conference Meeting Rooms

Workshop 1: HCW
Heterogeneity in Computing Workshop

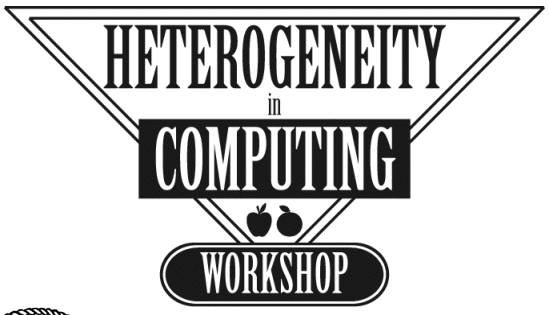

HCW 2010 is sponsored by the U.S. Office of Naval Research and by the IEEE Computer Society, through the Technical Committee on Parallel Processing (TCPP)

Description:
Today, most computing systems have elements of heterogeneity. Heterogeneity springs from the richness of environments where diversity and resource abundance prevail. Recognizing, capturing, and efficiently exploiting this diversity in an integrated and coherent manner are key goals of heterogeneous computing. With the increasing number of components in heterogeneous parallel and distributed systems, failure is becoming a critical factor that impacts application performance. In recent years, there also has been an increasing interest in robust design in parallel and distributed computing systems that must operate in an environment full of uncertainties. These uncertainties could include task execution times varying with data input sets and resources dynamically joining and leaving the system. This year, HCW specifically encouraged submissions that explore the capabilities of robust and fault-tolerant systems, paradigms, algorithms, and techniques for heterogeneous computing.

General Chair
David A. Bader, Georgia Institute of Technology, U.S.A.

Program Chair
Anne Benoit, Ecole Normale Supérieure de Lyon, France

Vice Program Chair
(for "Robust Scheduling and Fault-Tolerant Techniques")
Qin Zheng, Institute of High Performance Computing, Singapore

HCW Program

8:45-9:00: WELCOME

9:00-10:00: KEYNOTE TALK
Session Chair: H. J. Siegel

Architectural Considerations for a 500 TFLOPS Heterogeneous HPC
Richard Linderman, US Air Force Research Laboratory

10:00-10:30: BREAK

10:30-12:00: SESSION 1: Evaluation of Heterogeneous Environments
Session Chair: Anne Benoit

Characterizing Heterogeneous Computing Environments using Singular Value Decomposition
Abdulla Al-Qawasmeh and Tony Maciejewski and H. J. Siegel

Statistical Predictors of Computing Power in Heterogeneous Clusters
Ron Chi-Lung Chiang and Tony Maciejewski and Arnold L. Rosenberg and H. J. Siegel

A First Step to the Evaluation of SimGrid in the Context of a Real Application
Abdou Guermouche and Hélène Renard

12:00-1:30: LUNCH (on your own)

1:30-3:00: SESSION 2: Robust Scheduling and Fault-Tolerant Techniques
Session Chair: Ioana Banicescu

Runtime Adaptation for DAGs with Uncertainties in Heterogeneous Computing Systems
Qin Zheng

Unibus: Aspects of Heterogeneity and Fault Tolerance in Cloud Computing
Magdalena Slawinska and Jaroslaw Slawinski and Vaidy Sunderam

Robust Resource Allocation of DAGs in a Heterogeneous Multicore System
Luis D Briceño and Jay Smith and H. J. Siegel and Tony Maciejewski and Paul Maxwell and Abdulla Al-Qawasmeh and Ron Chi-Lung Chiang and Jiayin Li and Russ Wakefield

3:00-3:30: BREAK

3:30-5:30: SESSION 3: From Heterogeneous Multi-Cores to Applications
Session Chair: Qin Zheng

Decentralized Dynamic Scheduling across Heterogeneous Multi-core Desktop Grids
Jaehwan Lee and Pete Keleher and Alan Sussman

Custom Built Heterogeneous Multi-Core Architectures: Breaking the Conventions
Nagarajan Venkateswaran and Karthikeyan Saravanan and Nachiappan CN and Aravind Vasudevan and Balaji Subramaniam and Ravindhiran Mukundarajan

Improving MapReduce Performance via Data Placement in Heterogeneous Hadoop Clusters
Jiong Xie, Shu Yin, Xiaojun Ruan, Zhiyang Ding, Yun Tian, James Majors, Adam Manzanares and Xiao Qin

An Empirical Study of a Scalable Byzantine Agreement Algorithm
Olumuyiwa Oluwasanmi and Jared Saia and Valerie King

5:30: END OF WORKSHOP

Workshop 2: RAW Reconfigurable Architectures Workshop

Description:
As reconfigurable devices have grown in complexity and capabilities over the years, so has the complexity and variety of applications that must harness their features with unprecedented effectiveness. The key to effectively harnessing these features - heterogeneous fabrics, support for partial and run time reconfiguration (RTR), high speed IO, embedded microprocessors, etc. - includes research on areas such as algorithms, CAD tools, compilers, operating systems, and thermal/precision/fault management. Research challenges within these areas are expected to become even more complex with emerging trends in reconfigurable computing: 3D integration with other FPGA/memory chips, use of nano-technology devices as building blocks, operation at sub-threshold voltages for ultra-low power systems, etc. Furthermore, an appropriate mix of the theoretical foundations of reconfiguration, and practical considerations, including architectures, technologies and tools supporting soft/hard, partial/full, static/run-time reconfiguration is essential to fully reveal and exploit the possibilities created by this powerful computing paradigm. RAW 2010 aims to provide a venue to facilitate creative and productive interaction between all these disciplines.

Workshop Co Chairs
Jürgen Becker, Karlsruhe Institute of Technology - KIT, Germany
Eli Bozorgzadeh, UC Irvine, USA

Program Co Chairs
João M. P. Cardoso, University of Porto, Portugal
Aravind Dasu, Utah State University, USA

Steering Chair
Viktor K. Prasanna, University of Southern California, USA

Publicity Chair (USA)
Ramachandran Vaidyanathan, Louisiana State University, USA
Publicity Chair (Europe, Asia)
Reiner Hartenstein, Kaiserslautern University of Technology, Germany

RAW Program

(RAW will run two days, starting Monday and concluding Tuesday. The schedules for both days follow consecutively below and in the next pages.)

Monday 19th April Schedule for RAW

8:45-9:00 Chair's Welcome

9:00-10:00 Keynote
Paula Pingree, NASA Jet Propulsion Laboratory, USA
Advancing NASA¹s On-Board Processing Capabilities with Reconfigurable FPGA Technologies: Opportunities & Implications

10:00-10:45 Break

Session 1: Reconfiguration Techniques and Tools
Chair: João M. P. Cardoso

10:45-11:10
PATIS: Using Partial Configuration to Improve Static FPGA Design Productivity
Tannous Frangieh, Athira Chandrasekharan, Sureshwar Rajagopalan, Yousef Iskander, Stephen Craven and Cameron Patterson
11:10-11:35
Wirelength Driven Floor Placement for FPGA-Based Partial Reconfigurable Systems
Alessio Montone, Marco D Santambrogio and Donatella Sciuto
11:35-12:00
Fast dynamic and partial reconfiguration Data Path with low Hardware overhead on Xilinx FPGAs
Michael Hübner, Diana Göhringer, Juanjo Noguera and Jürgen Becker

12:00-13:30 Lunch

Session 2: Multicore/Multiprocessing Systems
Chair: Marco D. Santambrogio

13:30-13:55
A Reconfigurable Architecture for Multicore Systems
Annie Avakian, Jon Nafziger, Amayika Panda and Ranga Vemuri

13:55-14:20
A Shared Reconfigurable VLIW Multiprocessor System
Fakhar Anjam, Stephan Wong and M. Faisal Nadeem

14:20-14:45
TLP and ILP exploitation through Reconfigurable Multiprocessing System
Mateus Rutzig, Felipe Madruga, Marco Antonio Alves, Henrique Cota, Antonio C.S. Beck, Nicolas Maillard, Philippe O. A. Navaux and Luigi Carro

14:45-15:10
CAP-OS: Operating System for Runtime Scheduling, Task Mapping and Resource Management on Reconfigurable Multiprocessor Architectures
Diana Göhringer, Michael Hübner, Etienne Nguepi Zeutebouo and Jürgen Becker

15:10-17:00 Break and Poster Session

Poster Session, Monday 19th

Flexible IP cores for the k-NN classification problem and their FPGA implementation
Elias S. Manolakos and Ioannis Stamoulias

Automatic Mapping of Control-Intensive Kernels onto coarse-grained reconfigurable architecture with speculative execution
Ganghee Lee, Kyungwook Chang and Kiyoung Choi

Virtual Area Management: Multitasking on Dynamically Partially Reconfigurable Devices
Josef Angermeier, Sándor P. Fekete, Tom Kamphans, Nils Schweer and Jürgen Teich

Self-Configurable Architecture for Reusable Systems with Accelerated Relocation Circuit (SCARS-ARC)
Adarsha Sreeramareddy, Ramachandra Kallam, Aravind Dasu and Ali Akoglu

Reconfiguration-aware Spectrum Sharing for FPGA based Software Defined Radio
Hessam Kooti, Eli Bozorgzadeh, Shenghui Liao and Lichun Bao

Implementation of the Compression Function for Selected SHA-3 Candidates on FPGA
Ashkan Namin and Anwar Hasan

Improving application performance with hardware data structures
Ravikesh Chandra and Oliver Sinnen

Adaptive Traffic Scheduling Techniques for Mixed Real-Time and Streaming Applications on Reconfigurable Hardware
Tobias Ziermann and Jürgen Teich

Reconfigurable Architecture for Mathematical Morphology Using Genetic Programming and FPGAs
Emerson Carlos Pedrino, Osmar Ogashawara and Valentin Obac Roda

MU Decoders: A Class of Fast and Efficient Configurable Decoders
Matthew C. Jordan and Ramachandran Vaidyanathan

Analysis and validation of partially dynamically reconfigurable architecture based on Xilinx FPGAs
Marco D Santambrogio, Paolo Roberto Grassi, Davide Candiloro and Donatella Sciuto

Stack Protection Unit as a step towards securing MPSoCs
Slobodan Lukovic, Paolo Pezzino and Leandro Fiorin

Fast Smith-Waterman hardware implementation
Zubair Nawaz, Koen Bertels and H. Ekin Sümbül

Session 3: System Level Design
Chair: Michael Huebner

17:00-17:25
High-Level Synthesis Techniques for In-Circuit Assertion-Based Verification
John Curreri, Greg Stitt and Alan D. George

17:25-17:50
Support of Cross Calls between a Microprocessor and FPGA in CPU-FPGA Coupling Architecture
Giang Huong and Seon Wook Kim

17:50-18:15
An Architectural Space Exploration Tool for Domain Specific Reconfigurable Computing
Gayatri Mehta and Alex K. Jones

Tuesday 20th April Schedule for RAW

8:00-9:30 IPDPS Keynote (Burton Smith)

9:30-10:00 Break

10:00-11:00 RAW Keynote
Ivo Bolsens, CTO and a Senior VP at Xilinx Inc., USA
Programming Customized Parallel Architectures in FPGA

Session 4: Architectures and Components
Chair: R. Vaidyanathan

11:00-11:25
Memory Architecture Template for Fast Block Matching Algorithms on FPGAs
Shant Chandrakar, Abraham Clements, Arvind Sudarasanam and Aravind Dasu

11:25-11:50
A Low-Energy Approach for Context Memory in Reconfigurable Systems
Thiago Ló, Antonio Carlos Beck, Mateus Rutzig and Luigi Carro

11:50-12:15
Efficient Floating-Point Logarithm Unit For FPGAs
Nikolaos Alachiotis and Alexandros Stamatakis

12:15-13:45 Lunch

Session 4 Applications and Acceleration Engines
Chair: Eli Bozorgzadeh

13:45-14:10
A Configurable-Hardware Document-Similarity Classifier to Detect Web Attacks
Craig Ulmer and Maya Gokhale

14:10-14:35
A Configurable High-Throughput Linear Sorter System
Jorge Ortiz and David Andrews

14:35-15:00
Hardware Implementation for Scalable Lookahead Regular Expression Detection
Masanori Bando, Nabi Sertac Artan, Nishit Mehta, Yi Guan and H. Jonathan Chao

15:00-15:25
A GPU-Inspired Soft Processor for High-Throughput Acceleration
Jeffrey Kingyens and J. Gregory Steffan

15:25-15:40 Closing remarks

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

Description:
The 15th HIPS workshop focuses on high-level programming of multiprocessors, compute clusters, and massively parallel machines. Like previous workshops in the series, which was established in 1996, this event serves as a forum for research in the areas of parallel applications, language design, compilers, runtime systems, and programming tools. It provides a timely and lightweight forum for scientists and engineers to present the latest ideas and findings in these rapidly changing fields.

Workshop Chair
Felix Wolf, German Research School for Simulation Sciences, Germany

Steering Committee
Rudolf Eigenmann, Purdue University, USA
Michael Gerndt, Technische Universitaet Muenchen, Germany
Frank Mueller, North Carolina State University, USA
Craig Rasmussen, Los Alamos National Laboratory, USA
Martin Schulz, Lawrence Livermore National Laboratory, USA

HIPS Program

8:45 Welcome
Felix Wolf

9:00 Keynote
Sustained Petaflop/s Computing: An Application Developer's Perspective
Thomas Schulthess

10:00 Break

Session: High-level programming models

10:30
The Gozer Workflow System
Jason Madden, Nicolas Grounds, Jay Sachs, and John Antonio

11:00
Static Macro Data Flow: Compiling Global Control into Local Control
Pritish Jetley and Laxmikant Kale

11:30
False Conflict Reduction in the Swiss Transactional Memory (SwissTM) System
Aravind Natarajan and Neeraj Mittal

12:00 Lunch break

Session: Applications and architectures

1:30
Transforming Linear Algebra Libraries: From Abstraction to Parallelism
Ernie Chan, Jim Nagle, Robert van de Geijn, and Field Van Zee

2:00
AUTO-GC: Automatic Translation of Data Mining Applications to GPU Clusters
Wenjing Ma and Gagan Agrawal

2:30
Dense Linear Algebra Solvers for Multicore with GPU Accelerators
Stanimire Tomov, Rajib Nath, Hatem Ltaief, and Jack Dongarra

3:00 Break

Session: Program analysis and tools

3:30
Experiences of Using a Dependence Profiler to Assist Parallelization for Multi-cores
Dibyendu Das and Peng Wu

4:00
Integrating Parallel Application Development with Performance Analysis in Periscope
Ventsislav Petkov and Michael Gerndt

4:30
Handling Errors in Parallel Programs Based on "Happens Before" Relations
Nicholas Matsakis and Thomas Gross

5:00 Closing

Workshop 4: NIDISC
Workshop on Nature Inspired Distributed Computing

Description:
Techniques based on metaheuristics and nature-inspired paradigms can provide efficient solutions to a wide variety of problems. Moreover, parallel and distributed metaheuristics can be used to provide more powerful problem solving environments in a variety of fields, ranging, for example, from finance to bio- and health-informatics. This workshop seeks to provide an opportunity for researchers to explore the connection between metaheuristics and the development of solutions to problems that arise in operations research, parallel computing, telecommunications, and many others.

General Chair
Albert Y. Zomaya, The University of Sydney, Australia

Program Co-chairs
Enrique Alba, University of Málaga, Spain
Franciszek Seredynski, Polish Academy of Sciences, Poland
El-ghazali Talbi, University of Lille, INRIA, France

NIDISC Program

08:30-08:40
NIDISC 2010 Opening

08:40-10:00
Session 1: Application of bio-inspired algorithms

Evolving hybrid time-shuffled behavior of agents
P. Ediger, R. Hoffmann

Diagnosing permanent faults in distributed and parallel computing systems using artificial neural networks
M. Elhadef

CA-based generator of S-boxes for cryptography use
M. Szaban, F. Seredynski

Modeling memory resources distribution on multicore processors using games on cellular lattices
M-A. I. Tsompanas, G. Ch. Sirakoulis, I. Karafyllidis

10:00-10:30 Break

10:30-11:30
Session 2: Parallel metaheuristics

pALS: An object oriented framework for developing parallel cooperative metaheuristics
A. Bernal, H. Castro

A new parallel asynchronous cellular genetic algorithm for scheduling in Grids
F. Pinel, B. Dorronsoro, P. Bouvry

Heterogeneous parallel algorithms to solve Epistatic problems
C. Salto, E. Alba

11:30-13:40 LUNCH

13:40-15:00
Session 3: Swarm Intelligence

Particle swarm optimization to solve the vehicle routing problem with heterogeneous fleet, mixed Backauls and time windows
F. Belmecheri, C. Prins, F. Yalaoui, L. Amodeo

Particle swarm optimization under fuzzy logic controller for solving a hybrid re-entrant flow shop problem
N. Yalaoui, L. Amodeo, F. Yalaoui, H. Mahdi

A survey on bee colony algorithms
S. Bitam, M. Batouche, E-G. Talbi

A bio-inspired coverage-aware scheduling scheme for wireless networks
X. Fei, S. Samarah, A. Boukerche

Workshop 5: HiCOMB Workshop on High Performance Computational Biology

Description:
Computational Biology and related disciplines are fast emerging as an important area for academic research and industrial application. The large size of biological data sets, inherent complexity of biological problems and the ability to deal with error-prone data all result in large run-time and memory requirements. The goal of this workshop is to provide a forum for discussion of latest research in developing high-performance computing solutions to problems arising from molecular biology and related life sciences areas. We are especially interested in parallel algorithms, memory-efficient algorithms, large scale data mining techniques, and design of high-performance software.

