

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

New York University, NY, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Jack Dongarra Kaj Madsen
Jerzy Wasniewski (Eds.)

Applied Parallel Computing

State of the Art in Scientific Computing

7th International Workshop, PARA 2004
Lyngby, Denmark, June 20-23, 2004
Revised Selected Papers

Volume Editors

Jack Dongarra
University of Tennessee
Department of Computer Science
1122 Volunteer Blvd.
Knoxville, TN 37996-3450, USA
and
Oak Ridge National Laboratory
Computer Science and Mathematics Division
E-mail: dongarra@cs.utk.edu

Kaj Madsen
Jerzy Wasniewski
Technical University of Denmark
Informatics and Mathematical Modelling
Richard Petersens Plads, Building 321
2800 Kongens Lyngby, Denmark
E-mail: {km,jw}@imm.dtu.dk

Library of Congress Control Number: 2006920921

CR Subject Classification (1998): G.1-4, F.1-2, D.1-3, J.1

LNCS Sublibrary: SL 1 – Theoretical Computer Science and General Issues

ISSN	0302-9743
ISBN-10	3-540-29067-2 Springer Berlin Heidelberg New York
ISBN-13	978-3-540-29067-4 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Markus Richter, Heidelberg
Printed on acid-free paper SPIN: 11558958 06/3142 5 4 3 2 1 0

Preface

Introduction

The PARA workshops in the past were devoted to parallel computing methods in science and technology. There have been seven PARA meetings to date: PARA'94, PARA'95 and PARA'96 in Lyngby, Denmark, PARA'98 in Umeå, Sweden, PARA 2000 in Bergen, Norway, PARA 2002 in Espoo, Finland, and PARA 2004 again in Lyngby, Denmark. The first six meetings featured lectures in modern numerical algorithms, computer science, engineering, and industrial applications, all in the context of scientific parallel computing.

This meeting in the series, the PARA 2004 Workshop with the title “State of the Art in Scientific Computing”, was held in Lyngby, Denmark, June 20–23, 2004. The PARA 2004 Workshop was organized by Jack Dongarra from the University of Tennessee and Oak Ridge National Laboratory, and Kaj Madsen and Jerzy Waśniewski from the Technical University of Denmark. The emphasis here was shifted to high-performance computing (HPC). The ongoing development of ever more advanced computers provides the potential for solving increasingly difficult computational problems. However, given the complexity of modern computer architectures, the task of realizing this potential needs careful attention. For example, the failure to exploit a computer's memory hierarchy can degrade performance badly. A main concern of HPC is the development of software that optimizes the performance of a given computer.

The high cost of state-of-the-art computers can be prohibitive for many workplaces, especially if there is only an occasional need for HPC. A solution to this problem can be network computing, where remote computing facilities are exploited via the Internet.

PARA 2004 featured invited talks, contributed talks, minisymposia, and software and hardware vendors. The first day, June 20, was devoted to two parallel tutorials. The minisymposia and contributed talks during the main part of the workshop, June 21–23, were scheduled in parallel sessions. All invited and contributed talks were noncommercial. The workshop attracted 230 speakers from all over the world.

The PARA 2006 Workshop with the title “State-of-the-Art in Scientific and Parallel Computing” will be held in Umeå (Sweden) on June 17–21, 2006.

Tutorials

Validated Scientific Computing Using Interval Analysis was organized by *George F. Corliss* from Marquette University (USA). This tutorial gave an introduction to concepts and patterns of interval analysis. It was assumed that the participants had had a first course in scientific computation, including floating-point arithmetic, error analysis, automatic differentiation, Gaussian elimination, Newton's method, numerical optimization, and Runge-Kutta methods for ODEs. The tutorial included lectures, with examples in MATLAB and Sun's Fortran 95, and a set of supervised, hands-on exercises.

Automatic Differentiation was organized by *Andrea Walther* from the Technical University of Dresden (Germany). This tutorial gave a detailed introduction to the chain

rule based technique of automatic differentiation (AD) that provides first and higher derivatives without incurring truncation errors. Several examples illustrated the theoretical results. Some AD tools, selected as a reasonably representative sample, were tested in supervised, hands-on exercises.

Key Speakers

Richard P. Brent, Oxford University Computing Laboratory (UK), *Fast and Reliable Random Number Generators for Scientific Computing*. Fast and reliable pseudo-random number generators are required for simulation and other applications in scientific computing. Richard outlined the requirements for good uniform random number generators, and described a class of generators having very fast vector/parallel implementations with excellent statistical properties.

Bernd Dammann and **Henrik Madsen**, the Technical University of Denmark (Denmark), *High Performance Computing and the Importance of Code Tuning—Some Practical Experiences from Program Tuning at the DTU HPC Center*. This talk gave a short overview of the High Performance Computer installation at the Technical University of Denmark (DTU), as well as a summary of some code tuning experiments. It is easy to reduce the run time of an application for a given problem by buying a computer with a faster CPU (higher clock frequency). However, very often the same or even better speed-up of the code can be achieved by analyzing and tuning the code—without the need to invest in new hardware.