Workshop Co-Chairs
Srinivas Aluru, Iowa State University
David A. Bader, Georgia Institute of Technology

Program Chair
George Karypis, University of Minnesota

HiCOMB Program

8:45- 9:00: WELCOME
David A. Bader and Srinivas Aluru, Workshop Co-Chairs
George Karypis, Program Chair

9:00-10:00: KEYNOTE TALK
Analysis and Design of Organizational Computational Science
Christopher A. Lynberg
Centers for Disease Control and Prevention
IT Research & Development

10:00-10:30: BREAK

10:30-12:30: SESSION 1: Emerging HPC Architectures

Acceleration of Spiking Neural Networks in Emerging Multi-core and GPU Architectures
Mohammad A. Bhuiyan, Vivek K. Pallipuram, and Melissa Smith

GPU-Accelerated Multi-scoring Functions Protein Loop Structure Sampling
Yaohang Li and Weihang Zhu

A Tile-based Parallel Viterbi Algorithm for Biological Sequence Alignment on GPU with CUDA
Zhihui Du, Zhaoming Yin, and David A. Bader

Fast Binding Site Mapping using GPU and CUDA
Bharat Sukhwani and Martin C. Herbordt

12:30- 1:30: LUNCH (on your own)

1:30- 3:00: SESSION 2: Sequence Analysis

Exploring Parallelism in Short Sequence Mapping Using Burrows-Wheeler Transform
Doruk Bozdag, Ayat Hatem, and Umit V. Catalyurek

pFANGS: Parallel High Speed Sequence Mapping for Next Generation 454-Roche Sequencing Reads
Sanchit Misra, Ramanathan Narayanan, Wei-keng Liao, Alok Choudhary, and Simon Lin

Efficient and scalable parallel reconstruction of sibling relationships from genetic data in wild- populations
Saad Sheikh, Ashfaq Khokhar, and Tanya Berger-Wolf

3:00- 3:30: BREAK

3:30- 5:00: SESSION 3: HPC for Bioinformatics

Hybrid MPI/Pthreads Parallelization of the RAxML Phylogenetics Code
Wayne Pfeiffer and Alexandros Stamatakis

Measuring Properties of Molecular Surfaces Using Ray Casting
Mike Phillips, Iliyan Georgiev, Anna Dehof, Stefan Nickels, Lukas Marsalek, Hans-Peter Lenhof, Andreas Hildebrandt, and Philipp Slusallek

On the Parallelisation of MCMC-based Image Processing
Jonathan M R Byrd, Stephen A Jarvis and Abhir H. Bhalerao

5:00: END OF WORKSHOP

Workshop 6: APDCM
Advances in Parallel and Distributed Computing Models

Description:
The past twenty years have seen a flurry of activity in the arena of parallel and distributed computing. In recent years, novel parallel and distributed computational models have been proposed in the literature, reflecting advances in new computational devices and environments such as optical interconnects, programmable logic arrays, networks of workstations, radio communications, mobile computing, DNA computing, quantum computing, sensor networks etc. It is very encouraging to note that the advent of these new models has lead to significant advances in the resolution of various difficult problems of practical interest. The main goal of this workshop is to provide a timely forum for the exchange and dissemination of new ideas, techniques and research in the field of the parallel and distributed computational models. The workshop is meant to bring together researchers and practitioners interested in all aspects of parallel and distributed computing taken in an inclusive, rather than exclusive, sense. We are convinced that the workshop atmosphere will be conducive to open and mutually beneficial exchanges of ideas between the participants.

Workshop Chair
Oscar H. Ibarra, University of California, Santa Barbara, USA

Program Co-Chairs
Koji Nakano, Hiroshima University, Japan
Jacir L. Bordim, University of Brasilia, Brazil
Akihiro Fujiwara, Kyushu Institute of Technology, Japan

APDCM Program

Session 1: 8:20-10:00 (100 min)
Chair: Martti Forsell (University of Oulu)

Cross Layer Design of Heterogeneous Virtual MIMO Radio Networks with Multi-Optimization
Wei Chen, Heh Miao, Liang Hong, Jim Savage and Husam Adas

Performance Analysis and Evaluation of Random Walk Algorithms on Wireless Networks
Keqin Li

A random walk based clustering with local recomputations for mobile ad hoc networks
Alain Bui, Abdurusul Kudireti, and Devan Sohier

Distributed Tree Decomposition of Graphs and Applications to Verification
Stephane Grumbach and Zhilin Wu

Efficient Traffic Simulation using the GCA Model
Christian Schäck, Rolf Hoffmann, and Wolfgang Heenes

10:00-10:30 Break

Session 2: 10:30-12:10 (100 min)
Chair:Wei Chen (Tennessee State University)

Randomized Self-Stabilizing Leader Election in Preference-Based Anonymous Trees
Daniel Fajardo-Delgado, José Alberto Fernández-Zepeda and Anu G. Bourgeois

Self-Stabilizing Master-Slave Token Circulation and Efficient Size-Computation in a Unidirectional Ring of Arbitrary Size
Wayne Goddard and Pradip Srimani

Stability of a localized and greedy routing algorithm
Christelle Caillouet, Florian Huc, Nicolas Nisse, Stephane Perennes and Herve Rivano

A PRAM-NUMA Model of Computation for Addressing Low-TLP Workloads
Martti Forsell

Parallel Discrete Wavelet Transform using the Open Computing Language: a performance and portability study
Bharatkumar Sharma and Nagavijayalakshmi Vydyanathan

12:10-13:20 Lunch

Session 3: 13:20-15:00 (100 min)
Chair: Koji Nakano (Hiroshima University)

Throughput optimization for micro-factories subject to task and machine failures
Anne Benoit, Alexandru Dobrila, Jean-Marc Nicod, and Laurent Philippe

Collaborative Execution Environment for Heterogeneous Parallel Systems
Aleksandar Ilic and Leonel Sousa

Modeling and Analysis of Real-Time Systems with Mutex Components
Guoqiang Li, Xiaojuan Cai and Shoji Yuen

Detecting and Using Critical Paths at Runtime in Message Driven Parallel Programs
Isaac Dooley and Laxmikant Kale

Polylogarithmic Time Simulation of Reconfigurable Row/Column Buses by Static Buses
Susumu Matsumae

15:00-15:30 Break

Session 4: 15:30-17:10 (100 min)
Chair: Susumu Matsumae (Saga University)

OpenCL - An effective programming model for data parallel computations at the Cell Broadband Engine
Jens Breitbart and Claudia Fohry

Parallel external sorting for CUDA-enabled GPUs with load balancing and low transfer overhead
Hagen Peters, Ole Schulz-Hildebrandt and Norbert Luttenberger

An Efficient GPU Implementation of the Revised Simplex Method
Jakob Bieling,Patrick Peschlow and Peter Martini

Accelerating Mutual-Information-Based Registration on Multi-Core Systems
Jian Shen, Yurong Chen, He Li, Yimin Zhang and Yinlong Xu

Efficient Exhaustive Verification of the Collatz Conjecture using DSP48E blocks of Xilinx Virtex-5 FPGAs
Yasuaki Ito and Koji Nakano

Note: All presentations will be of 20 minutes duration, which includes a 15-minute presentation and 5-minute question and discussion time.

Workshop 7: CAC Communication Architecture for Clusters

Description:
Many of the world's largest and fastest computer systems are workstation or server clusters. Numerous research groups in academia, industry, and government are currently engaged in cluster research, seeking new ways to advance the state of the art of cluster communication. The goal of the CAC workshop is to bring together researchers and practitioners working in the areas of communication hardware and software to discuss their latest findings as well as future trends in the design of scalable, high-performance, and cost-effective communication architectures for clusters.

Co-chairs
Scott Pakin (LANL)
Craig Stunkel (IBM)
José Flich (UPV)

CAC Program

8:30 - 8:45
Welcome and workshop introduction

8:45 - 10:00 Session I
Topic: MPI and PGAS
Chair: Taisuke Boku, University of Tsukuba

Optimizing MPI Communication Within Large Multicore Nodes with Kernel Assistance
Stéphanie Moreaud, Brice Goglin, Dave Goodell, Raymond Namyst

Acceleration for MPI Derived Datatypes Using an Enhancer of Memory and Network
Noboru Tanabe, Hironori Nakajo

Efficient Hardware Support for the Partitioned Global Address Space
Holger Fröning, Heiner Litz

10:00 - 10:30 Break

10:30 - 11:45 Session II
Topic: InfiniBand
Chair: T.B.D.

Overlapping Computation and Communication: Barrier Algorithms and Connectx-2 CORE-Direct Capabilities
Richard Graham, Stephen W. Poole, Pavel Shamis, Gil Bloch, Noam Boch, HIllel Chapman, Michael Kagan, Arial Shahar, Ishai Rabinovitz, Gilad Shainer

Designing Topology-Aware Collective Communication Algorithms for Large Scale Infiniband Clusters: Case Studies with Scatter and Gather
Krishna Chaitanya Kandalla, Hari Subramoni, Dhabaleswar Panda

Designing High-Performance and Resilient Message Passing on Infiniband
Matthew J. Koop, Pavel Shamis, Ishai Rabinovitz, Dhabaleswar Panda

11:45 - 1:00 Lunch (on your own)

1:00 - 3:00 Session III
Topic: Scalable stream processing systems
Chair: Henrique Andrade, IBM

Index Tuning for Adaptive Multi-Route Data Stream Systems
Karen Works, Elke Rundensteiner, Emmanuel Agu

Towards Execution Guarantees for Stream Queries
Rafael Fernández-Moctezuma, David Maier, Kristin Tufte

Exploiting Constraints to Build a Flexible and Extensible Data Stream Processing Middleware
Nazario Cipriani, Carlos Lübbe, Alexander Moosbrugger

Distributed Monitoring of Conditional Entropy for Anomaly Detection in Streams
Chrisil Arackaparambil, Sergey Bratus, Joshua Brody, Anna Shubina

3:00 Workshop closing

Workshop 8: HPPAC
High-Performance, Power-Aware Computing

Description:
High-performance computing is and has always been performance-oriented. However, a consequence of the push towards maximum performance is increased energy consumption, especially in datacenters and supercomputing centers. Moreover, as peak performance is rarely attained, some of this energy consumption results in little or no performance gain. In addition, large energy consumption costs datacenters and supercomputing centers a significant amount of money and wastes natural resources. The main goal of this workshop is to provide a timely forum for the exchange and dissemination of new ideas, techniques, and research in high-performance, power-aware computing (HPPAC). HPPAC presents research that reduces (1) power consumption, (2) energy consumption, or (3) heat generation with little or no performance penalty in high-performance computing systems. In effect, the workshop aims to move towards "greener" solutions for datacenters and supercomputing centers.

Workshop Co-Chairs
Rong Ge, Marquette University, USA
Karsten Schwan, Georgia Tech, USA

HPPAC Program

9:00-9:15 Welcome

9:15-10:15 Keynote Address:
Laurent Lefevre, INRIA
Between Gaps and Bursts: Supporting GREEN-Approaches to Reduce Energy Consumption in Large Scale Distributed Systems

10:15-10:30 Break

10:30-12:00 Session: Power/performance measurement, analysis, and modeling

The Green500 List: Year Two
Wu-chun Feng, Heshan Lin

Characterizing Energy Efficiency of I/O Intensive Parallel Applications on Power-Aware Clusters
Rong Ge, Xizhou Feng, Sindhu Subramanya, Xian-he Sun.

VMeter: Power Modelling for Virtualized Clouds
Ata Bohra, Vipin Chaudhary

12:00-1:30 Lunch

1:30-3:30 Session: Power management in large-scale systems

Reducing Grid Energy Consumption through Choice of Resource Allocation Method
Timothy M. Lynar, Ric D. Herbert, Simon, and William J. Chivers

BSLD Threshold Driven Power Management Policy for HPC Centers
Maja Etinski, Julita Corbalan, Jesús Labarta, Mateo Valero

Scheduling Parallel Tasks on Multiprocessor Computers with Efficient Power Management
Keqin Li

Performance Evaluation of a Green Scheduling Algorithm for Energy Savings in Cloud Computing
Truong Vinh Truong Duy, Yukinori Sato and Yasushi Inoguchi

3:30-3:45 Break

3:45-5:15 Session: Power aware microarchitecture

T-NUCA - A Novel Approach to Non-Uniform Access Latency Cache Architectures for 3D CMPs
Konrad Malkowski, Padma Raghavan, Mahmut Kandemir, Mary Jane Irwin

Integrated Energy-Aware Cyclic and Acyclic Scheduling for Clustered VLIW Processors
Jimmy Bahuleyan, Rahul Nagpal, Srikant YN

Dynamic Core Partitioning for Energy Efficiency
Yang Ding, Mahmut Kandemir, Mary Jane Irwin, Padma Raghavan

5:15-5:45 Discussion

Workshop 9: HPGC
High Performance Grid Computing

Description:
This is the seventh annual workshop which provides a forum for researchers and engineers to present their results in grid and distributed computing. Special areas of interest will be grid middleware, grid applications, grid benchmarking, data distribution and replication on the grid, fault tolerance and efficiency of grid applications, novel models of grid computing, including cloud and mobile computing.

Workshop Organizers
Eric Aubanel, University of New Brunswick, Canada
Virendra C. Bhavsar, University of New Brunswick, Canada
Michael Alex Frumkin, Google, USA

HPGC Program

9:00 - 10:00
Welcome to HPGC
(Chair: Michael Frumkin, Google)

Keynote talk
Computational Science and Engineering & Clouds
Manish Parashar (Rutgers, The State University of New Jersey)

10:00 - 10:30 Break

10:30 - 12:00 pm
Session 1: Grid Infrastructures and Standards
(Chair: Eric Aubanel, University of New Brunswick, Canada)

An Interoperable & Optimal Data Grid Solution for Heterogeneous and SOA based Grid-GARUDA
Payal Saluja, B.B. Prahlada Rao, V. Shashidhar, Neetu Sharma, and A. Paventhan (CDAC, India)

Improvements of Common Open Grid Standards to Increase High Throughput and High Performance Computing Effectiveness on Large-scale Grid and e-Science Infrastructures
M. Riedel, M.S. Memon, A.S. Memon, A. Streit, F. Wolf, Th. Lippert, A. Konstaninov, M. Marzolla, B. Konya, O. Smirnova, L. Zangrando, J. Watzl and D. Kranzlmüller (F. Jülich, GRS, LMU, Germany; U. Lund, Sweden; U. di Bologna, Italy)

A Distributed Diffusive Heuristic for Clustering a Virtual P2P Supercomputer
Joachim Gehweiler and Henning Meyerhenke (University of Paderborn, Germany)

12:00 - 1:30 Lunch

1:30 - 3:00
Session 2: Grid Resource Allocation
(Chair: Virendra Bhavsar, University of New Brunswick, Canada)

How Algorithm Definition Language (ADL) Improves the Performance of SmartGridSolve Applications
Michele Guidolin, Thomas Brady, and Alexey Lastovetsky (University College Dublin, Ireland)

GridP2P: Resource Usage in Grids and Peer-to-Peer Systems
Sergio Esteves, Luís Veiga, and Paulo Ferreira (INESC-ID, Portugal)

A Grid Simulation Framework to Study Advance Scheduling Strategies for Complex Workflow Applications
Adan H Carbajal and Andrie Tchernykh; Thomas Röblitz and Ramin Yahyapour (Autonomous U. of Baja California, CICESE, Mexico; T.U. Dortmund, Germany)

3:00 - 3:30 Break

3:30 - 5:00
Session 3: Grid Scheduling and Fault Tolerance
(Chair: Michael Frumkin, Google)

Meta-Scheduling in Advance using Red-Black Trees in Heterogeneous Grids
Luis Tomás, Carmen Carrión, Blanca Caminero and Agustín Caminero (U. of Castilla la Mancha)

SPSE: A Flexible QoS-based Service Scheduling Algorithm in Service-Oriented Grid
Laiping Zhao, Yizhi Ren, Mingchu Li, and Kouichi Sakurai (Kyushu U., Japan)

Fault-Tolerance for PastryGrid Middleware
Heithem Abbes, Christophe Cerin, Mohamed Jemni, and Yazid Missaoui (UTIC, Tunisia; U. Paris XIII, France)

5:00 Concluding remarks

Workshop 10: SMTPS
Workshop on System Management Techniques, Processes and Services

Description:
This workshop is intended to bring together researchers and practitioners to define and discuss the issues and challenges and share early findings in systems management in the emerging cloud computing paradigm. With advent of virtualization technologies, efficient and effective systems management plays an even more critical role to deal with the ever-increasing computing resources and their connectivity and complexity. The cloud computing paradigm has been introduced to better respond to dynamic computation demand by emerging applications and workload, and to better utilize computing resources and energy in aggregation to serve a large group of users. Many issues in systems management within a traditional computing environment remain challenging for cloud computing. For instance, system administrators for both traditional and cloud systems are required to provide satisfactory and continuous services to applications even in the presence of interrupts and system failures. On the other hand, there are new challenges and opportunities posed by cloud computing, which call for different solutions Management of computing infrastructure, platform and software are redefined to be part of online service offerings, commonly termed as "infrastructure as a Service" and "software as a service". Management of future datacenters with tens of thousands of severs, a wide variety of middleware, and emerging web portals brings in new challenges for tomorrow's system administrators and architects.