Jack Dongarra, the University of Tennessee and Oak Ridge National Laboratory (USA), *High Performance Computing Trends and Self Adapting Numerical Software (SANS)—Effort*. In this talk Jack looked at how high performance computing has changed over the last 10 years and predicted future trends. In addition, he advocated the need for self adapting software.

Iain Duff, the Rutherford Appleton Laboratory (UK) and CERFACS (France), *Partitioning and Parallelism in the Solution of Large Sparse Systems*. Iain first reviewed the various levels of parallelism that are available in the direct solution of large sparse linear systems. He also briefly considered iterative as well as direct methods in this study.

Fred Gustavson, the IBM T.J. Watson Research Center (USA), *Ideas for High Performance Linear Algebra Software*. In this talk Fred presented several ideas for the development of sequential and parallel HPC dense linear algebra software. The main results were obtained from the algorithms and architecture approach.

Per Christian Hansen, the Technical University of Denmark (Denmark), *Large-Scale Methods in Inverse Problems*. Inverse problems arise in geophysics, tomography, image deblurring and many other areas where the goal is to compute interior or hidden information from exterior data. This talk presented a survey of numerical methods and paradigms suited for large-scale inverse problems.

Bo Kågström, the University of Umeå (Sweden), *Recursive Blocked Algorithms and Hybrid Data Structures for Dense Matrix Library Software*. Matrix computations are both fundamental and ubiquitous in computational science and its vast application areas. Along with the development of more advanced computer systems with complex memory

hierarchies, there is a continuing demand for new algorithms and library software that efficiently utilize and adapt to new architecture features.

John K. Reid, the Rutherford Appleton Laboratory (UK), *Fortran Is Getting More and More Powerful*. There is much happening just now with respect to Fortran. The features of Fortran 2003 have been chosen and the standard is very near completion. John is the convener of the ISO Fortran Committee.

Peter Sloot, the University of Amsterdam (The Netherlands). *Scientific Computing in the Grid: A Biophysical Case Study*. Workers at the University of Amsterdam conducted computer simulation experiments in pre-operative planning of vascular reconstruction with a physician in the experimental loop. Peter showed new results from numerical simulations of blood flow with 3D cellular automata.

Zahari Zlatev, National Environmental Research Institute (Denmark), *Large-Scale Computations with the Unified Danish Eulerian Model*. The Unified Danish Eulerian Model (UNI-DEM) is a mathematical model for performing different comprehensive studies related to damaging effects from high pollution levels in Denmark and Europe. The model is described by a system of partial differential equations (PDEs).

Minisymposia

Interval Methods, organized by *Luke Achenie*, University of Connecticut (USA), *Vladik Kreinovich*, University of Texas at El Paso (USA), and *Kaj Madsen*, Technical University of Denmark (Denmark). In many practical problems there is a need to (a) solve systems of equations and inequalities, and/or (b) optimize some performance measure. The results obtained by conventional algorithms are either local or cannot be guaranteed. Interval analysis provides guaranteed approximations of the set of all the actual solutions of the problem. This ensures that no solution is missed. There were 21 speakers in this minisymposium.

Trends in Large Scale Computing, organized by *Scott B. Baden*, University of California at San Diego (USA). Software infrastructures for large scale computation often fail to realize the full potential afforded by technological advances, and the result is lost opportunities for making scientific discoveries. This minisymposium examined two important issues in software infrastructure for large scale computation: achieving scalability, and optimization through specialization. There were 5 speakers in this minisymposium.

High Performance Linear Algebra Algorithms, organized by *Fred G. Gustavson*, IBM T.J. Watson Research Center (USA), and *Jerzy Wasniewski*, Technical University of Denmark (Denmark). The algorithms of Linpack and Eispack and later LAPACK and ScaLAPACK have stood the test of time in terms of robustness and accuracy. The focus of this minisymposium was on explaining high performance versions of these algorithms. There were 7 speakers in this minisymposium.

Substructuring, Dimension Reduction and Applications, organized by *Zhaojun Bai*, University of California (USA) and *Rencang Li*, University of Kentucky USA. There are a variety of reasons to go for substructuring and dimension reduction in scientific computations and applications. Substructuring makes it possible to solve large and seemingly intractable computational problems by some kind of divide-and-conquer technique. It

also offers a general methodology for parallelization. There were 12 speakers in this minisymposium.

Parallel Processing in Science and Engineering, organized by *Adam W. Bojańczyk*, Cornell University (USA). This minisymposium concerned selected aspects of parallel and distributing computing as they arise in engineering. Both non-traditional applications as well as relevant software tools were presented. There were 9 speakers in this minisymposium.

Distributed Computing: Tools, Paradigms and Infrastructures, organized by *Beniamino Di Martino*, *Rocco Aversa*, Second University of Naples (Italy), and *Laurence Tianruo Yang*, Francis Xavier University (Canada). The minisymposium presented recent advances in distributed computing technology, methodology and tools. The presentations featured a variety of topics ranging from mobile and location-aware computing to skeletons and high-level parallel languages, from programming environments and tools for Grid applications development and tuning, to distributed monitoring and security issues. There were 9 speakers in this minisymposium.