Workshop General Co-Chairs
Yanyong Zhang, Rutgers Univ.
Walfredo Cirne, Google

Program Co-Chairs
Kyung Dong Ryu, IBM Research
Ramendra K. Sahoo, Citigroup

SMTPS Program

8:30 –8:45 Welcoming Remarks

8:45 –9:45 Keynote Address
Managing Large-scale Utility Cloud
Karsten Schwan (Georgia Tech)

9:45 –10:00 Break

Session 1: Distributed Systems (10:00 – 12:00)

10:00 – 10:30
ROME: Road Monitoring and Alert System through Geocache
Bin Zan, Tingting Sun, Marco Gruteser, Yanyong Zhang (Winlab, Rutgers University)
10:30 – 11:00
Streaming, Low-latency Communication in On-line Trading Systems
Hari Subramoni, Fabrizio Petrini, Virat Agarwal (IBM T. J. Watson), Davide Pasetto (IBM, Ireland)
11:00 –11:30
Autonomic Management of Distributed Systems using Online Clustering
Andres Quiroz, Manish Parashar, Ivan Rodero (Rutgers University)
11:30 – 12:00
Business-Driven Capacity Planning of a Cloud-based IT Infrastructure for the Execution of Web Applications
Raquel Lopes; Francisco Brasileiro; Paulo Maciel Jr. (Universidade Federal de Campina Grande)

12:00 – 2:00 Lunch

Session 2: Cloud Optimization (2:00 – 3:30)

2:00 –2:30
Desktop Workload Study with Implications for Desktop Cloud Resource Optimization
Andrzej Kochut, Kirk Beaty, Hidayatullah Shaikh, Dennis Shea (IBM T. J. Watson)
2:30 –3:00
Simplifying solution deployment on a Cloud through composite appliances
Trieu Chieu, Alexei Karve, Ajay Mohindra and Alla Segal (IBM)

3:00 –3:30
Combining Virtualization, Resource Characterization, and Resource Management to Enable Efficient High Performance Compute Platforms Through Intelligent Dynamic Resource Allocation
Jim Brandt; Frank Chen; Vincent De Sapio; Ann Gentile; Jackson Mayo; Philippe Pebay; Diana Roe; David Thompson; Matthew Wong (Sandia National Lab)

3:30 –4:00 Break

Session 3: Performance and Scalability (4:00 –5:30)

4:00 – 4:30
Scalability Analysis of Embarassingly Parallel Applications on Large Clusters
Fabrício A. B. Silva and Hermes Senger (Universidade de Lisboa)
4:30 –5:00
Automation and Management of Scientific Workflows in Distributed Network Environments
Qishi Wu (Univ. of Memphis); Mengxia Zhu (Southern Illinois Univ.); Xukang Lu (Univ. of Memphis); Patrick Brown (Southern Illinois Univ.); Yunyue Lin (Univ. of Memphis); Yi Gu (Univ. of Memphis); Fei Cao (Southern Illinois Univ.); Michael Reuter (Oak Ridge National Lab.)
5:00 –5:30
Initial Characterization of Parallel NFS Implementations
Weikuan Yu (Auburn University); Jeffrey S. Vetter (Oak Ridge National Lab.)

5:30–5:45 Closing Remarks

TCPP Membership Meeting & Reception

6:00 – 7:00 PM Reception

7:00 PM Brief meeting followed by

TCPP Invited Speaker:
Craig Stunkel
IBM T. J. Watson Research Center

Title of talk:
Exascale: Parallelism Gone Wild!

IPDPS Sponsor
IEEE Computer Society
Technical Committee on Parallel Processing
Sushil K. Prasad, Chair

Open to all IPDPS Attendees

Opening Session 8:00 AM - 8:30 AM
Chair David Bader

Keynote Session 8:30 AM - 9:30 AM
Chair Andrew Lumsdaine

Keynote Speech: Operating System Resource Management

Speaker: Burton Smith
Technical Fellow, Microsoft Corporation

Abstract: Resource management is the dynamic allocation and de-allocation by an operating system of processor cores, memory pages, and various types of bandwidth to computations that compete for those resources. The objective is to allocate resources so as to optimize responsiveness subject to the finite resources available. Historically, resource management solutions have been relatively unsystematic, and now the very assumptions underlying the traditional strategies fail to hold. First, applications increasingly differ in their ability to exploit resources, especially processor cores. Second, application responsiveness is approximately two-valued for "Quality-Of-Service" (QOS) applications, depending on whether deadlines are met. Third, power and battery energy have become constrained. This talk will propose a scheme for addressing the operating system resource management problem.

Morning Break 9:30 AM - 10:00 AM

Tuesday 20th April Schedule

Commercial Book Exhibits - All Day

Parallel Technical Sessions 1, 2, 3, 4, & 5
10:00 AM - 12:00 PM

Lunch 12 Noon – 1:30 PM (on your own)

PhD Forum Posters Start – 12 Noon
(on display until Wednesday evening)

Parallel Technical Sessions 6, 7, 8, 9, & 10
1:30 PM – 3:30 PM

Afternoon Break 3:30 PM – 4:00 PM

Parallel Technical Sessions 11, 12, 13, 14, & 15
4:00 PM - 6:00 PM

IEEE-TPDS Meeting
6:00 PM - 7:00 PM

All Symposium Tutorial
MapReduce Programming with Apache Hadoop
7:00 PM – 10:00 PM

Session 1
Algorithms for Network Management
Chair: Anne Benoit

Distributed Advance Network Reservation with Delay Guarantees
Niloofar Fazlollahi (Boston University, USA); David Starobinski (Boston University, USA)

A General Algorithm for Detecting Faults under the Comparison Diagnosis Model
Iain A Stewart (Durham University, UK)

On the Importance of Bandwidth Control Mechanisms for Scheduling on Large Scale Heterogeneous Platforms
Olivier Beaumont (INRIA, FR); Hejer Rejeb (LaBRI-INRIA, FR)

Broadcasting on Large Scale Heterogeneous Platforms under the Bounded Multi-Port Model
Olivier Beaumont (INRIA, FR); Lionel Eyraud-Dubois (INRIA Bordeaux Sud-Ouest, FR); Shailesh Kumar Agrawal (INRIA, FR)

Session 2
Scientific Computing with GPUs
Chair: Ling Zhou

Improving Numerical Reproducibility and Stability in Large-Scale Numerical Simulations on GPUs
Michela Taufer (University of Delaware, US); Philip Saponaro (University of Delaware, US); Omar Padron (Kean University, US); Sandeep Patel (University of Delaware, US)

Implementing the Himeno Benchmark with CUDA on GPU Clusters
Everett Phillips (NVIDIA, US); Massimiliano Fatica (NVIDIA, US)

Direct Self-Consistent Field Computations on GPU Clusters
Guochun Shi, Volodymyr Kindratenko (National Center for Supercomputing Applications, US); Ivan Ufimtsev, Todd Martinez (Stanford University, US)

Parallelization of Tau-Leap Coarse-Grained Monte Carlo Simulations on GPUs
Lifan Xu (University of Delaware, US); Michela Taufer (University of Delaware, US); Stuart Collins (University of Delaware, US); Dionisios Vlacho (University of Delaware, US)

Session 3
Data Storage and Memory Systems
Chair: Bradley Kuszmaul

DEBAR: A Scalable High-Performance De-duplication Storage System for Backup and Archiving
Tianming Yang (Huazhong University of Science and Technology, PRC); Hong Jiang (University of Nebraska, US); Dan Feng (Huazhong University of Science and Technology, PRC); Zhongying Niu (Huazhong University of Science and Technology, PRC); Ke Zhou (Huazhong University of Science and Technology, PRC); Yaping Wan (Huazhong University of Science and Technology, PRC)

HPDA: A Hybrid Parity-based Disk Array for Enhanced Performance and Reliability
Bo Mao (Huazhong University of Science and Technology, PRC); Hong Jiang (University of Nebraska, US); Dan Feng (Huazhong University of Science and Technology, PRC); Suzhen Wu (Huazhong University of Science and Technology, PRC); Jianxi Chen (Huazhong University of Science and Technology, PRC); Lingfang Zeng (Huazhong University of Science and Technology, PRC); Lei Tian (Huazhong University of Science and Technology, PRC)

Fine-Grained QoS Scheduling for PCM-based Main Memory Systems
Ping Zhou (University of Pittsburgh, US); Yu Du (University of Pittsburgh, US); Youtao Zhang (University of Pittsburgh, US); Jun Yang (University of Pittsburgh, US)

Performance Impact of Resource Contention in Multicore Systems
Robert Hood (CSC-NASA Ames, US); Haoqiang Jin (NASA Ames Research Center, US); Piyush Mehrotra (NASA Ames Research Center, US); Johnny Chang (CSC-NASA Ames Research Center, US); Jahed Djomehri (NASA Ames Research Center, US); Sharad Gavali (NASA Ames Research Center, US); Dennis Jespersen (NASA Ames Research Center, US); Kenichi Taylor (Silicon Graphics International, US); and Rupak Biswas (NASA Ames Research Center, US)

Session 4
Fault Tolerance
Chair: Almadena Chtchelkanova

Improving the Performance of Hypervisor-Based Fault Tolerance
Jun Zhu (Peking University, PRC); Wei Dong (Peking University, PRC); ZheFu Jiang (Peking University, PRC); Xiaogang Shi (Peking University, PRC); Zhen Xiao (Peking University, PRC); XiaoMing Li (Peking University, PRC)

Supporting Fault Tolerance in a Data-Intensive Computing Middleware
Tekin Bicer (The Ohio State University, US); Wei Jiang (The Ohio State University, US); Gagan Agrawal (The Ohio State University, US)

A High-Performance Fault-Tolerant Software Framework for Memory on Commodity GPUs
Naoya Maruyama (Tokyo Institute of Technology, JPN); Akira Nukada (Tokyo Institute of Technology, JPN); Satoshi Matsuoka (Tokyo Institute of Technology, JPN)

Scalable Failure Recovery for High-performance Data Aggregation
Dorian Arnold (University of New Mexico, US); Barton Miller (University of Wisconsin, US)

Session 5
Sorting
Chair: George Biros

High Performance Comparison-Based Sorting Algorithm on Many-Core GPUs
Xiaochun Ye (Chinese Academy of Sciences, PRC); Dongrui Fan (Chinese Academy of Sciences, PRC); Wei Lin (Chinese Academy of Sciences, PRC); Nan Yuan (Chinese Academy of Sciences, PRC); Paolo Ienne (EPFL, Switzerland)

GPU Sample Sort
Vitaly Osipov (Karlsruhe Institute of Technology, Germany); Peter Sanders (University of Karlsruhe, Germany); Nikolaj Leischner (University of Karlsruhe, Germany)

Highly Scalable Parallel Sorting
Edgar Solomonik (University of Illinois at Urbana-Champaign, US); Laxmikant Kale (University of Illinois at Urbana-Champaign, US)

Session 6
Scheduling
Chair: David Bunde

A Scheduling Framework for Large-Scale, Parallel, and Topology-Aware Applications
Valentin Kravtsov (Technion - Israel Institute of Technology, Israel); Pavel Bar (Technion - Israel Institute of Technology, Israel); David Carmeli (Technion - Israel Institute of Technology, Israel); Assaf Schuster (Technion - Israel Institute of Technology, Israel); Martin Swain (Technion - Israel Institute of Technology, Israel);

Load Regulating Algorithm for Static-Priority Task Scheduling on Multiprocessors
Risat Pathan (Chalmers University of Technology, Sweden); Jan Jonsson (Chalmers University of Technology, Sweden)

Scheduling Algorithms for Linear Workflow Optimization
Kunal Agrawal (Washington University in St. Louis, US); Anne Benoit (Ecole Normale Superieure de Lyon Lyon, FR); Loic Magnan (Ecole Normale Superieure de Lyon Lyon, FR); Yves Robert (Ecole Normale Superieure de Lyon, FR)

Hypergraph-Based Task-Bundle Scheduling Towards Efficiency and Fairness in Heterogeneous Distributed Systems
Han Zhao (Oklahoma State University, US); Xinxin Liu (Oklahoma State University, US); Xiaolin (Andy) Li (Oklahoma State University, US)

Session 7
Performance/Scalability Improvement for Scientific Applications
Chair: Srinivas Aluru

Improving the Performance of Uintah: A Large-Scale Adaptive Meshing Computational Framework
Justin Luitjens (University of Utah, US); Martin Berzins (University of Utah, US)

Optimizing and Tuning the Fast Multipole Method for State-of-the-Art Multicore Architectures
Aparna Chandramowlishwaran (Georgia Institute of Technology, US); Samuel W Williams (Lawrence Berkeley National Laboratory, US); Leonid Oliker (Lawrence Berkeley National Laboratory, US); Ilya Lashuk (Georgia Institute of Technology, US); George Biros (Georgia Institute of Technology, US); Richard Vuduc (Georgia Institute of Technology, US)

Parallelization of DQMC Simulation for Strongly Correlated Electron Systems
Che-Rung Lee (National Tsing Hua University, Taiwan); I-Hsin Chung (IBM T.J. Watson Research Center, US); Zhaojun Bai (University of California, Davis, US)

Parallel I/O Performance: From Events to Ensembles
Andrew Uselton (Lawrence Berkeley National Laboratory, US); Mark Hawison (Lawrence Berkeley National Laboratory, US); Nicholas J. Wright (Lawrence Berkeley National Laboratory, US); David Skinner (Lawrence Berkeley National Laboratory, US); Noel Keen (Lawrence Berkeley National Laboratory, US); John Shalf (Lawrence Berkeley National Laboratory, US); Karen L Karavanic (Portland State University, US); Leonid Oliker (Lawrence Berkeley National Laboratory, US)

Session 8
Network Architecture and Algorithms
Chair: Neeraj Mittal

Achieve Constant Performance Guarantees using Asynchronous Crossbar Scheduling without Speedup
Deng Pan (Florida International University, US); Kia Makki (Florida International University, US); Niki Pissinou (Florida International University, US)

Distributive Waveband Assignment in Multi-granular Optical Networks
Yang Wang (Georgia State University, US); Xiaojun Cao (Georgia State University, US)

QoS Aware BiNoC Architecture
Shih-Hsin Lo (National Taiwan University, Taiwan); Ying-Cherng Lan (National Taiwan University, Taiwan); Hsin-Hsien Yeh (National Taiwan University, Taiwan); Wen-Chung Tsai (National Taiwan University, Taiwan); Yu Hen Hu (National Taiwan University, Taiwan); Sao-Jie Chen (National Taiwan University, Taiwan)

First Experiences with Congestion Control in InfiniBand Hardware
Ernst Gran (Simula Research Laboratory, Norway); Magne Eimot (Simula Research Laboratory, Norway); Sven-Arne Reinemo (Simula Research Laboratory, Norway); Tor Skeie (Simula Research Laboratory, Norway); Olav Lysne (Simula Research Laboratory, Norway); Lars Paul Huse (Simula Research Laboratory, Norway)

Session 9
Software Support for Using GPUs
Chair: Anne Elster

Object-Oriented Stream Programming using Aspects
Mingliang Wang (Rutgers University, US); Manish Parashar (Rutgers University, US)

Optimal Loop Unrolling for GPGPU Programs
Giridhar Sreenivasa Murthy (The Ohio State University, US); Muthu Ravishankar (The Ohio State University, US); Muthu Manikandan Baskaran (The Ohio State University, US); Ponnuswamy Sadayappan (The Ohio State University, US);