High-Performance Computing in Earth and Space Science, organized by *Peter Messmer*, Tech-X Corporation at Boulder (USA). High-performance computing facilities enable simulations of physical phenomena with ever increasing fidelity and accuracy. The range of resolved scales in a single simulation, as well as the number of physical processes included, yield results that can be directly compared with observational data. There were 7 speakers in this minisymposium.

Advanced Algorithms and Software Components for Scientific Computing, organized by *Padma Raghavan*, Pennsylvania State University (USA). This minisymposium concerned algorithms for sparse linear systems solution and function approximation and their implementation using advanced software architectures. Discussions emphasized the role of such techniques for improving the performance of long-running PDE-based simulations. There were 7 speakers in this minisymposium.

Software Engineering and Problem Solving Environments for Scientific Computing, organized by *José C. Cunha*, Universidade Nova de Lisboa (Portugal) and *Omer F. Rana*, Cardiff University (UK). The emergence of computational grids in the last few years provides new opportunities for the scientific community to undertake collaborative and multi-disciplinary research. The aim of this minisymposium was to bring together experts who have experience in developing software tools to support application scientists, and those who make use of these tools. There were 5 speakers in this minisymposium.

Runtime Software Techniques for Enabling High-Performance Applications, organized by *Masha Sosonkina*, Iowa State University (USA). Parallel computing platforms are advancing rapidly, both in speed and size. However, often only a fraction of the peak hardware performance is achieved by high-performance scientific applications. One way to cope with the changeability of hardware is to start creating applications able to adapt themselves “on-the-fly”. The talks of the minisymposium discussed this issue from both the application-centric and system-centric viewpoints. There were 6 speakers in this minisymposium.

Sparse Direct Linear Solvers, organized by *Sivan Toledo*, Tel Aviv University (Israel). The matrices of most of the systems of linear algebraic equations arising from scientific and engineering applications are sparse. This minisymposium dealt with some modern algorithms for sparse direct linear solvers. There were 12 speakers in this minisymposium.

Treatment of Large Scientific Models, organized by *Krassimir Georgiev*, Bulgarian Academy of Science (Bulgaria) and *Zahari Zlatev*, National Environmental Research Institute (Denmark). The exploitation of new fast computers in the effort to avoid non-physical assumptions and, thus, to develop and run more reliable and robust large scientific models was the major topic of this minisymposium. There were 9 speakers in this minisymposium.

Performance Evaluation and Design of Hardware-Aware PDE Solvers, organized by *Markus Kowarschik* and *Frank Hülsemann*, University of Erlangen-Nuremberg (Germany). In an ideal situation, all performance optimization of computationally intensive software would take place automatically, allowing the researchers to concentrate on the development of more efficient methods rather than having to worry about performance. However, for the time being, the need to identify and remove the performance bottlenecks of computationally intensive codes remains. As an example of a class of computationally intensive problems, this minisymposium concentrated on the numerical solution of PDEs. There were 7 speakers in this minisymposium.

Computationally Expensive Methods in Statistics, organized by *Wolfgang Hartmann*, SAS Institute Inc. (USA) and *Paul Somerville*, University of Central Florida (USA). A two-dimensional data set with N observations (rows) and n variables (columns) and large scale data requires intensive computational work. Of course there may be even more dimensions of the data set. There were 5 speakers in this minisymposium.

Approaches or Methods of Security Engineering (AMSE), organized by *Taihoon Kim* and *Ho Yeol Kwon*, Kangwon National University (Korea). Security engineering software is needed for reducing security holes. The talks presented a number of methods for designing such software. There were 16 speakers in this minisymposium.

Contributed Talks

Some contributed talks were added to the minisymposium sessions. The others were organized in the following independent sessions: two sessions of “Grid and Network”, two sessions of “HPC Applied to Security Problems”, two sessions of “Clusters and Graphics”, one session of “HPC Applied to Cryptology”, one session of “ODEs, PDEs and Automatic Differentiation”, one session of “Computer Tools”, and a special session of “Computer Vendors”.

Workshop Proceedings

The proceedings of the PARA 2004 Workshop are divided into two complementary publications, this Springer volume and the following report:

- J. Dongarra, K. Madsen and J. Waśniewski (Eds.)

- ▶ Complementary proceedings of the PARA 2004 Workshop on State-of-the-Art in Scientific Computing, Lyngby, Denmark, June, 2004.
- ▶ IMM-Technical report-2005-09.
- ▶ Informatics and Mathematical Modelling, Technical University of Denmark, DK-2800 Lyngby, Denmark.
- ▶ URL: http://www2.imm.dtu.dk/pubdb/views/publication_details.php?id=3927

A list of those papers appearing in the report is given in this volume in Chapter II of the contributed talks.

Acknowledgments

The organizers are indebted to the PARA 2004 sponsors, whose support was vital to the success of the workshop. The main sponsor was the Department of Informatics and Mathematical Modelling of the Technical University of Denmark. The other sponsors were:

- HP High Performance Computing USA,
- NAG Numerical Algorithms Group Ltd. UK,
- Comsol A/S Denmark (MATLAB distributor),
- Sun Microsystems Denmark,
- UNI•C Danish Computing Center Denmark,
- Microsoft Denmark, and
- IBM (International Business Machines) Denmark.

The organizers would like to thank Kirsten Probst for her excellent work as PARA 2004 secretary. Dorthe Thøgersen, Henrik Krogh and other staff of the conference also provided very valuable help.

We are very grateful to Prof. Ho Yeol Kwon from the Kangwon National University, Electrical and Computer Engineering Department for taking many photos during the PARA 2004 conference. These are available at the PARA 2004 URL

<http://www.imm.dtu.dk/~jw/para04/>.

Thanks are also due to Vincent A. Barker for his kind assistance in the preparation of both the workshop and these proceedings.

The PARA 2004 conference ended on June 23, 2004. The evening of June 23, the eve of St. Hans Day, is celebrated in Denmark by the lighting of bonfires. We are indebted to Arriva Denmark A/S for making available two boats for the PARA 2004 participants, from which we could see some of the bonfires and admire beautiful Copenhagen.

Finally, we would like to thank the PARA 2004 referees for their careful evaluation of the workshop papers.

Jack Dongarra
Kaj Madsen
Jerzy Waśniewski

Table of Contents

Invited Talks

Fast and Reliable Random Number Generators for Scientific Computing	1
<i>Richard P. Brent</i>	
New Generalized Data Structures for Matrices Lead to a Variety of High Performance Dense Linear Algebra Algorithms	11
<i>Fred G. Gustavson</i>	
Management of Deep Memory Hierarchies – Recursive Blocked Algorithms and Hybrid Data Structures for Dense Matrix Computations	21
<i>Bo Kågström</i>	
Fortran Is Getting More and More Powerful	33
<i>John K. Reid</i>	
Large-Scale Computations with the Unified Danish Eulerian Model	43
<i>Zahari Zlatev</i>	

Minisymposia

I Interval Methods

Interval Methods: An Introduction	53
<i>Organizers: Luke E.K. Achenie, Vladik Kreinovich, and Kaj Madsen</i>	
A Chemical Engineering Challenge Problem That Can Benefit from Interval Methods	57
<i>Luke E.K. Achenie and Gennadi M. Ostrovsky</i>	
Performance of Taylor Model Methods for Validated Integration of ODEs	65
<i>Martin Berz and Kyoko Makino</i>	
On the Use of Intervals in Scientific Computing: What Is the Best Transition from Linear to Quadratic Approximation?	75
<i>Martine Ceberio, Vladik Kreinovich, and Lev Ginzburg</i>	

HPC-ICTM: The Interval Categorizer Tessellation-Based Model for High Performance Computing	83
<i>Marilton S. de Aguiar, Graçaliz P. Dimuro, Fábila A. Costa, Rafael K.S. Silva, César A.F. De Rose, Antônio C.R. Costa, and Vladik Kreinovich</i>	
Counting the Number of Connected Components of a Set and Its Application to Robotics	93
<i>Nicolas Delanoue, Luc Jaulin, and Bertrand Cottenceau</i>	
Interval-Based Markov Decision Processes for Regulating Interactions Between Two Agents in Multi-agent Systems	102
<i>Graçaliz P. Dimuro and Antônio C.R. Costa</i>	
A Domain Theoretic Account of Euler's Method for Solving Initial Value Problems	112
<i>Abbas Edalat and Dirk Pattinson</i>	
Reliable Computation of Equilibrium States and Bifurcations in Nonlinear Dynamics	122
<i>C. Ryan Gwaltney and Mark A. Stadtherr</i>	
A Verification Method for Solutions of Linear Programming Problems	132
<i>Ismail I. Idriss</i>	
Compressing 3D Measurement Data Under Interval Uncertainty	142
<i>Olga Kosheleva, Sergio Cabrera, Brian Usevitch, and Edward Vidal Jr.</i>	
Computing Interval Bounds for Statistical Characteristics Under Expert-Provided Bounds on Probability Density Functions	151
<i>Victor G. Krymsky</i>	
Interval Parallel Global Optimization with Charm++	161
<i>José A. Martínez, Leocadio G. Casado, José A. Alvarez, and Inmaculada García</i>	
On the Approximation of Interval Functions	169
<i>Klaus Meer</i>	
The Distributed Interval Geometric Machine Model	179
<i>Renata H.S. Reiser, Antônio C.R. Costa, and Graçaliz P. Dimuro</i>	
New Algorithms for Statistical Analysis of Interval Data	189
<i>Gang Xiang, Scott A. Starks, Vladik Kreinovich, and Luc Longpré</i>	
On Efficiency of Tightening Bounds in Interval Global Optimization	197
<i>Antanas Žilinskas and Julius Žilinskas</i>	