Speculative Execution on Multi-GPU Systems
Gregory Diamos (Georgia Institute of Technology, US); Sudakhar Yalamanchili (Georgia Institute of Technology, US)

Dynamic Load Balancing on Single- and Multi-GPU Systems
Long Chen (University of Delaware, US); Oreste Villa (Pacific Northwest National Laboratory, US); Sriram Krishnamoorthy (Pacific Northwest National Laboratory, US); Guang Gao (University of Delaware, US)

Session 10
Performance Prediction and Benchmarking Tools
Chair: George Bosilca

Servet: A Benchmark Suite for Autotuning on Multicore Clusters
Jorge González-Domínguez (University of A Coruna, Spain); Guillermo Lopez Taboada (University of A Coruna, Spain); Basilio Fraguela (University of A Coruna, Spain); María J. Martín (University of A Coruna, Spain); Juan Touriño (University of A Coruna, Spain);

KRASH: Reproducible CPU Load Generation on Many-Cores Machines
Swann Perarnau (INRIA Moais Research Team, FR); Guillaume Huard (ID Laboratory, FR)

Power-aware MPI Task Aggregation Prediction for High-End Computing Systems
Dong Li (Virginia Tech, US); Dimitrios Nikolopoulos (Foundation of Research and Technology Hellas, Greece); Kirk Cameron (Virginia Tech, US); Bronis R. de Supinski (Lawrence Livermore National Laboratory, US); Martin Schulz (Lawrence Livermore National Laboratory, US)

Session 11
Resource Allocation
Chair: Anne Benoit

Varying Bandwidth Resource Allocation Problem with Bag Constraints
Venkatesan Chakaravarthy (IBM Research, India); Vinayaka Pandit (IBM Research, India); Yogish Sabharwal (IBM Research, India); Deva Seetharam (IBM Research, India)

Decentralized Resource Management for Multi-core Desktop Grids
Jaehwan Lee (University of Maryland, College Park, US); Pete Keleher (University of Maryland, US); Alan Sussman (University of Maryland, US)

Dynamic Fractional Resource Scheduling for HPC Workloads
Mark Lee Stillwell (University of Hawaii at Manoa, US); Frédéric Vivien (INRIA, FR); Henri Casanova (University of Hawaii at Manoa)

ADEPT Scalability Predictor in Support of Adaptive Resource Allocation
Arash Deshmeh (University of Windsor, Canada); Jacob Machina (University of Windsor, Canada); Angela Sodan (University of Windsor, Canada)

Session 12
Image Processing and Data Mining
Chair: David Konerding

Exploiting the Forgiving Nature of Applications for Scalable Parallel Execution
Jiayuan Meng (University of Virginia); Anand Raghunathan (NEC Research Labs, US); Srimat Chakradhar (NEC Research Labs, US); Surendra Byna (NEC Research Labs, US)

Fisheye Lens Distortion Correction on Multicore and Hardware Accelerator Platforms
Konstantis Daloukas (University of Thessaly, Greece); Christos Antonopoulos (University of Thessaly, Greece); Nikos Bellas (University of Thessaly, Greece); Sek Chai (Motorola, US)

Large-Scale Multi-Dimensional Document Clustering on GPU Clusters
Yongpeng Zhang (North Carolina State University, US); Frank Mueller (North Carolina State University, US); Xiaohui Cui (Oak Ridge National Laboratory, US); Thomas Potok (Oak Ridge National Laboratory, US)

eScience in the Cloud: A MODIS Satellite Data Reprojection and Reduction Pipeline in Windows Azure Platform
Jie Li (University of Virginia, US); Deb Agarwal (Lawrence Berkeley National Laboratory, US); Marty Humphrey (University of Virginia, Charlottesville, US); Catharine van Ingen (Microsoft Research); Keith Jackson (Lawrence Berkeley National Laboratory, US); Youngryel Ryu (University of California at Berkeley, US)

Session 13
Transactional Memory
Chair: Anne Elster

Locality-Aware Adaptive Grain Signatures for Transactional Memories
Woojin Choi (University of Southern California, US); Jeffrey Draper (University of Southern California, US)

Dynamic Analysis of the Relay Cache-Coherence Protocol for Distributed Transactional Memory
Bo Zhang (Virginia Tech, US); Binoy Ravindran (Virginia Tech, US)

Runtime Checking of Serializability in Software Transactional Memory
Arnab Sinha (Princeton University, US); Sharad Malik (Princeton University, US)

Consistency in Hindsight, A Fully Decentralized STM Algorithm
Annette Bieniusa (University of Freiburg, US); Thomas Fuhrmann (Technische Universitat Munchen, Germany)

Session 14
Tools for Performance and Correctness Analysis
Chair: Almadena Chtchelkanova

Identifying Ad-hoc Synchronization for Enhanced Race Detection
Ali Jannesari (University of Karlsruhe, Germany); Water F. Tichy (University of Karlsruhe, Germany)

Improving the Performance of Program Monitors with Compiler Support in Multi-Core Environment
Guojin He (University of Minnesota, US); Antonia Zhai (University of Minnesota, US)

On-Line Detection of Large-Scale Parallel Application's Structure
German Llort (Barcelona Supercomputing Center, Spain); Juan Gonzalez Garcia (Universitat Politècnica de Catalunya, Spain); Harald Servat (Barcelona Supercomputing Center, Spain); Judit Gimenez (Barcelona Supercomputing Center, Spain); Jesus Labarta (Barcelona Supercomputing Center, Spain)

Adaptive Sampling-Based Profiling Techniques for Optimizing the Distributed JVM Runtime
King Tim Lam (The University of Hong Kong, Hong Kong); Yang Luo (The University of Hong Kong, Hong Kong); Cho-Li Wang (The University of Hong Kong, Hong Kong)

Session 15
Parallel Linear Algebra I
Chair: Esmond Ng

Algorithmic Cholesky Factorization Fault Recovery
Douglas Hakkarinen (Colorado School of Mines, US); Zizhong Chen (Colorado School of Mines, US)

Analyzing the Soft-Error Resiliance of Linear Solvers on Multicore Multiprocessors
Konrad Malkowski (The Pennsylvania State University); Padma Raghavan (The Pennsylvania State University); Mahmut Taylan Kandemir (The Pennsylvania State University)

A Parallel Architecture for Meaning Comparison
Suneil Mohan (Texas A&M University, US); Amitava Biswas (Texas A&M University, US); Aalap Tripathy (Texas A&M University, US); Jagannath Panigraphy (Texas A&M University, US); Rabi Mahapatra (Texas A&M University, US)

IPDPS 2010 PhD Forum

Description

The IPDPS PhD Forum focuses on current graduate students working toward a PhD in broadly defined areas related to parallel and distributed processing. Twenty-one students who submitted proposals for the IPDPS 2010 event have been selected to display a poster describing their dissertation research. Brief papers describing their research are included in the proceedings. The posters will go **on display at 12 Noon on Tuesday**, 20 April in the conference foyer and will be available for viewing **until 6:00 PM on Wednesday, at which time all of the students presenters will be on hand** to describe their work to conference attendees and answer questions regarding their research.

PhD Forum Posters

Memory Affinity Management for Numerical Scientific Applications over Multi-core Multiprocessors with Hierarchical Memory
(Christiane Pousa Ribeiro and Alexandre Carissimi and Jean-François Méhaut - Computer Science and Applied Mathematics Institute of Grenoble, France)

Performance Improvements of Real-Time Crowd Simulations
(Guillermo Vigueras and Juan Manuel Orduña and Miguel Lozano - University of Valencia, Spain)

Parallel Applications on Emerging Architectures
(Abhinav Sarje and Srinivas Aluru - Iowa State University, USA)

Fault Tolerant Linear Algebra: Recovering from Fail-Stop Failures without Checkpointing
(Teresa Davies and Zizhong Chen - Colorado School of Mines, USA)

Highly Scalable Checkpointing for Exascale Computing
(Christer Karlsson and Zizhong Chen - Colorado School of Mines, USA)

Performance Modeling of Heterogeneous Systems
(Jan Meyer and Anne C. Elster - Norwegian University of Science and Technology, Norway)

Large-Scale Distributed Storage for Highly Concurrent MapReduce Applications
(Diana Moise - INRIA-IRISA, France)

Scalable Verification of MPI Programs
(Anh Vo and Ganesh Gopalakrishnan - University of Utah, USA)

Ensuring Deterministic Concurrency through Compilation
(Nalini Vasudevan and Stephen A. Edwards - Columbia University, USA)

Use of Peer-To-Peer Technology in Internet Access Networks and its Impacts
(Peter Danielis and Dirk Timmermann - Universität Rostock, Germany)

A Path Based Reliable Middleware Framework for RFID Devices
(Nova Ahmed - Georgia Institute of Technology, USA)

Improving Topological Mapping on NoCs
(Rafael Tornero Gavilá and Juan Manuel Orduña - University of Valencia, Spain)

Coping with Uncertainty in Scheduling Problems
(Louis-Claude Canon - LaBRI-INRIA-Bordeaux, France)

AuctionNet: Market Oriented Task Scheduling in Heterogeneous Distributed Environments
(Han Zhao and Xiaolin Andy Li - Oklahoma State University, USA)

Towards Dynamic Reconfigurable Load-balancing for Hybrid Desktop Platforms
(Alécio Binotto and Dieter Fellner and Carlos E. Pereira - INF Federal University of Rio Grande do Sul, Brazil)

Dynamic Fractional Resource Scheduling for Cluster Platforms
(Mark Lee Stillwell - University of Hawaii, USA)

Energy-aware Joint Scheduling of Tasks and Messages in Wireless Sensor Networks
(Benazir Fateh and Manimaran Govindarasu - Iowa State University, USA)

BlobSeer: Efficient Data Management for Data-Intensive Applications Distributed at Large-Scale
(Bogdan Nicolae - INRIA-IRISA, France)

Extendable Storage Framework for Reliable Clustered Storage Systems
(Sumit Narayan and John Chandy - University of Connecticut, USA)

The Effects on Branch Prediction when Utilizing Control Independence
(Chris Michael - Louisiana State University, USA)

High Performance Reconfigurable Multi-Processor-Based Computing on FPGAs
(Diana Göhringer and Juergen Becker - Fraunhofer FOM, Germany)

Note: In cases where multiple authors are listed, the first one is the PhD student and the other is the student's advisor.

IEEE-TPDS Open Meeting 6:00 PM - 7:00 PM

IEEE Transactions on Parallel & Distributed Systems

All IPDPS attendees are invited to attend this open meeting.

Hosted by:
Ivan Stojmenovic, University of Ottawa, Canada
Editor-in-Chief, IEEE TPDS

All Symposium Tutorial
7:00 PM - 10:00 PM

(Open to all IPDPS 2010 attendees)

Title:
MapReduce Programming with Apache Hadoop

Presenter
Milind Bhandarkar
Hadoop Solutions Architect
Yahoo! Inc.

Brief
Apache Hadoop has become the platform of choice for developing large-scale data-intensive applications. In this tutorial, we will discuss design philosophy of Hadoop, and describe how to design and develop Hadoop applications and higher-level application frameworks to crunch several terabytes of data, using anywhere from four to 4,000 computers. We will discuss solutions to common problems encountered in maximizing Hadoop application performance. We will also describe several frameworks and utilities developed using Hadoop that increase programmer productivity and application performance.

Keynote Session 8:30 AM - 9:30 AM
Chair: Bradley Kuszmaul

Keynote Speech: Chip Multiprocessor Architecture: A Programmability-Driven Approach

Speaker: Kunle Olukotun
Stanford University

Abstract: Chip multiprocessors (CMPs) are now the dominant architecture in microprocessor design. However, in many software environments, due to the difficulty of writing correct and high performing parallel programs, the capability of CMPs is underutilized. In this talk, I will argue that to enable more parallel programs to be written more easily, the design of CMPs should be driven by the needs of programmability. To demonstrate the benefits of this approach, I will describe example CMP designs where a focus on programmability has resulted in a simpler programming model and excellent performance. I will also describe a programming environment where programmability can be used to drive the development of future CMP architectures.

Morning Break 9:30 AM - 10:00 AM

Wednesday 21st April Schedule

Commercial Book Exhibits - All Day
PhD Forum Posters Continue

Best Papers Plenary Session
10:00 AM - 12:00 PM

Lunch 12 Noon – 1:30 PM (on your own)

Parallel Technical Sessions 16, 17, 18, & 19
1:30 PM – 3:30 PM

Afternoon Break 3:30 PM – 4:00 PM

Parallel Technical Sessions 20, 21, 22, & 23
4:00 PM - 6:00 PM

PhD Forum
6:00 PM - 7:30 PM
Student Authors Available to Discuss Poster Research with Conference Attendees

Pre-banquet Reception (in Forum area)
6:30 PM

Symposium Banquet (Buffet Dinner)
7:30 PM

Plenary Session – Best Papers
Chair: Cynthia Phillips

Extreme Scale Computing: Modeling the Impact of System Noise in Multicore Clustered Systems
Seetharami R Seelam (IBM Research, US); Liana Fong (IBM T.J. Watson Research Center, US); Asser Tantawi (IBM T.J. Watson Research Center, US); John Lewars (IBM Systems and Technology Group, US); John Divirgilio (IBM, US); Kevin Gildea (IBM, US)

Oblivious Algorithms for Multicores and Network of Processors
Rezaul Chowdhury (University of Texas at Austin, US); Francesco Silvestri (University of Padova, Italy); Brandon Blakeley (University of Texas, US); Vijaya Ramachandran (University of Texas at Austin, US)

Analyzing and Adjusting User Runtime Estimates to Improve Job Scheduling on the Blue Gene/P
Wei Tang (Illinois Institute of Technology, US); Narayan Desai (Argonne National Laboratory, US), Daniel Buettner (Argonne National Laboratory, US); Zhiling Lan (Illinois Instititue of Technology, US)

Performance Evaluation of Concurrent Collections on High-Performance Multicore Computing Systems
Aparna Chandramowlishwaran (Georgia Institute of Technology, US); Kathleen Knobe (Intel, US); Richard W. Vuduc (Georgia Institute of Technology, US)

Session 16
P2P Algorithms
Chair: Amitabha Bagchi

A Hybrid Interest Management Mechanism for Peer-to-Peer Networked Virtual Environments
Ke Pan (Nanyang Technological University, Singapore); Wentong Cai (Nanyang Technological University, Singapore); Xueyan Tang (Nanyang Technological University, Singapore); Suiping Zhou (Nanyang Technological University, Singapore); Stephen John Turner (Nanyang Technological University, Singapore)

Attack-Resistant Frequency Counting
Bo Wu (University of New Mexico, US); Valerie King (University of Victoria, Canada); Jared Saia (University of New Mexico, US)

Overlays with preferences: Approximation algorithms for matching with preference lists
Giorgos Georgiadis (Chalmers University of Technology, Sweden); Marina Papatriantafilou (Chalmers University of Technology, Sweden)

Analysis of Durability in Replicated Distributed Storage Systems
Joseph Pasquale (University of California, San Diego, US); Sriram Ramabhadran (University of California, San Diego, US)

Session 17
Parallel Solutions for String and Sequence Problems
Chair: Ruppa Thulasiram

Scalable Multi-Pipeline Architecture for High Performance Multi-Pattern String Matching
Weirong Jiang (University of Southern California, US); Yi-Hua Yang (University of Southern California, US); Viktor K. Prasanna (University of Southern California, US)

Head-Body Partitioned String Matching for Deep Packet Inspection with Scalable and Attack-Resilient Performance
Yi-Hua Yang (University of Southern California, US); Viktor K. Prasanna (University of Southern California, US); Chenqian Jiang (University of Southern California, US)

Parallel de novo Assembly of Large Genomes from High-Throughput Short Reads
Benjamin G. Jackson (AOL, US); Matthew Regennitter (Iowa State University, US); Xiao Yang (Iowa State University, US); Patrick Schnable (Iowa State University, US); Srinivas Aluru (Iowa State University, US)

Efficient Parallel Algorithms for Maximum-Density Segment Problem
Xue Wang (Georgia State University, US); Fasheng Qiu (Georgia State University, US); Sushil Prasad (Georgia State University, US); Guantao Chen (Georgia State University, US)

Session 18
Energy-aware Task Management
Chair: David Bunde

Hybrid MPI/OpenMP Power-aware Computing
Dong Li (Virginia Tech, US); Bronis R. de Supinski (Lawrence Livermore National Laboratory, US); Martin Schulz (Lawrence Livermore National Laboratory, US); Kirk Cameron (Virginia Tech, US); Dimitrios S. Nikolopoulos (Foundation for Research and Technology Hellas, Greece)