II Trends in Large Scale Computing

Trends in Large Scale Computing: An Introduction	206
<i>Organizer: Scott B. Baden</i>	
Ygdrasil: Aggregator Network Toolkit for Large Scale Systems and the Grid	207
<i>Susanne M. Balle, John Bishop, David LaFrance-Linden, and Howard Rifkin</i>	
Enabling Coupled Scientific Simulations on the Grid	217
<i>Alan Sussman and Henrique Andrade</i>	

III High Performance Linear Algebra Algorithms

High Performance Linear Algebra Algorithms: An Introduction	225
<i>Organizers: Fred G. Gustavson and Jerzy Waśniewski</i>	
Applying Software Testing Metrics to Lapack	228
<i>David J. Barnes and Tim R. Hopkins</i>	
A Matrix-Type for Performance–Portability	237
<i>N. Peter Drakenberg</i>	
A New Array Format for Symmetric and Triangular Matrices	247
<i>John A. Gunnels and Fred G. Gustavson</i>	
A Family of High-Performance Matrix Multiplication Algorithms	256
<i>John A. Gunnels, Fred G. Gustavson, Greg M. Henry, and Robert A. van de Geijn</i>	

IV Substructuring, Dimension Reduction and Applications

Substructuring, Dimension Reduction and Applications: An Introduction	266
<i>Organizers: Zhaojun Bai and Ren-Cang Li</i>	
Parallel Algorithms for Balanced Truncation Model Reduction of Sparse Systems	267
<i>José M. Badía, Peter Benner, Rafael Mayo, and Enrique S. Quintana-Ortí</i>	
Towards an Optimal Substructuring Method for Model Reduction	276
<i>Zhaojun Bai and Ben-Shan Liao</i>	
Model Reduction for RF MEMS Simulation	286
<i>David S. Bindel, Zhaojun Bai, and James W. Demmel</i>	

A Model-Order Reduction Technique for Low Rank Rational Perturbations of Linear Eigenproblems	296
<i>Frank Blömeling and Heinrich Voss</i>	
Parallel Global Optimization of Foundation Schemes in Civil Engineering	305
<i>Raimondas Čiegis, Milda Baravykaite, and Rimantas Belevičius</i>	
A Combined Linear and Nonlinear Preconditioning Technique for Incompressible Navier-Stokes Equations	313
<i>Feng-Nan Hwang and Xiao-Chuan Cai</i>	
Structure-Preserving Model Reduction	323
<i>Ren-Cang Li and Zhaojun Bai</i>	
A Comparison of Parallel Preconditioners for the Sparse Generalized Eigenvalue Problems by Rayleigh-Quotient Minimization	333
<i>Sangback Ma and Ho-Jong Jang</i>	
Theoretical Relations Between Domain Decomposition and Dynamic Substructuring	342
<i>Daniel J. Rixen</i>	
Model Order Reduction for Large Scale Engineering Models Developed in ANSYS	349
<i>Evgenii B. Rudnyi and Jan G. Korvink</i>	
Rational Krylov for Large Nonlinear Eigenproblems	357
<i>Axel Ruhe</i>	
Algebraic Sub-structuring for Electromagnetic Applications	364
<i>Chao Yang, Weiguo Gao, Zhaojun Bai, Xiaoye S. Li, Lie-Quan Lee, Parry Husbands, and Esmond G. Ng</i>	
V Parallel Processing in Science and Engineering	
Parallel Processing in Science and Engineering: An Introduction	374
<i>Organizer: Adam W. Bojańczyk</i>	
Rapid Development of High-Performance Linear Algebra Libraries	376
<i>Paolo Bientinesi, John A. Gunnels, Fred G. Gustavson, Greg M. Henry, Margaret Myers, Enrique S. Quintana-Ortí, and Robert A. van de Geijn</i>	
Automatic Derivation of Linear Algebra Algorithms with Application to Control Theory	385
<i>Paolo Bientinesi, Sergey Kolos, and Robert A. van de Geijn</i>	
Cluster Computing for Financial Engineering	395
<i>Shirish Chinchalkar, Thomas F. Coleman, and Peter Mansfield</i>	

Semi-automatic Generation of Grid Computing Interfaces for Numerical Software Libraries	404
<i>Erik Elmroth and Rikard Skelander</i>	
Rapid Development of High-Performance Out-of-Core Solvers	413
<i>Thierry Joffrain, Enrique S. Quintana-Ortí, and Robert A. van de Geijn</i>	
ALPS: A Software Framework for Parallel Space-Time Adaptive Processing	423
<i>Kyusoon Lee and Adam W. Bojańczyk</i>	
Hybrid Parallelization of CFD Applications with Dynamic Thread Balancing	433
<i>Alexander Spiegel, Dieter an Mey, and Christian Bischof</i>	
VI Distributed Computing: Tools, Paradigms and Infrastructures	
Distributed Computing: Tools, Paradigms and Infrastructures. An Introduction	442
<i>Organizers: Beniamino Di Martino, Rocco Aversa, and Laurence Tianruo Yang</i>	
Parallelization of GSL: Performance of Case Studies	444
<i>José Aliaga, Francisco Almeida, José M. Badía, Sergio Barrachina, Vicente Blanco, María Castillo, U. Dorta, Rafael Mayo, Enrique S. Quintana, Gregorio Quintana, Casiano Rodríguez, and Francisco de Sande</i>	
Design of Policy-Based Security Mechanisms in a Distributed Web Services Architecture	454
<i>Valentina Casola, Antonino Mazzeo, Nicola Mazzocca, and Salvatore Venticinqu</i>	
Supporting Location-Aware Distributed Applications on Mobile Devices	464
<i>Cristiano di Flora, Massimo Ficco, and Stefano Russo</i>	
Grid Application Development on the Basis of Web Portal Technology	472
<i>Gábor Dózsa, Péter Kacsuk, and Csaba Németh</i>	
A Distributed Divide and Conquer Skeleton	481
<i>Juan R. González, Coromoto León, and Casiano Rodríguez</i>	
A Tool to Display Array Access Patterns in OpenMP Programs	490
<i>Oscar R. Hernandez, Chunhua Liao, and Barbara M. Chapman</i>	

A Model Analysis of a Distributed Monitoring System Using a Multi-formalism Approach	499
<i>Mauro Iacono, Stefano Marrone, Nicola Mazzocca, Francesco Moscato, and Valeria Vittorini</i>	
Performance Oriented Development and Tuning of GRID Applications	509
<i>Emilio Mancini, Massimiliano Rak, Roberto Torella, and Umberto Villano</i>	
Towards a Bulk-Synchronous Distributed Shared Memory Programming Environment for Grids	519
<i>Håkan Mattsson and Christoph Kessler</i>	
VII HPC in Earth and Space Science	
High-Performance Computing in Earth- and Space-Science: An Introduction	527
<i>Organizer: Peter Messmer</i>	
Applying High Performance Computing Techniques in Astrophysics	530
<i>Francisco Almeida, Evencio Mediavilla, Alex Osoez, and Francisco de Sande</i>	
Statistical Properties of Dissipative MHD Accelerators	538
<i>Kaspar Arzner, Loukas Vlahos, Bernard Knaepen, and Nicolas Denewet</i>	
A Simulation Model for Forest Fires	546
<i>Gino Bella, Salvatore Filippone, Alessandro De Maio, and Mario Testa</i>	
MHD Modeling of the Interaction Between the Solar Wind and Solar System Objects	554
<i>Andreas Ekenbäck and Mats Holmström</i>	
Implementing Applications with the Earth System Modeling Framework	563
<i>Chris Hill, Cecelia DeLuca, V. Balaji, Max Suarez, Arlindo da Silva, William Sawyer, Carlos Cruz, Atanas Trayanov, Leonid Zaslavsky, Robert Hallberg, Byron Boville, Anthony Craig, Nancy Collins, Erik Kluzek, John Michalakes, David Neckels, Earl Schwab, Shepard Smithline, Jon Wolfe, Mark Iredell, Weiyu Yang, Robert Jacob, and Jay Larson</i>	
Parallel Discrete Event Simulations of Grid-Based Models: Asynchronous Electromagnetic Hybrid Code	573
<i>Homa Karimabadi, Jonathan Driscoll, Jagrut Dave, Yuri Omelchenko, Kalyan Perumalla, Richard Fujimoto, and Nick Omid</i>	
Electromagnetic Simulations of Dusty Plasmas	583
<i>Peter Messmer</i>	

VIII Advanced Algorithms and Software Components for Scientific Computing

Advanced Algorithms and Software Components for Scientific Computing: An Introduction	590
<i>Organizer: Padma Raghavan</i>	
Extending PSBLAS to Build Parallel Schwarz Preconditioners	593
<i>Alfredo Buttari, Pasqua D'Ambra, Daniela di Serafino, and Salvatore Filippone</i>	
A Direct Orthogonal Sparse Static Methodology for a Finite Continuation Hybrid LP Solver	603
<i>Pablo Guerrero-García and Ángel Santos-Palomo</i>	
Applying Parallel Direct Solver Techniques to Build Robust High Performance Preconditioners	611
<i>Pascal Hénon, François Pellegrini, Pierre Ramet, Jean Roman, and Yousef Saad</i>	
The Design of Trilinos	620
<i>Michael A. Heroux and Marzio Sala</i>	
Software Architecture Issues in Scientific Component Development	629
<i>Boyana Norris</i>	
Parallel Hybrid Sparse Solvers Through Flexible Incomplete Cholesky Preconditioning	637
<i>Keita Teranishi and Padma Raghavan</i>	
Parallel Heuristics for an On-Line Scientific Database for Efficient Function Approximation	644
<i>Ivana Veljkovic and Paul E. Plassmann</i>	

IX Software Engineering and Problem Solving Environments for Scientific Computing

Software Engineering and Problem Solving Environments for Scientific Computing: An Introduction	654
<i>Organizers: Jose C. Cunha and Omer F. Rana</i>	
A General Architecture for Grid-Based PSE Toolkits	656
<i>Mario Cannataro, Carmela Comito, Antonio Congiusta, Gianluigi Folino, Carlo Mastroianni, Andrea Pugliese, Giandomenico Spezzano, Domenico Talia, and Pierangelo Veltri</i>	
An Expert Assistant for Computer Aided Parallelization	665
<i>Gabriele Jost, Robert Chun, Haoqiang Jin, Jesus Labarta, and Judit Gimenez</i>	