Performance and Energy Optimization of Concurrent Pipelined Applications
Anne Benoit (Ecole Normale Supérieure de Lyon, FR); Paul Renaud-Goud (Ecole Normale Supérieure de Lyon, FR); Yves Robert (Ecole Normale Supérieure de Lyon, FR)

Robust Control-theoretic Thermal Balancing for Server Clusters
Yong Fu (Washington University in St. Louis, US); Chenyang Lu (Washington University in St. Louis, US); Hongan Wang (Washington University in St. Louis, US)

A Simple Thermal Model for Multi-core Processors and Its Application to Slack Allocation
Zhe Wang (University of Florida, US); Sanjay Ranka (University of Florida, US)

Session 19
Parallel Operating Systems and System Software
Chair: George Bosilca

GenerOS: An Asymmetric Operating System Kernel for Multi-core Systems
Qingbo Yuan (Institute of Compute Technology, PRC); Jianbo Zhao (Institute of Compute Technology, PRC); Mingyu Chen (Institute of Compute Technology, PRC); Ninghui Sun (Institute of Compute Technology, PRC)

Palacios and Kitten: New High Performance Operating Systems for Scalable Virtualized and Native Supercomputing
John Lange (Northwestern University, US); Kevin Pedretti (Sandia National Laboratories, US); Trammell Hudson (Sandia National Laboratories, US); Peter Dinda (Northwestern University, US); Zheng Cui (University of New Mexico, US); Lei Xia (Northwestern University, US); Patrick Bridges (University of New Mexico, US); Andy Gocke (Northwestern University, US); Steven Jaconette (Northwestern University, US); Michael Levenhagen (Sandia National Laboratories, US); and Ron Brightwell (Sandia National Laboratories, US)

MMT: Exploiting Fine-Grained Parallelism in Dynamic Memory Management
Devesh Tiwari (North Carolina State University, US); Sanghoon Lee (North Carolina State University, US); James Tuck (North Carolina State University, US); Yan Solihin (North Carolina State University, US)

Optimization of Applications with Non-blocking Neighborhood Collectives via Multisends on the Blue Gene/P Supercomputer
Sameer Kumar (IBM Research, US); Philip Heidelberger (IBM Research, USA); Dong Chen (IBM Research, US); Michael Hines (IBM Research, US)

Session 20
Parallel Graph Algorithms I
Chair: Cynthia Phillips

A Multi-Source Label-Correcting Algorithm for the All-Pairs Shortest Paths Problem
Hiroki Yanagisawa (IBM, Japan)

Parallel Computation of Best Connections in Public Transportation Networks
Daniel Delling (Microsoft Research, Germany); Bastian Katz (Karlsruhe Institute of Technology, Germany); Thomas Pajor (Universitat Karlsruhe)

Dynamically Tuned Push-Relabel Algorithm for the Maximum Flow Problem on CPU-GPU-Hybrid Platforms
Zhengyu He (Georgia Institute of Technology, US); Bo Hong (Georgia Institute of Technology, US)

A Novel Application of Parallel Betweenness Centrality to Power Grid Contingency Analysis
Shuangshuang Jin (Pacific Northwest National Laboratory, US); Zhenyu Huang (Pacific Northwest National Laboratory, US); Yousu Chen (Pacific Northwest National Laboratory, US); Daniel Gerardo Chavarria (Pacific Northwest National Laboratory, US); John Feo (Pacific Northwest National Laboratory, US); Pak Wong (Pacific Northwest National Laboratory, US)

Session 21
Parallel Linear Algebra II
Chair: Esmond Ng

Adapting Communication-Avoiding LU and QR Factorizations to Multicore Architectures
Laura Grigori (INRIA, FR); Simplice Donfack (INRIA, FR); Alok Kumar Gupta (BCCS, Norway)

QR Factorization of Tall and Skinny Matrices in a Grid Computing Environment
Emmanuel Agullo (University of Tennessee, US); Camille Coti (INRIA, Saclay-Ile de France, FR); Jack Dongarra (University of Tennessee, Knoxville, US); Thomas Herault (Universite Paris Sud (LRI), FR); Julien Langou (University of Colorado Denver, US)

Tile QR Factorization with Parallel Panel Processing for Multicore Architectures
Bilel Hadri (University of Tennessee, US); Hatem Ltaief (University of Tennessee, US); Emmanuel Agullo (University of Tennessee, US); Jack Dongarra (University of Tennessee, Knoxville, US)

Linpack Evaluation on a Supercomputer with Heterogenous Accelerators
Toshio Endo (Tokyo Institute of Technology, Japan); Akira Nukada (Tokyo Institute of Technology, Japan); Satoshi Matsuoka (Tokyo Institute of Technology, Japan); Naoya Maruyama (Tokyo Institute of Technology, Japan)

Session 22
Caches and Caching
Chair: Richard Murphy

Adapting Cache Partitioning Algorithms to Pseudo-LRU Replacement Policies
Kamil Kędzierski (Technical University of Catalonia, UPC, Spain); Miquel Moreto (Universitat Politecnica de Catalunya, Spain); Francisco Cazorla (Barcelona Supercomputing Center); Mateo Valero (Technical University of Catalonia, Spain)

Exploiting Set-Level Non-Uniformity of Capacity Demand to Enhance CMP Cooperative Caching
Dongyuan Zhan (University of Nebraska at Lincoln, US); Hong Jiang (University of Nebraska at Lincoln, US); Sharad Seth (University of Nebraska at Lincoln, US)

Masking I/O Latency using Application Level I/O Caching and Prefetching on Blue Gene System
Seetharami Seelam (IBM T.J. Watson Research Center); I-Hsin Chung (IBM T.J. Watson Research Center); John Bauer (IBM T.J. Watson Research Center); Hui-Fang Wen (IBM T.J. Watson Research Center)

Intra-Application Cache Partitioning
Sai Prashanth Muralidhara (The Pennsylvania State University, US); Mahmut Taylan Kandemir (The Pennsylvania State University, US); Padma Raghavan (The Pennsylvania State University, US)

Session 23
Thread Scheduling
Chair: Guang Gao

SLAW: a Scalable Locality-aware Adaptive Work-stealing Scheduler
Yi Guo (Rice University); Jlsheng Zhao (Rice University); Vincent Cave (Rice University); Vivek Sarkar (Rice University)

Executing Task Graphs Using Work-Stealing
Kunal Agrawal (Washington University in St. Louis, US); Charles Leiserson (Massachusetts Institute of Technology, US); Jim Sukha (Massachusetts Institute of Technology, US)

Structuring Execution of OpenMP Applications for Multicore Architectures
François Broquedis (University of Bordeaux, FR); Olivier Aumage (University of Bordeaux, FR); Brice Goglin (INRIA Bordeaux - Sud Ouest, FR); Samuel Thibault (University of Bordeaux, FR); Pierre-Andre Wacrenier (University of Bordeaux, FR); Raymond Namyst (University of Bordeaux, FR)

Oversubscription on Multicore Processors
Costin Iancu (Lawrence Berkeley National Laboratory); Steven Hofmeyr (Lawrence Berkeley National Laboratory); Yili Zheng (Lawrence Berkeley National Laboratory); Filip Blagojević (Lawrence Berkeley National Laboratory)

6:00 PM **PhD Forum**

6:30 PM **Pre-banquet Reception (Co-located with PhD Forum)**

7:30 PM **Symposium Banquet**

Keynote Session 8:30 AM - 9:30 AM
Chair: Rajmohan Rajaraman

Keynote Speech: Engineering an Algorithm for Parallel External Sorting of Massive Data Sets

Speaker: Peter Sanders
Universität Karlsruhe

Abstract: This talk describes algorithm engineering (AE) as a methodology for algorithmic research where design, analysis, implementation and experimental evaluation of algorithms form a feedback cycle driving the development of efficient algorithms. Additional important components of the methodology include realistic models, algorithm libraries, and collections of realistic benchmark instances. We use one main example throughout this talk: sorting huge data sets using many multi-core processors and disks. The described system is the current record holder for the GraySort and MinuteSort sorting benchmarks.

Morning Break 9:30 AM - 10:00 AM

Thursday 22nd April Schedule

Commercial Book Exhibits - All Day

Parallel Technical Sessions 24, 25, 26, 27, & 28
10:00 AM - 12:00 PM

Lunch 12 Noon – 1:30 PM (on your own)

Parallel Technical Sessions 29, 30, 31, & 32
1:30 PM – 3:30 PM

Afternoon Break 3:30 PM – 4:00 PM

Symposium Panel Discussion
Plenary Session
4:00 PM - 6:00 PM

NVIDIA Symposium Tutorial
PM Parallel Computing with CUDA
7:00 PM – 10:00 PM

Session 24
Distributed Algorithms
Chair: Amitabha Bagchi

A Scalable Algorithm for Maintaining Perpetual System Connectivity in Dynamic Distributed Systems
Tarun Bansal (The Ohio State University, US); Neeraj Mittal (The University of Texas at Dallas, US)

Algorithmic Mechanisms for Internet-based Master-Worker Computing with Untrusted and Selfish Workers
Antonio Fernández Anta (Universidad Rey Juan Carlos, Spain); Chryssis Georgiou (University of Cyprus, Cyprus); Miguel Mosteiro (Rutgers University, US and Universidad Rey Juan Carlos, Spain)

Stabilizing Pipelines for Streaming Applications
Andrew Berns (The University of Iowa, US); Anurag Dasgupta (The University of Iowa, US); Sukumar Ghosh (The University of Iowa, US)

A Dynamic Approach for Characterizing Collusion in Desktop Grids
Louis-Claude Canon (Nancy University); Emmanuel Jeannot (INRIA Bordeaux Sud-Ouest, FR): Jon Weissman (University of Minnesota, Twin Cities, US)

Session 25
Automatic Tuning and Automatic Parallelization
Chair: Guang Gao

Offline Library Adaptation Using Automatically Generated Heuristics
Frédéric de Mesmay (Carnegie Mellon University, US); Yevgen Voronenko (Carnegie Mellon University, US); Markus Pueschel (Carnegie Mellon University, US)

An Auto-Tuning Framework for Parallel Multicore Stencil Computations
Shoaib Kamil (Lawrence Berkeley National Laboratory, US); Cy Chan (Massachusetts Institute of Technology, US); Leonid Oliker (Lawrence Berkeley National Laboratory, US); John Shalf (Lawrence Berkeley National Laboratory, US); Samuel Williams (Lawrence Berkeley National Laboratory, US)

DynTile: Parametric Tiled Loop Generation for Parallel Execution on Multicore Processors
Albert Hartono (Ohio State University, US); Muthu Manikandan Baskaran (Ohio State University, US); J. Ram Ramanujan (Louisiana State University); Ponnuswamy Sadayappan (Ohio State University, US)

Using Focused Regression For Accurate Time-Constrained Scaling of Scientific Applications
Bradley Barnes (University of Georgia, US); Jeonifer Garren (University of Georgia, US); David Lowenthal (University of Arizona, US); Jaxk Reeves (University of Georgia, US); Bronis R. de Supinski (Lawrence Livermore National Laboratory, US); Martin Schulz (Lawrence Livermore National Laboratory, US); Barry Rountree (University of Georgia, US)

Session 26
Architectural Support for Runtime Systems
Chair: Arun Rodrigues

A Low Cost Split-Issue Technique to Improve Performance of SMT Clustered VLIW Processors
Manoj Gupta (Universitat Politècnica de Catalunya, Spain); Fermín Sánchez (Universitat Politècnica de Catalunya, Spain); Josep Llosa (Universitat Politècnica de Catalunya, Spain)

Exploiting Inter-thread Temporal Locality for Chip Multithreading
Jiayuan Meng (University of Virginia, US); Jeremy Sheaffer (NVIDIA, US); Kevin Skadron (University of Virginia, US)

Profitability-Based Power Allocation for Speculative Multithreaded Systems
Polychronis Xekalakis (University of Edinburgh, UK); Nikolas Ioannou (University of Edinburgh, UK); Salman Khan (University of Edinburgh, UK); Marcelo Cintra (University of Edinburgh, UK)

Evaluating Standard-Based Self-Virtualizing Devices: A Performance Study on 10 GbE NICs with SR-IOV Support
Jiuxing Liu (IBM T.J. Watson Research Center, US)

Session 27
Client-Server System Management and Analysis
Chair: Chen Ding

QoS Assessment of WS-BPEL Processes through non-Markovian Stochastic Petri Nets
Dario Bruneo (Universita di Messina, Italy); Salvatore Distefano (Universita di Messina, Italy); Francesco Longo (Universita di Messina, Italy); Marco Scarpa (Universita di Messina, Italy)

Power-aware Resource Provisioning in Cluster Computing
Kaiqi Xiong (North Carolina State University, US)

Using the Middle Tier to Understand Cross-Tier Delay in a Multi-tier Application
Haichuan Wang (IBM Research, PRC); Qiming Teng (IBM Research, PRC); Xiao Zhong (IBM Research, PRC); Peter Sweeney (IBM T.J. Watson Research Center, US)

Service and Resource Discovery in Cycle-Sharing Environments with a Utility Algebra
João Nuno Silva (Technical University of Lisbon, Portugal); Paulo Ferreira (Technical University of Lisbon, Portugal); Luís Veiga (Technical University of Lisbon, Portugal)

Session 28
Parallel Graph Algorithms II
Chair: Padma Raghavan

Optimization of Linked List Prefix Computations on Multithreaded GPUs Using CUDA
Zheng Wei (University of Maryland, US); Joseph Jaja (University of Maryland, College Park, US)

Parallel External Memory Graph Algorithms
Lars Arge (Aarhus University, Denmark); Michael Goodrich (University of California, Irvine, US); Nodari Sitchinava (Aarhus University, Denmark)

Engineering a Scalable High Quality Graph Partitioner
Mauel HoltGrewe (University of Karlsruhe, Germany); Peter Sanders (University of Karlsruhe, Germany); Christian Schulz (University of Karlsruhe, Germany)

Session 29
Algorithms for Wireless Networks
Chair: Neeraj Mittal

Sparse Power-Efficient Topologies for Wireless Ad Hoc Sensor Networks
Amitabha Bagchi (Indian Institute of Technology, Delhi, India)

Contention-based Georouting with Guaranteed Delivery, Minimal Communication Overhead, and Shorter Paths in Wireless Sensor Networks
Stefan Rührup (OFFIS - Institute for Information Technology, Germany); Ivan Stojmenović (University of Ottawa, Canada)

Midpoint Routing Algorithms for Delaunay Triangulations
Albert Zomaya (University of Sydney, Australia); Weisheng Si (University of Sydney, Australia)

A Local, Distributed Constant-Factor Approximation Algorithm for the Dynamic Facility Location Problem
Bastian Degener (University of Paderborn, Germany); Barbara Kempkes (University of Paderborn, Germany); Peter Pietrzyk (University of Paderborn, Germany)

Session 30
Analysis of heterogeneity and future platforms
Chair: Richard Murphy

Toward Understanding Heterogeneity in Computing
Arnold Rosenberg (Colorado State University, US); Ron Chi-Lung Chiang (Colorado State University, US)

Balls into Non-uniform Bins
Petra Berenbrink (Simon Fraser University, Canada); André Brinkmann (University of Paderborn, Germany); Tom Friedetzky (Durham University, UK); Lars Nagel (Durham University, UK)

An Introductory Exascale Feasibility Study for FFTs and Multigrid
Hormozd Gahvari (University of Illinois at Urbana-Champaign, US); William Gropp (Argonne National Laboratory, US)

Session 31
Data Management
Chair: Zhihui Du

A Cost-Effective Strategy for Intermediate Data Storage in Scientific Cloud Workflow Systems
Dong Yuan (Swinburne University of Technology, Australia); Yun Yang (Swinburne University of Technology, Australia); Xiao Liu (Swinburne University of Technology, Australia); Jinjun Chen (Swinburne University of Technology, Australia)

BlobSeer: Bringing High Throughput under Heavy Concurrency to Hadoop Map/Reduce Applications
Bogdan Nicolae (University of Rennes, FR), Diana Moise (INRIA, Rennes, FR); Gabriel Antoniu (INRIA Rennes-Bretagne, FR); Luc Bougé (IRISA/Ecole Normale Superieure Cachan Brittany, FR); Matthieu Dorier (Ecole Normale Superieure Cachan, FR)

PreDatA - Preparatory Data Analytics on Peta-Scale Machines
Fang Zheng (Georgia Institute of Technology, US); Hasan Abbasi (University of Sydney, Australia); Ciprian Docan (Rutgers University, US); Jay Lofstead (Georgia Institute of Technology, US); Qing Liu (Oak Ridge National Laboratory, US); Scott Klasky (Oak Ridge National Laboratory, US); Manish Prashar (Rutgers University, US); Norbert Podhorszki (Oak Ridge National Laboratory, US); Karsten Schwan (Georgia Institute of Technology, US); Matt Wolf (Georgia Institute of Technology, US)

Reconciling Scratch Space Consumption, Exposure, and Volatility to Achieve Timely Staging of Job Input Data
Henry Monti (Virginia Tech, US); Ali R Butt (Virginia Tech, US); Sudharshan S Vazhkudai (Oak Ridge National Laboratory, US)

Session 32
Synchronization
Chair: Chen Ding

Hierarchical Phasers for Scalable Synchronization and Reductions in Dynamic Parallelism
Jun Shirako (Rice University, US); Vivek Sarkar (Rice University, US)

Clustering JVMs with Software Transactional Memory Support
Christos Kotselidis (University of Manchester, UK); Mikel Luján (University of Manchester, UK); Behram Khan (University of Manchester, UK); Mohammad Ansari (University of Manchester, UK); Konstantinos Malakasis (University of Manchester, UK); Chris Kirkham (University of Manchester, UK); Ian Watson (University of Manchester, UK)

Inter-Block GPU Communication via Fast Barrier Synchronization
Shucai Xiao (Virginia Tech, US); Wu-chun Feng (Virginia Tech, US)

A Lock-Free, Cache-Efficient Multi-Core Synchronization Mechanism for Line-Rate Network Traffic Monitoring
Patrick Pak-Ching Lee (The Chinese University of Hong Kong, Hong Kong); Tian Bu (Bell Labs, Lucent, US); Girish Chandranmenon (Lucent Technologies, US)

Symposium Panel Discussion - Plenary
4:00 PM – 6:00 PM

Topic: Unconventional Wisdom in Multicore Computing

Panel Organizer:
Richard (Rich) Vuduc, Georgia Institute of Technology
Panelists:
Thomas Cormen, Dartmouth College
Markus Pueschel, Carnegie Mellon University
Karthikeyan (Karu) Sankaralingam, University of Wisconsin
Vivek Sarkar, Rice University
Jeffrey Vetter, Oak Ridge National Laboratory / Georgia Tech

Description:
The aim of this panel is to assess the prevailing assumptions in today's multicore software R&D and computing education, and ask in particular whether the parallel computing community's current trajectory will produce the "right" algorithmic and software infrastructure for tomorrow's applications.