Scalable Middleware Environment for Agent-Based Internet Applications	675
<i>Benno J. Overeinder and Frances M.T. Brazier</i>	

Automatic Generation of Wrapper Code and Test Scripts for Problem Solving Environments	680
<i>Andreas Schreiber</i>	

X Runtime Software Techniques for Enabling High-Performance Applications

Runtime Software Techniques for Enhancing High-Performance Applications: An introduction	690
<i>Masha Sosonkina</i>	

Efficient Execution of Scientific Computation on Geographically Distributed Clusters	691
<i>Eduardo Argollo, Dolores Rexachs, Fernando G. Tinetti, and Emilio Luque</i>	

Improving the Performance of Large-Scale Unstructured PDE Applications	699
<i>Xing Cai</i>	

A Runtime Adaptive Load Balancing Algorithm for Particle Simulations	709
<i>Matthew F. Dixon</i>	

Evaluating Parallel Algorithms for Solving Sylvester-Type Matrix Equations: Direct Transformation-Based Versus Iterative Matrix-Sign-Function-Based Methods	719
<i>Robert Granat and Bo Kågström</i>	

Performance Analysis for Parallel Adaptive FEM on SMP Clusters	730
<i>Judith Hippold and Gudula Rünger</i>	

Performance Tuning of Matrix Triple Products Based on Matrix Structure	740
<i>Eun-Jin Im, Ismail Bustany, Cleve Ashcraft, James W. Demmel, and Katherine A. Yelick</i>	

Adapting Distributed Scientific Applications to Run-Time Network Conditions	747
<i>Masha Sosonkina</i>	

XI Sparse Direct Linear Solvers

Sparse Direct Linear Solvers: An Introduction	756
<i>Organizer: Sivan Toledo</i>	

Oblio: Design and Performance	758
<i>Florin Dobrian and Alex Pothén</i>	

Performance Analysis of Parallel Right-Looking Sparse LU Factorization on Two Dimensional Grids of Processors	768
<i>Laura Grigori and Xiaoye S. Li</i>	
A Shared- and Distributed-Memory Parallel Sparse Direct Solver	778
<i>Anshul Gupta</i>	
Simple and Efficient Modifications of Elimination Orderings	788
<i>Pinar Heggernes and Yngve Villanger</i>	
Optimization of a Statically Partitioned Hypermatrix Sparse Cholesky Factorization	798
<i>José R. Herrero and Juan J. Navarro</i>	
Maximum-Weighted Matching Strategies and the Application to Symmetric Indefinite Systems	808
<i>Stefan Röllin and Olaf Schenk</i>	
An Evaluation of Sparse Direct Symmetric Solvers: An Introduction and Preliminary Findings	818
<i>Jennifer A. Scott, Yifan Hu, and Nicholas I.M. Gould</i>	

XII Treatment of Large Scale Models

Treatment of Large Scientific Problems: An Introduction	828
<i>Organizers: Zahari Zlatev and Krassimir Georgiev</i>	
Towards a Parallel Multilevel Preconditioned Maxwell Eigensolver	831
<i>Peter Arbenz, Martin Bečka, Roman Geus, and Ulrich Hetmaniuk</i>	
On Improvement of the Volcano Search and Optimization Strategy	839
<i>Venansius Baryamureeba and John Ngubiri</i>	
Aggregation-Based Multilevel Preconditioning of Non-conforming FEM Elasticity Problems	847
<i>Radim Blaheta, Svetozar Margenov, and Maya Neytcheva</i>	
Efficient Solvers for 3-D Homogenized Elasticity Model	857
<i>Ronald H.W. Hoppe and Svetozara I. Petrova</i>	
Performance Evaluation of a Parallel Algorithm for a Radiative Transfer Problem	864
<i>Paulo B. Vasconcelos and Filomena d'Almeida</i>	

XIII Performance Evaluation and Design of Hardware-Aware PDE Solvers

Performance Evaluation and Design of Hardware-Aware PDE Solvers: An Introduction	872
<i>Organizers: Frank Hülsemann and Markus Kowarschik</i>	