All Symposium Tutorial
7:00 PM - 10:00 PM

(Open to all IPDPS 2010 attendees)

Title: Parallel Computing with CUDA

Presenter:
Michael Garland, NVIDIA

Abstract: NVIDIA's CUDA architecture provides a powerful platform for writing highly parallel programs. By providing simple abstractions for hierarchical thread organization, memories, and synchronization, the CUDA programming model allows programmers to write scalable programs without the burden of learning a multitude of new programming constructs. The CUDA architecture can support many languages and programming environments, including C, Fortran, OpenCL, and DirectX Compute. In this tutorial, I will provide an overview of modern GPU processor design and its implications for successful parallel programming models. I will present the programming model defined by the CUDA architecture, and demonstrate how this is exposed in the C/C++ language. Finally, I will sketch some techniques for implementing common data-parallel algorithms in the CUDA model.

Workshop 11: PDSEC
Workshop on Parallel and Distributed Scientific and Engineering Computing

Description:
The field of high performance computing has earned prominence through advances in electronic and integrated technologies beginning in the 1940s. Current times are very exciting and the years to come will witness a proliferation in the use of parallel and distributed systems. The scientific and engineering application domains have a key role in shaping future research and development activities in academia and industry, especially when the solution of large and complex problems must cope with tight timing schedules. This special workshop is to bring together computer scientists, applied mathematicians and researchers to present, discuss and exchange ideas, results, work in progress and experiences in the area of parallel and distributed computing for problems in science and engineering applications and inter-disciplinary applications.

General Co-Chairs
Raphael Couturier, Universite de Franche-Comte, France
Gudula Runger, Chemnitz University of Technology, Germany

Program Co-Chairs
Thomas Rauber, University of Bayreuth, Germany
Enrique Quintana-Orti, University of Jaime I, Spain

Steering Co-Chairs
Laurence T. Yang, St Francis Xavier University, Canada
Yi Pan, Georgia State University, USA

PDSEC Program

9:00-9:05
Welcome and Opening Remarks

9:05-10:05
Session 1: Applications on GPUs

Designing Scalable Many-core Parallel Algorithms for Min Graphs using CUDA
Quoc-Nam Tran

CUDA-based AES Parallelization with Fine-Tuned GPU Memory Utilization
Chonglei Mei, Hai Jiang, Jeff Jenness

Performance Study of Mapping Irregular Computations on GPUs
Steven Solomon, Parimala Thulasiraman

Simulating Anomalous Diffusion on Graphics Processing Units
Karl Heinz Hoffmann, Michael Hofmann, Jens Lang, Gudula Rünger, Steffen Seeger

10:05-10:30 Break

10:30-11:30
Session 2: Numerical Analysis and Computations

Solving large sparse linear systems in a grid environment using Java
Raphaël Couturier, Fabienne Jézéquel

Issues in Adaptive Mesh Refinement
William Dai

Solving the advection PDE on the Cell Broadband Engine
Georgios Rokos, Gerassimos Peteinatos, Georgia Kouveli, Georgios Goumas, Kornilios Kourtis, Nectarios Koziris

Storage Space Reduction for the Solution of Systems of Ordinary Differential Equations by Pipelining and Overlapping of Vectors
Matthias Korch, Thomas Rauber

11:30-12:30
Session 3: Application Scheduling and Load Balancing

Enhancing Adaptive Middleware for Quantum Chemistry Applications with a Database Framework
Lakshminarasimhan Seshagiri, Meng-Shiou Wu, Masha Sosonkina, Zhao Zhang, Mark Gordon, Michael W. Schmidt

Scheduling instructions on hierarchical machines
Florent Blachot, Guillaume Huard, Johnatan Pecero, Erik Saule, Denis Trystram

Mapping Asynchronous Iterative Applications on Heterogeneous Distributed Architectures
Raphaél Couturier, David Laiymani, Sébastien Miquée

Investigating the robustness of adaptive dynamic loop scheduling on heterogeneous computing systems
Srishti Srivastava, Ioana Banicescu, Florina Ciorba

12:30-13:30 Lunch Break (on your own)

13:30-15:00
Session 4: Performance Optimization Techniques

Prototype for a Large-Scale Static Timing Analyzer Running on an IBM Blue Gene
Akintayo Holder, Christopher Carothers, Kerim Kalafala

Performance Prediction of Weather Forecasting Software on Multicore Systems
Javier Delgado, Masoud Sadjadi, Marlon Bright, Malek Adjouadi, Hector Duran

Restructuring Parallel Loops to Curb False Sharing on Multicore Architectures
Santash Sarangkar, Apan Qasem

Parallel Task for parallelizing object-oriented desktop applications
Nasser Giacaman, Oliver Sinnen

Application Tuning through Bottleneck-driven Refactoring
Guogjing Cong, I-Hsin Chung, Huifang Wen, David Klepacki, Hiroki Murata, Yasushi Negishi and Takao Moriyama

The Pilot Approach to Cluster Programming in C
John Carter, William Gardner, Gary Grewal

15:00-15:30 Break

15:30-16:30
Session 5: Distributed Systems and Applications

A Framework for FPGA Functional Units in High Performance Computing
Andreas Koltes, John O'Donnell

FG-MPI: Fine-grain MPI for Multicore and Clusters
Humaira Kamal, Alan Wagner

Processor Affinity and MPI Performance on SMP-CMP Clusters
Chi Zhang, Xin Yuan, Ashok Srinivasan

The Resource Locating Strategy Based on Sub-domain Hybrid P2P Network Model
Yuhua Liu, Yuling Li, Laurence T. Yang, Naixue Xiong, Longquan Zhu, Kaihua Xu

16:30-17:15
Session 6: PDCoF
Special Session for Workshop on Parallel and Distributed Computing in Finance

Parallelizing a Black-Scholes Solver based on Finite Elements and Sparse Grids
Hans-Joachim Bungartz, Alexander Heinecke, Dirk Pflüger, Stefanie Schraufstetter

Pricing of Cross-Currency Interest Rate Drivatives on Graphics Processing Units
Duy Minh Dang

A Parallel Particle Swarm Optimization Algorithm for Opting Pricing
Hari Prasain, Girish K. Jha, Parimala Thulasiraman and Ruppa K. Thulasiram

17:15-17:20 Closing

Workshop 12: PMEO
Performance Modeling, Evaluation, and Optimisation of Ubiquitous Computing and Networked Systems

Description:
The performance modeling, evaluation, and optimization of ubiquitous computing and networked systems have been an important research topic over the past years and poses challenging problems that require new tools and methods to keep up with the rapid evolution and increasing complexity of such systems. This workshop will bring together scientists, engineers, practitioners, and computer users to share and exchange their experiences, discuss challenges, and report state-of-the-art and in-progress research on all aspects of performance modeling, evaluation, and optimization of ubiquitous computing and networked systems.

Workshop Co-Chairs
Geyong Min, University of Bradford, U.K.
Mohamed Ould-Khaoua, Sultan Qaboos University, Oman

Program Co-Chairs
Xiaolong Jin, University of Bradford, U.K.
Ahmed Y. Al-Dubai , Edinburgh Napier University, U.K.

PMEO Program

8:30-8:40
Welcome and Workshop Introduction

8:40-9:30
Keynote Speech
(Chair: Liang Zhao, Edinburgh Napier University, UK)

Caching from the Internet to the Cloud Professor
Keqiu Li (Dalian University of Technology, China)

9:30-10:30
Session 1
(Chair: Keqin Li, State University of New York at New Paltz, USA)

Multicore-Aware Reuse Distance Analysis
Derek L. Schuff, Benjamin S. Parsons and Vijay S. Pai (Purdue University, USA)

Clairvoyant Site Allocation of Jobs with Highly Variable Service Demands in a Computational Grid
Stylianos Zikos and Helen D. Karatza (Aristotle University of Thessaloniki, Greece)

Resource Management of Enterprise Cloud Systems Using Layered Queuing and Historical Performance Models
David A. Bacigalupo (University of Southampton, UK), Jano van Hemert (University of Edinburgh, UK), Asif Usmani (University of Edinburgh, UK), Donna Dillenberger (IBM T.J. Watson Research Centre, USA), Gary B. Wills (University of Southampton, UK) and Stephen A. Jarvis (University of Warwick, UK)

10:30-10:50 Break

10:50-12:30
Session 2
(Chair: Helen D. Karatza, Aristotle University of Thessaloniki, Greece)

Cross Layer Neighbourhood Load Routing for Wireless Mesh Networks
Liang Zhao (Edinburgh Napier University, UK), Ahmed Y. Al-Dubai (Edinburgh Napier University, UK) and Geyong Min (University of Bradford, UK)

Adapting to NAT timeout values in P2P Overlay Networks
Richard Price and Peter Tino (University of Birmingham, UK)

Power Assignment and Transmission Scheduling in Wireless Networks
Keqin Li (State University of New York at New Paltz, USA)

Agent Placement in Wireless Embedded Systems: Memory Space and Energy Optimizations
Nikos Tziritas (Center for Research and Technology, Thessaly, Greece; University of Thessaly, Greece), Thanasis Loukopoulos (Center for Research and Technology, Thessaly, Greece; Technological & Educational Institute of Lamia, Greece), Spyros Lalis (Center for Research and Technology, Thessaly, Greece; University of Thessaly, Greece) and Petros Lampsas (Center for Research and Technology, Thessaly, Greece; Technological & Educational Institute of Lamia, Greece)

A Supplying Partner Strategy for Mobile Networks-based 3D Streaming - A Proof of Concept
Haifa Raia Maamar, Richard W. Pazzi, Azzedine Boukerche and Emil Petriu (SITE-University of Ottawa, Canada)

12:30-14:00 Lunch

14:00-15:20
Session 3
(Chair: Liang Zhao, Edinburgh Napier University, UK)

A stochastic framework to depict viral propagation in wireless heterogeneous networks
Hoai-Nam Nguyen, Yasuhiro Ohara, and Yoichi Shinoda (Japan Advanced Institute of Science and Technology, Japan)

An Adaptive I/O Load Distribution Scheme for Distributed Systems
Xin Chen (Tennessee Technological University, USA), Jeremy Langston (Tennessee Technological University, USA), Xubin He (Tennessee Technological University, USA) and Fengjiang Mao (Shenzhen Polytechnic, China)

A New Probabilistic Linear Exponential Backoff Scheme for MANETs
Muneer Bani Yassein (Jordan University of Science and Technology, Jordan) Saher Manaseer (University of Glasgow, UK), Asmahan Abu Al-hassan (Jordan University of Science and Technology, Jordan), Zeinab Abu Taye (Jordan University of Science and Technology, Jordan) and Ahmed Y. Al-Dubai (Edinburgh Napier University, UK)

A Design Aid and Real-Time Measurement Framework for Virtual Collaborative Simulation Environment
Ming Zhang, Hengheng Xie and Azzedine Boukerche (University of Ottawa, Canada)
15:20-15:40 Break

15:40-17:20
Session 4
(Chair: Liang Zhao, Edinburgh Napier University, UK)

Analytical Performance Comparison of 2D Mesh, WK-Recursive, and Spidergon NoCs
M. Bakhouya (UTBM, France), S. Suboh (GWU, USA), J. Gaber (UTBM, France) and T. El-Ghazawi (GWU, USA)

A Markov Chain Based Method for NoC End-to-End Latency Evaluation
Sahar Foroutan (ST-Microelectronics, USA; CEA-Leti, USA), Yvain Thonnart (CEA-Leti, USA), Richard Hersemeule (ST-Microelectronics, USA) and Ahmed Jerraya (CEA-Leti, USA)

Performance Impact of SMP-Cluster on the On-Chip Large-scale Parallel Computing Architecture
Shenggang Chen, Shuming Chen and Yaming Yin (National University of Defense Technology, China)

Parallel Isolation-Aggregation Algorithms to Solve Markov Chains Problems with Application to Page Ranking
Abderezak Touzene (Sultan Qaboos University, Oman)

Predictibility of Inter-component latency in a Software Communications Architecture Operating Environment
Gael Abgrall (UEB, ENSIETA/DTN, France), Frédéric Le Roy (UEB, ENSIETA/DTN, France), Jean-Philippe Diguety (UEB, UBS/Lab-STICC, France), Guy Gogniat (UEB, UBS/Lab-STICC, France) and Jean-Philippe Delahaye (DGA/CELAR, France)

Please note that all presentations are scheduled for 20 minutes, including presentation and discussion.

Workshop 13: DPDNS
Dependable Parallel, Distributed and Network-Centric Systems

Description:
Increasingly large and complex parallel, distributed and network-centric computing systems provide unique challenges to the researchers in dependable computing, especially because of the high failure rates intrinsic to these systems. The goal of this workshop in continuation of the FTPDS (Fault-Tolerant Parallel and Distributed Systems) workshop series is to provide a forum for researchers and practitioners to discuss all aspects of dependability including reliability, availability, safety and security for parallel, distributed and network-centric systems.
All aspects of design, theory and realization are of interest.

Program Chair
Barry Johnson, University of Virginia, USA

Program Co-Chair
Carl Elks, University of Virginia, USA

DPDNS Program

8:15 am Opening Remarks

8:30 System Dependability and Fault Tolerance

8:30
Failure Prediction for Autonomic Management of Networked Computer Systems with Availability Assurance
Ziming Zhang and Song Fu, New Mexico Tech University.
9:00
J2EE Instrumentation for software aging root cause application component determination with AspectJ
Javier Alonso, Jordi Torres, Josep L. Berral and Ricard Gavaldà, Technical University of Catalonia.

9:30-10:00 Break.

10:00
Improving MapReduce Fault Tolerance in the Cloud
Qin Zheng, Institute of High Performance Computing.

10:30 Distributed Systems

10:30
Tackling Consistency Issues for Runtime Updating Distributed Systems
Filippo Bannò, Daniele Marletta, Giuseppe Pappalardo and Emiliano Tramontana, University of Catania.

11:00
Achieving Information Dependability in Grids through GDS²
Vincenzo Daniele Cunsolo, Salvatore Distefano, Antonio Puliafito and Marco Scarpa, University of Messina.
11:30
Evaluating Database-oriented Replication Schemes in Software Transactional Memory Systems
Roberto Palmieri, Francesco Quaglia, Paolo Romano and Nuno Carvalho, Sapienza Universitá di Roma.

12:00 Keynote Speech
Recent Results in Checkpointing and Failure Recovery in Distributed Systems and Wireless Networks
Mukesh Singhal, Department of Computer Science, University of Kentucky.

1:00 pm Lunch Break

2:00 Reliable Storage
Optimizing RAID for Long Term Data Archives
Henning Klein and Joerg Keller, Fujitsu Technology Solutions.

2:30 Networks

Experimental Responsiveness Evaluation of Decentralized Service Discovery
Andreas Dittrich and Felix Salfner, Humboldt-Universität zu Berlin.
3:00
Analysis of Network Topologies and Fault-Tolerant Routing Algorithms using Binary Decision Diagrams
Andreas Döring, IBM Research - Zurich.

3:30-4:00 Break

4:00
Incentive Mechanisms in Peer-to-Peer Networks
Pedro da Silva Rodrigues, Carlos Ribeiro and Luis Veiga, Technical University of Lisbon.

4:30
Lessons Learned During the Implementation of the BVR Wireless Sensor Network Protocol on SunSPOTs
Ralph Robert Erdt and Martin Gergeleit, Hochschule RheinMain.

5:00 pm Concluding remarks

Workshop 14: HOTP2P International Workshop on Hot Topics in Peer-to-Peer Systems

Description:
Peer-to-Peer (P2P) systems are decentralized, self-organizing distributed systems that cooperate to exchange data. These systems have emerged as the dominant consumer of residential Internet subscribers' bandwidth, and are being increasingly used in many different application domains. In the last few years, research on P2P systems has been quite intensive, and has produced remarkable results in scalability, robustness, location, distributed storage, and system measurements. Consequently, P2P systems continue to evolve, differentiating today's state-of-the-art from earlier instantiations such as Napster, KaZaA, Gnutella, and Morpheus. The International Workshop on Hot Topics in Peer-to-Peer Systems (Hot-P2P), aims to bring together researchers and practitioners, from both industry and academia, in the fields of systems, networking, and theory, and to represent an occasion to share latest research results and ideas on P2P systems, thereby promoting research activities in this area.

Program Chair
Julien Bourgeois, LIFC, University of Franche-Comte (France)

HOTP2P Program

08:30 - Workshop opening

08:45 - **Session 1: P2P Multimedia Applications**
Estimating Operating Conditions in a Peer-to-Peer Session Initiation Protocol Overlay Network
Jouni Mäenpää, Gonzalo Camarillo
09:10
Adaptive Server Allocation for Peer-assisted Video-on-Demand
Konstantin Pussep, Osama Abboud, Thorsten Strufe, Ralf Steinmetz
09:35
Heterogeneity in Data-Driven Live Streaming: Blessing or Curse?
Fabien Mathieu

10:00 - Break

10:30 – **Session 2: P2P Multimedia applications**
Techniques for Low-latency proxy selection in Wide-area P2P networks
Arijit Ganguly, Oscar Boykin, Renato Figueiredo
10:55
Mobile-Friendly Peer-to-Peer Client Routing Using Out-of-Band Signaling
Wei Wu, Jim Womack, Xinhua Ling
11:20
Deetoo: Scalable Unstructured Search Built on a Structured Overlay
Tae Woong Choi, Oscar Boykin
11:45
Using query transformation to improve Gnutella search performance
Surendar Chandra, William F. Acosta
12:10
Analysis of Random Time-Based Switching for File Sharing in Peer-to-Peer Networks
Keqin Li

12:35 - Lunch

13:35 - Invited Talk
Challenges and results in secured fleets of P2P mobile and autonomous communicating systems
Serge Chaumette

14:10 - **Session 3: P2P Optimization II**
Tagging with DHARMA
Luca M Aiello, Marco Milanesio, Giancarlo Ruffo, Rossano Schifanella
14:35
Degree Hunter: on the Impact of Balancing Node Degrees in de Bruijn-Based Overlay Networks
Pierre Fraigniaud, Hoang-Anh Phan
15:00
BitTorrent, fountain codes: friends or foes?
Salvatore Spoto, Rossano Gaeta, Marco Grangetto, Matteo Sereno

15:25 - Break

15:55 - **Session 4: P2P Applications & Security**
Modeling and Analyzing the Effects of Firewalls, NATs in P2P Swarming Systems
Lucia D'Acunto, Michel Meulpolder, Rameez Rahman, Johan A. Pouwelse, Henk J. Sips
16:20
Efficient DHT attack mitigation through peers' ID distribution
Thibault Cholez, Isabelle Chrisment, Olivier Festor
17:10
High Performance Peer-to-Peer Distributed Computing with Application to Obstacle Problem
The Tung Nguyen, Didier El Baz, Pierre Spitéri, Guillaume Jourjon, Ming Chau High

17:35 - Workshop closing

Workshop 15: MTAAP
Workshop on Multi-Threaded Architectures and Applications

Description:
Multithreading (MT) programming and execution models are starting to permeate the high-end and mainstream computing scene. This trend is driven by the need to increase processor utilization and deal with the memory-processor speed gap. Recent and upcoming examples architectures that fit this profile are Cray's XMT, IBM Cyclops, and several SMT processors from Sun (Victoria Falls), IBM (Power6, Power7), or Intel, as well as heterogeneous systems with accelerators such as GPGPUs from ATI, NVIDIA, and Intel. The underlying rationale to increase processor utilization is a varying mix of new metrics that take performance improvements as well as better power and cost budgeting into account. Yet, it remains a challenge to identify and productively program applications for these architectures with a resulting substantial performance improvement. This workshop intends to identify applications that are amenable to MT and the MT programming and execution models as well as the underlying architectures on which they can thrive. The workshop seeks to explore programming frameworks in the form of languages and libraries, compilers, analysis and debugging tools to increase the programming productivity.

Chairs
Luiz DeRose (Cray)
Jeffrey Vetter (ORNL and Georgia Tech)

MTAAP Program

08:00 - 08:15 : MTAAP 2010 Welcome

08:15 - 10:15 : Systems and Tools

Modeling Bounds on Migration Overhead for a Traveling Thread Architecture
Patrick La Fratta and Peter Kogge (University of Notre Dame, USA)

TiNy Threads on BlueGene/P: Exploring Many-Core Parallelisms Beyond The Traditional OS
Handong Ye, Robert S Pavel, Aaron Landwehr, Guang Gao (University of Delaware, USA)

Scheduling complex streaming applications on the Cell processor
Matthieu Gallet (Ecole normale superieure de Lyon, France), Mathias Jacquelin (ENS Lyon, France), Loris Marchal (CNRS, France)

User Level DB: a Debugging API for User-Level Thread Libraries
Kevin Pouget, Marc Pérache, Patrick Carribault, Hervé Jourdren (CEA/DAM Ile de France, France)

10:15 - 10:45 : Break I

10:45 - 11:45 : Keynote
The POWER7 Processor
Dr. Ram Rajamony, IBM Austin Research Laboratory

11:45 - 13:00 : Lunch

13:00 - 15:00 : Applications

A Multi-Threaded Approach for Data-Flow Analysis
Marcus Edvinsson and Welf Lowe (Vaxjo University, Sweden)

Experimental Comparison of Emulated Lock-free vs. Fine-grain Locked Data Structures on the Cray XMT
Rob Farber (Pacific Northwest Laboratory, USA) and David W. Mizell (Cray, USA)

Large Scale Complex Network Analysis using the Hybrid Combination of a MapReduce cluster and a Highly Multithreaded System
Seunghwa Kang, David A. Bader (Georgia Institute of Technology, USA)

On the Parallelisation of of MCMC by Speculative Chain Execution
Jonathan M R Byrd, Stephen Jarvis and Abhir H. Bhalerao (University of Warwick, United Kingdom)

15:00 - 15:30 : Break II

15:30 - 17:00 : Programming Framework

Out-of-Core Distribution Sort in the FG Programming Environment
Priya Natarajan (Dartmouth College, USA); Thomas H. Cormen (Dartmouth College, USA); Elena Riccio Strange (A9.com, USA)

Massive Streaming Data Analytics: A Case Study with Clustering Coefficients
David Ediger, Karl Jiang, Jason Riedy, David A. Bader (Georgia Institute of Technology, USA)

Hashing Strategies for the Cray XMT
Eric Goodman (Sandia National Laboratories, USA); David J Haglin, Chad Scherrer, Daniel Gerardo Chavarria, Jace Mogill, John Feo (Pacific Northwest National Laboratory, USA)

Adjourn

Workshop 16: PDCoF Workshop on Parallel and Distributed Computing in Finance

Description:
In the last decade Computational Finance (CF) has influenced the market place extensively with enormous impact on wealth building, employment opportunities, and tremendous economic growth. This field forms an ever-expanding part of the financial sector, in numerous ways today. The time is, therefore, ripe now to bring together researchers in the areas of finance using complicated financial models to solve computationally intensive problems and computer scientists having the resources and solution methodologies to solve such problems. The main goal of this workshop is to provide a timely forum for these two groups to exchange and disseminate new ideas, techniques, and research in computational finance. Strong discussions will follow the presentations and the experience could lead them into the formulation, implementation of the models used by the practitioners in financial sector. This workshop will be a first such effort in bringing together researchers in the areas of finance and advanced computing (i) who develop and employ parallel and distributed computing extensively (ii) at a venue where parallel, distributed, high performance computing is the fundamental thread of discussions and arguments.

General Chair:
Ruppa K. Thulasiram,
University of Manitoba, Canada

Program Chairs:
Alan J. King, IBM TJ Watson, USA
Stephane Vialle, Supelec, France

PDCoF Program

Note that the papers accepted for the PDCoF workshop, listed below, will be presented as **a special session in the PDSEC workshop**, listed earlier in the program.

16:30-17:15
PDSEC Workshop - Session 6:
Special Session for Workshop on Parallel and Distributed Computing in Finance

Parallelizing a Black-Scholes Solver based on Finite Elements and Sparse Grids
Hans-Joachim Bungartz, Alexander Heinecke, Dirk Pfluger, Stefanie Schraufstetter

Pricing of Cross-Currency Interest Rate Drivatives on Graphics Processing Units
Duy Minh Dang

A Parallel Particle Swarm Optimization Algorithm for Option Pricing
Hari Prasain, Girish K. Jha, Parimala Thulasiraman and Ruppa K. Thulasiram

Workshop 17: LSPP
Workshop on Large-Scale Parallel Processing

Description:
The workshop on Large-Scale Parallel Processing is a forum that focuses on computer systems that utilize thousands of processors and beyond. This is a very active area given the goals of many researchers world-wide to enhance science-by-simulation through installing large-scale multi-petaflop systems at the start of the next decade. Large-scale systems, referred to by some as extreme-scale and Ultra-scale, have many important research aspects that need detailed examination in order for their effective design, deployment, and utilization to take place. These include handling the substantial increase in multi-core on a chip, the ensuing interconnection hierarchy, communication, and synchronization mechanisms. The workshop aims to bring together researchers from different communities working on challenging problems in this area for a dynamic exchange of ideas.

Workshop Co-chairs
Darren J. Kerbyson Los Alamos National Laboratory
Ram Rajamony IBM Austin Research Lab
Charles Weems University of Massachusetts

Additional Steering Committee Members
Johnnie Baker Kent State University
Alex Jones University of Pittsburgh
H.J. Siegel Colorado State University

LSPP Program

Invited Keynote Presentation
8:30
Blue Waters: An Extraordinary Computing Resource for Advanced Science & Engineering
Thom H. Dunning, Jr. National Center for Supercomputing Applications, Illinois.

Session 1: Computing with GPUs
9:30
Efficient Lists Intersection by CPU-GPU Cooperative Computing
D. Wu, F. Zhang, N. Ao, G. Wang, X. Liu, Jing Liu, Nankai University

10:00 Break

Session 1: Computing with GPUs (continued)
10:30
High Precision Integer Multiplication with a Graphics Processing Unit
N. Emmart, C. Weems, University of Massachusetts
11:00
Large Neighborhood Local Search Optimization on Graphics Processing Units
T. Van Luong, N. Melab, E. Talbi, INRIA Lille Nord Europe, CNRS/LIFL University of Lille
11:30
A Fast GPU Algorithm for Graph Connectivity
J. Soman, K. Kothapalli, P.J. Narayanan, Int. Institute of Information Technology Hyderabad

12:00 Lunch Break

Session 2: Node Level Issues
1:30
An Efficient Associative Processor Solution to an Air Traffic Control Problem
M. Yuan, J.W. Baker, L. Neiman, F. Drews, W. Meilander, Kent State University
2:00
Analyzing the Trade-off between Multiple Memory Controllers and Memory Channels on Multi-core Processor Performance
J.C. Sancho, M. Lang, D.J. Kerbyson, Los Alamos
2:30
Multicore-aware parallel temporal blocking of stencil codes for shared and distributed memory
M. Wittman, G. Hager, G. Wellein, University of Erlangen

3:00 Break

Session 3: Large-scale Systems
3:30
Scalable Parallel I/O Alternatives for Massively Parallel Partitioned Solver Systems
J. Fu, N. Liu, O. Sahni, K.E. Jansen, M.S. Shephard, C.D. Carothers, Rensselaer Polytechnic Institute
4:00
Performance analysis of Sweep3D on Blue Gene/P with the Scalasca toolset
B.J.N. Wylie, D. Boehme, B. Mohr, Z. Szebenyi, F. Wolf, Jülich Supercomputing Centre
4:30
To Upgrade or not to Upgrade? Catamount vs. Cray Linux Environment
S.D. Hammond, G.R. Mudalige, J.A. Smith, J.A. Davis, S.A. Jarvis, J. Holt, I. Miller, A. Herdman, A. Vadgama, University of Warwick, Atomic Weapons Establishment, UK

5:00 Closing Remarks

Workshop 18: JSSPP
Workshop on Job Scheduling Strategies for Parallel Processing

Description:
This annual series of workshops has been held since 1995 in conjunction with various other conferences, including IPDPS (formerly IPPS), the main parallel processing symposium sponsored by the IEEE, the ACM SIGMETRICS and ICS conferences, and HPDC/GGF, the premier annual meeting of the Global Grid Forum. This workshop series is often ranked among the top venues in Computer Science in terms of impact by CiteSeerX.

Founding Co-organizers
Dror Feitelson, The Hebrew University
Larry Rudolph, MIT

Workshop Organizers
Eitan Frachtenberg, Facebook
Uwe Schwiegelshohn, Technische Universität Dortmund

JSSPP Program*

Session 1 (8:00-9:30)
Opening and Keynote
Eitan Frachtenberg

Dan Tsafrir: ***Using Inaccurate Estimates Accurately***

Session 2 (10:00-12:00) – Grid Scheduling

Resource Provisioning in SLA-based Cluster Computing
Kaiqi Xiong and Sang Suh

An Advance Reservation-based Co-Allocation Algorithm for Distributed Computers and Network Bandwidth on QoS-guaranteed Grids
Atsuko Takefusa, Hidemoto Nakada, Tomohiro Kudoh, and Yoshio Tanaka

A Greedy Double Auction Mechanism for Grid Resource Allocation
Ding Ding

Risk Aware Overbooking for Commercial Grids
Georg Birkenheuer, Andre Brinkmann, and Holger Karl

Session 3 (13:30-15:30) – Cluster Scheduling

The Gain of Resource Delegation in Distributed Computing Environments
Alexander Fölling, Christian Grimme, Joachim Lepping, and Alexander Papaspyrou

A Moldable Online Scheduling Algorithm and Its Application to Parallel Short Sequence Mapping
Erik Saule, Doruk Bozdag, and Umit V. Catalyurek

Dynamic Proportional Share Scheduling in Hadoop
Thomas Sandholm and Kevin Lai

The Importance of Complete Data Sets for Job Scheduling Simulations
Dalibor Klusacek and Hana Rudová

Session 4 (16:00-17:30) – Thread Scheduling

Hierarchical Scheduling of DAG Structured Computations on Manycore Processors with Dynamic Thread Grouping
Yinglong Xia, Viktor K. Prasanna, and James Li

Weaving Low and High QoS Workloads in Virtual Environments
Sam Verboven, Kurt Vanmechelen, and Jan Broeckhove

Proposal and Evaluation of APIs for utilizing Inter-Core Time Aggregation Scheduler
Satoshi Yamada and Shigeru Kusakabe

***JSSPP Proceedings:**
Note that the proceedings for this workshop are not included in the regular IPDPSW 2010 volume. Interim proceedings containing a collection of the papers presented will be distributed at the workshop. It is planned to also publish a post-workshop proceedings in the Springer-Verlag Lecture Notes on Computer Science series, as was done in previous years.