A Cache-Aware Algorithm for PDEs on Hierarchical Data Structures	874
<i>Frank Günther, Miriam Mehl, Markus Pögl, and Christoph Zenger</i>	
Constructing Flexible, Yet Run Time Efficient PDE Solvers	883
<i>Frank Hülsemann and Benjamin Bergen</i>	
Analyzing Advanced PDE Solvers Through Simulation	893
<i>Henrik Johansson, Dan Wallin, and Sverker Holmgren</i>	
Towards Cache-Optimized Multigrid Using Patch-Adaptive Relaxation	901
<i>Markus Kowarschik, Iris Christadler, and Ulrich Rüde</i>	
Hierarchical Partitioning and Dynamic Load Balancing for Scientific Computation	911
<i>James D. Teresco, Jamal Faik, and Joseph E. Flaherty</i>	
Cache Optimizations for Iterative Numerical Codes Aware of Hardware Prefetching	921
<i>Josef Weidendorfer and Carsten Trinitis</i>	
XIV Computationally Expensive Methods in Statistics	
Computationally Expensive Methods in Statistics: An Introduction	928
<i>Organizer: Wolfgang M. Hartmann</i>	
Dimension Reduction vs. Variable Selection	931
<i>Wolfgang M. Hartmann</i>	
Reproducible Statistical Analysis in Microarray Profiling Studies	939
<i>Ulrich Mansmann, Markus Ruschhaupt, and Wolfgang Huber</i>	
Step-Down FDR Procedures for Large Numbers of Hypotheses	949
<i>Paul N. Somerville</i>	
XV Approaches or Methods of Security Engineering (AMSE)	
Applying Security Engineering to Build Security Countermeasures: An Introduction	957
<i>Organizers: Tai-hoonn Kim and Ho-yeol Kwon</i>	
CC-SEMS: A CC Based Information System Security Evaluation Management System	964
<i>Young-whan Bang, Yeun-hee Kang, and Gang-soo Lee</i>	
A Secure Migration Mechanism of Mobile Agents Under Mobile Agent Environments	974
<i>Dongwon Jeong, Young-Gab Kim, Young-Shil Kim, Lee-Sub Lee, Soo-Hyun Park, and Doo-Kwon Baik</i>	

A Flexible Privilege Management Scheme for Role Graph Model	983
<i>Yuna Jung and Eenjun Hwang</i>	
The System Modeling for Detections of New Malicious Codes	992
<i>EunYoung Kim, CheolHo Lee, HyungGeun Oh, and JinSeok Lee</i>	
Information Hiding Method Using CDMA on Wave Files	1000
<i>Young-Shil Kim, Sang Yun Park, Suk-Hee Wang, and Seung Lee</i>	
Efficient Key Distribution Protocol for Electronic Commerce in Mobile Communications	1009
<i>Jin Kwak, Soohyun Oh, and Dongho Won</i>	
A Framework for Modeling Organization Structure in Role Engineering	1017
<i>HyungHyo Lee, YoungLok Lee, and BongNam Noh</i>	
An Efficient Pointer Protection Scheme to Defend Buffer Overflow Attacks	1025
<i>Yongsu Park and Yookun Cho</i>	

Contributed Talks

I Contributed Talks in this Volume

Parallel Hierarchical Radiosity: The PIT Approach	1031
<i>Fabrizio Baiardi, Paolo Mori, and Laura Ricci</i>	
Optimizing Locationing of Multiple Masters for Master-Worker Grid Applications	1041
<i>Cyril Banino</i>	
An OGSA-Based Bank Service for Grid Accounting Systems	1051
<i>Erik Elmroth, Peter Gardfjäll, Olle Mulmo, and Thomas Sandholm</i>	
A Grid Resource Broker Supporting Advance Reservations and Benchmark-Based Resource Selection	1061
<i>Erik Elmroth and Johan Tordsson</i>	
The Dragon Graph: A New Interconnection Network for High Speed Computing	1071
<i>Jywe-Fei Fang</i>	
Speeding up Parallel Graph Coloring	1079
<i>Assefaw H. Gebremedhin, Fredrik Manne, and Tom Woods</i>	
On the Efficient Generation of Taylor Expansions for DAE Solutions by Automatic Differentiation	1089
<i>Andreas Griewank and Andrea Walther</i>	

Edge-Disjoint Hamiltonian Cycles of WK-Recursive Networks	1099
<i>Chien-Hung Huang, Jywe-Fei Fang, and Chin-Yang Yang</i>	
Simulation-Based Analysis of Parallel Runge-Kutta Solvers	1105
<i>Matthias Korch and Thomas Rauber</i>	
A Novel Task Scheduling Algorithm for Distributed Heterogeneous Computing Systems	1115
<i>Guan-Joe Lai</i>	
Study of Load Balancing Strategies for Finite Element Computations on Heterogeneous Clusters	1123
<i>Kalyani Munasinghe and Richard Wait</i>	
Parallel Algorithms for the Determination of Lyapunov Characteristics of Large Nonlinear Dynamical Systems	1131
<i>Günter Radons, Gudula Rünger, Michael Schwind, and Hong-liu Yang</i>	
Online Task Scheduling on Heterogeneous Clusters: An Experimental Study	1141
<i>Einar M.R. Rosenvinge, Anne C. Elster, and Cyril Banino</i>	
A Parallel Method for Large Sparse Generalized Eigenvalue Problems by OmniRPC in a Grid Environment	1151
<i>Tetsuya Sakurai, Kentaro Hayakawa, Mitsuhsa Sato, and Daisuke Takahashi</i>	
An Implementation of Parallel 3-D FFT Using Short Vector SIMD Instructions on Clusters of PCs	1159
<i>Daisuke Takahashi, Taisuke Boku, and Mitsuhsa Sato</i>	
II Contributed Talks Appearing Elsewhere	
Other Para'04 Contributed Talks	1168
Author Index	1171