GENERAL CHAIR
David A. Bader, Georgia Institute of Technology, USA

GENERAL VICE CHAIR
Alan Sussman, University of Maryland, USA

PROGRAM CHAIR
Cynthia Phillips, Sandia National Laboratories, USA

PROGRAM VICE-CHAIRS
ALGORITHMS
Rajmohan Rajaraman, Northeastern University, USA
APPLICATIONS
Jeffrey Vetter, Oak Ridge National Laboratory, USA
ARCHITECTURES
Bradley C. Kuszmaul, Massachusetts Institute of Technology, USA
SOFTWARE
Andrew Lumsdaine, Indiana University, USA

STEERING CO-CHAIRS
Viktor K. Prasanna, University of Southern California, USA
George Westrom, Discovery Science Center & FSEA, USA

STEERING COMMITTEE
David A. Bader, Georgia Institute of Technology, USA
K. Mani Chandy, California Institute of Technology, USA
Allan Gottlieb, New York University, USA
Ali R. Hurson, Pennsylvania State University, USA
Joseph JaJa, University of Maryland, USA
F. Tom Leighton, MIT, USA
Alessandro Mei, Sapienza University of Rome, Italy
Cynthia Phillips, Sandia National Laboratories, USA
Viktor K. Prasanna, University of Southern California, USA
José D.P. Rolim, University of Geneva, Switzerland
Arnold L. Rosenberg, Colorado State University, USA
Sartaj Sahni, University of Florida, USA
Behrooz Shirazi, Washington State University, USA
H.J. Siegel, Colorado State University, USA
Per Stenström, Chalmers University, Sweden
Hal Sudborough, University of Texas at Dallas, USA
Alan Sussman, University of Maryland, USA
George Westrom, Discovery Science Center & FSEA, USA

WORKSHOPS CHAIR
Ümit V. Çatalyürek, Ohio State University, USA

PROCEEDINGS CO-CHAIRS
Yuanyuan Yang, State University of New York, Stony Brook, USA
Xin Yuan, Florida State University, USA

PhD FORUM GENERAL CHAIR
Manish Parashar, Rutgers University, USA

PhD FORUM PROGRAM CHAIR
Luc Bougé, ENS Cachan, France

TCPP CHAIR
Sushil K. Prasad, Georgia State University

FINANCE CHAIR
Bill Pitts, IEEE Conferences Committee

LOCAL ARRANGEMENTS CO-CHAIRS
Susamma Barua, California State Fullerton, USA
Anu Bourgeois, Georgia State University, USA

PRODUCTION CHAIR
Sally Jelinek Westrom, Electronic Design Associates, Inc., USA

PUBLICITY CO-CHAIRS
Michela Taufer, University of Delaware, USA
Gudula Rünger, Chemnitz University of Technology, Germany
Zhihui Du, Tsinghua University, China

IPDPS 2010 PROGRAM COMMITTEE

Ittai ABRAHAM (Microsoft Research Silicon Valley) USA
Kunal AGRAWAL (Washington University in St. Louis) USA
Srinivas ALURU, (Iowa State University) USA
Elliot ANSHELEVICH (Rensselaer Polytechnic Institute) USA
Antonio FERNANDEZ ANTA (Universidad Rey Juan Carlos) Spain
Amitabha BAGCHI (Indian Institute of Technology, Delhi) India
Roscoe BARTLETT (Sandia National Laboratories) USA
Sanjoy BARUAH (University of North Carolina, Chapel Hill) USA
Michael BENDER (Stony Brook University) USA
Anne BENOIT (Ecole Normale Supérieure de Lyon) France
Petra BERENBRINK (Simon Fraser University) Canada
Martin BERZINS (University of Utah) USA
Gianfranco BILARDI (University of Padova) Italy
George BIROS (Georgia Institute of Technology) USA
Christian BISCHOF (Aachen University) Germany
Arndt BODE (Technical University of Munich) Germany
George BOSILCA (The University of Tennessee) USA
Sonja BUCHEGGER (Deutsche Telekom Laboratories, Berlin), Germany
David BUNDE (Knox College) USA
Costas BUSCH (Louisiana State University) USA
Tiziana CALAMONERI (University of Rome "Sapienza") Italy
Franck CAPPELLO (INRIA) France
Francisco J. CAZORLA (Barcelona Supercomputing Center) Spain
Barbara CHAPMAN (University of Houston) USA
Arun CHAUHAN (Indiana University) USA
Edmond CHOW (D.E. Shaw) USA
Almadena CHTCHELKANOVA (National Science Foundation) USA
Marcelo CINTRA (University of Edinburgh) United Kingdom
Guojing CONG, (IBM T.J. Watson Research Center) USA
Artur CZUMAJ (University of Warwick) England
Andre DEHON (University of Pennsylvania) USA
Chen DING (University of Rochester) USA
Pedro DINIZ (University of Southern California /Information Sciences Institute - USC/ISI) USA
Jack DONGARRA (University of Tennessee and ORNL) USA
Zhihui DU (Tsinghua University) China
Rudolf EIGENMANN (Purdue University) USA
Anne ELSTER (Norwegian University of Science and Technology) Norway
Babak FALSAFI (Ecole Polytechnique Fédérale de Lausanne) Switzerland
Panagiota FATOUROU (University of Ioannina) Greece
Pascal FELBER (University of Neuchatel) Switzerland
Jeremy FINEMAN (Massachusetts Institute of Technology) USA
Robert FOWLER (Renaissance Computing Institute) USA
Guang GAO (University of Delaware) USA
Rahul GARG (IBM, T.J. Watson Research Center) USA
John GILBERT (University of California, Santa Barbara) USA
Christos GKANTSIDIS (Microsoft Research, Cambridge) UK
Eldad HABER (Emory University) USA
Jin HAI (Huazhong University of Science and Technology) China
Tim HARRIS (Microsoft) United Kingdom
Aaron HARWOOD (University of Melbourne) Australia
Jim P. HELD (Intel) USA
Michael HEROUX (Sandia National Laboratories) USA
Torsten HOEFLER (Indiana University) USA
Paul HOVLAND (Argonne National Laboratory) USA
George KARYPIS (University of Minnesota) USA
Kate KEAHEY (University of Chicago/ANL) USA
Paul KELLY (Imperial College) UK
Thilo KIELMANN (Vrije Universiteit) Netherlands
Matt KNEPLEY (Argonne National Laboratory) USA
David KONERDING (Google) USA
Goran KONJEVOD (Arizona State University) USA
Fabian KUHN (University of Lugano) Switzerland
Miroslaw KUTYLOWSKI (Wroclaw University of Technology) Poland
Nikolaos LAOUTARIS (Telefonica) Spain
Zvi LOTKER (Ben Gurion University) Israel
Kamesh MADDURI (Lawrence Berkeley National Laboratory) USA
Grzegorz MALEWICZ (Google Labs) USA
Satoshi MATSUOKA (Tokyo Institute of Technology) Japan
Neeraj MITTAL (University of Texas, Dallas) USA
Thomas MOSCIBRODA (Microsoft Research, Redmond) USA
Richard MURPHY (Sandia National Laboratories) USA
Kengo NAKAJIMA (University of Tokyo) Japan
Esmond NG (Lawrence Berkeley National Laboratory)
Gopal PANDURANGAN (Purdue University) USA
Srinivasan PARTHASARATHY (IBM Research) USA
Fabrizio PETRINI (IBM T.J. Watson Research Center) USA
Andrea PIETRACAPRINA (University of Padova) Italy
Kirk PRUHS (University of Pittsburgh) USA
Dan QUINLAN (Lawrence Livermore National Laboratory) USA
Padma RAGHAVAN (Pennsylvania State University) USA
Lawrence RAUCHWERGER (Texas A&M University) USA
Jose RENAU (University of California, Santa Cruz) USA
Arch ROBISON (Intel) USA
Arun RODRIGUES (Sandia National Laboratories) USA
Adi ROSEN (CNRS and Universite Paris Sud) France
Jared SAIA (University of New Mexico) USA
Alessandra SALA (University of California, Santa Barbara) USA
Yanos SAZEIDES (University of Cyprus) Cyprus
Christian SCHEIDELER (University of Paderborn) Germany
Thomas SCHULTHESS (Swiss National Supercomputing Centre) Switzerland
Simha SETHUMADHAVAN (Columbia University) USA
André SEZNEC (IRISA) France
John SHALF (Lawrence Berkeley National Laboratory) USA
Federico SILLA (Universidad Politécnica de Valencia) Spain
Paul SPIRAKIS (University of Patras) Greece
Michelle STROUT (Colorado State University) USA
Mike TAYLOR (Univeristy of California, San Diego) USA
Ruppa THULASIRAM (University of Manitoba) Canada
Srikanta TIRTHAPURA (Iowa State University) USA
Sivan TOLEDO (Tel Aviv University) Israel
Denis TRYSTRAM (Institut Polytechnique de Grenoble) France
Anil VULLIKANTI (Virginia Polytechic Institute) USA
Joel WEIN (Polytechnic Institute of NYU and Akamai Technologies) USA
Hsu WEN-JING (Nanyang Technological University) Singapore
Theresa WINDUS (Iowa State University) USA
Haifeng YU (National University of Singapore) Singapore
Mohamed ZAHRAN (City University of New York) USA
Lisa ZHANG (Lucent Bell Labs) USA
Ling ZHUO (Chevron) USA

IPDPS 2010 • Reviewers

The names listed here were compiled from multiple sources. Apologies to anyone whose name did not transfer correctly to this list.

- Jaume Abella
- Sachin Agarwal
- Virat Agarwal
- Bushra Ahsan
- Jesús Alastruey
- Chanaka J. Liyana Arachchige
- Sergio Arévalo
- Chen Avin
- Mani Azimi
- Jonathan Babb
- Hansang Bae
- Tongxin Bai
- Oleg Bakun
- Nick Baltas
- Robert Banz
- Dan Barcay
- Brian Barrett
- Novella Bartolini
- Ayon Basumallik
- Daniel Beck
- Julien Bernard
- Boris Bierbaum
- Aart Bik
- Punyashloka Biswal
- Michael Borokhovich
- Marin Bougeret
- Slim Bouguera
- Aurélien Bouteiller
- Marco Bressan
- Rainer Buchty
- Martin Bücker
- Aydin Buluc
- Ariel Burton
- Antal Buss
- Alfredo Buttari
- Vladimir Čakarević
- Saverio Caminiti
- Bugra Caskurlu
- John Cavazos
- Daniel Cederman
- Luis Ceze
- Ho-Leung Chan
- Binbin Chen
- Jen-Yeu Chen
- Yongwook Choi
- Jacek Cichoń
- Kypros Constantinides
- Gennaro Cordasco
- Tom Crick
- Ruben Cuevas
- Atish Das Sarma
- Chirag Dave
- Ed D'Azevedo
- Francesco De Cassai
- Cees de Laat
- Pedro de las Heras Quirós
- John Demme
- Mathijs den Burger
- Pedro Diaz
- Stephanie Dinkins
- Shlomi Dolev
- Alastair Donaldson
- David Donofrio
- Niels Drost
- Marián Dvorský
- Nicholas Edmonds
- Brett Estrade
- Carlo Fantozzi
- Chris Fensch
- Daniel Ford
- Andrew Friedley
- Michael Frumkin
- Emanuele Fusco
- Saturnino Garcia
- Vijay K. Garg
- Maciej Gebala
- Alexandros Gerbessiotis
- Roberto Gioiosa
- Luis Goes
- Brice Goglin
- Zbigniew Gołebiewski
- Nathan Goulding
- Douglas Gregor
- Laura Grigori
- Xiaoming Gu
- Anshuman Gupta
- Thomas Hacker
- Derin Harmanci
- Rami Hasanov
- Ameya Hate
- David Hay
- Xubin He
- Scott Hemmert
- Thomas Herault
- Ezra Hoch
- Oliver Hohlfeld
- Matthew Holiday
- Atsushi Hori
- Lee Howes
- Michael Huang
- George Iosifidis
- Christian Iwainsky
- Keith Jackson
- Arpith Jacob
- Rohit Jain
- Jaakko Jarvi
- Donghwan Jeon
- Victor Jimenez
- Will Jones
- Fahed Jubair
- Prabhanjan Kambadur
- Yoshikazu Kamoshida
- Yossi Kanizo
- Nick Karonis
- Shiva Kasiwiswanathan
- Kamil Kedzierski
- Maleq Khan
- Salman Khan
- Rohit Khandekar
- Arzad Kherani
- Marcin Kik
- Marios Kleanthous
- Sebastian Kniesburges
- Will Knottenbelt
- Steven Ko
- Przemysław Kobylański
- Bjoern Kolbeck
- Athanasios Konstantinidis
- Nitish Korula
- Mirosław Korzeniowski
- Eleftherios Kosmas
- Michał Koza
- Bill Kramer
- Christopher Krieger
- Konstantinos Krikellas
- Chandra Krintz
- Peter Kropf
- Tomohiro Kudoh
- Santosh Kulkarni
- Miroslaw Kutylowski
- Okwan Kwon
- James LaGrone
- Alan LaMielle
- Joe Lancaster
- Jakub Lemiesz
- Tamir Levy
- Yang Liao
- Shih-Lien Lu
- Sangik Lee
- Seyong Lee
- Sven Leyffer
- Min Yeol Lim
- Andrzej Lingas
- Anton Lokhmotov
- Mikel Lujan
- Jason Maassen
- Krzysztof Majcher
- Anirban Mandal
- Loris Marchal
- Graham Markall
- Gaia Maselli
- Annalisa Massini
- Matt Masten
- Petar Maymounkov
- Michael Mesnier
- Ulrich Meyer
- Alessia Milani
- Antonio Sanz Montemayor
- Jose E Moreira
- Miquel Moretó
- Miguel A. Mosteiro
- Dheya Mustafa
- Ram Nanjegowda
- Danupon Nanongkai
- Wahid Nasri
- Aravind Natarajan
- Dushyanth Narayanan
- Yanik Ngoko
- Jalani Nelson
- Chrysostomos Nicopoulos
- Dimitris Nikolopoulos
- Ramon Novales
- Dan Nussbaum
- Cosmin Oancea
- Satoshi Ohshima
- Konstantinos Oikonomou
- Stephen Olivier
- Svetlana Olonetsky
- Shrideep Pallickara
- Saurav Pandit
- Marco Paolieri
- Ioannis Papadopoulos
- Manish Parashar
- Dongkook Park
- Davide Pasetto
- Jonathan Pecero-Sanchez
- Maikel Pennings
- Alberto Pettarin
- Marlon Pierce
- Natascia Piroso
- Imran Pirwani
- Jelena Pjesivac-Grbovic
- Beth Plale
- James Plank
- Allan Porterfield
- Polyvios Pratikakis
- Geppino Pucci
- Eduardo Quiñones
- Jean-Noel Quintin
- Petar Radojković
- Sanjay Rajopadhye
- Charles Reiss
- Michał Ren
- Vinay Ribeiro
- Scott Ricketts
- Etienne Riviere
- Yves Robert
- Luis Rodero-Merino
- German Rodríguez
- Jon Roelofs
- Adel Safi
- Ramendra Sahoo
- Alfred Samman
- Peter Sanders
- Oliverio J. Santana
- Agustín Santos
- Abhinav Sarje
- Thomas Sauerwald
- Rahul Savani
- Lee H. Savoie
- Fabio Schifano
- Stefan Schmid
- Fabian Schneider
- Michele Scquizzato
- Cigdem Sengul
- Jie Shao
- Gokarna Sharma
- Francesco Silvestri
- Riccardo Silvestri
- Simone Silvestri
- Vasu Singh
- Asia Slowinska
- Georgios Smaragdakis
- Yannis Smaragdakis
- Tim Smith
- Kyle Spafford
- Mike Spear
- Srivathsan Srinivas
- Alexandros Stamatakis
- Julinda Stefa
- Andrea Sterbini
- Andrew Stone
- Hari Subramoni
- Jim Sukha
- Hongyang Sun
- Sayantan Sur
- Gabriel Tanase
- Jijun Tang
- Marc Tchiboukdjian
- Christian Terboven
- Nathan Thomas
- Christopher Thraves
- Sunil Thulasidasan
- Parimala Thulasiraman
- Alexander Tiskin
- Amitabh Trehan
- Gopi Tummala
- Bora Ucar
- Fabio Vandin
- Rob van Nieuwpoort
- Kees van Reeuwijk
- Thanes Vela
- Sravanthi Kota Venkata
- Ganesh Venkatesh
- Francesco Versaci
- Frederic Vivien
- Berthold Vöcking
- Vinod Vokkarane
- Zhenlin Wang
- Mike Waychison
- Yaniv Weizman
- Tong Wen
- Adam Wierman
- Jeremiah Willcock
- Nicholas Wright
- Yuqing Wu
- Polychronis Xekalakis
- Xiaoya Xiang
- Bojian Xu
- Hao Yang
- Xiao Yang
- Shuai Ye
- Qing Yi
- Andy Yoo
- Weikuan Yu
- Marcin Zawada
- Yanyan Zeng
- Cong Zhang
- Fa Zhang
- Yanyong Zhang
- Shaojun Zhao
- Paweł Zieliński
- Jaroslaw Zola

NOTES