

Lecture Notes in Computer Science

1947

Tor Sørsvik Fredrik Manne Randi Moe
Assefaw Hadish Gebremedhin (Eds.)

Applied Parallel Computing

New Paradigms for HPC
in Industry and Academia

5th International Workshop, PARA 2000
Bergen, Norway, June 2000
Proceedings



Springer

Tor Sørøvik Fredrik Manne Randi Moe
Assefaw Hadish Gebremedhin (Eds.)

Applied Parallel Computing

New Paradigms for HPC
in Industry and Academia

5th International Workshop, PARA 2000
Bergen, Norway, June 18-20, 2000
Proceedings



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany

Juris Hartmanis, Cornell University, NY, USA

Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Tor Sørevik

Fredrik Manne

Assefaw Hadish Gebremedhin

University of Bergen, Department of Informatics

5020 Bergen, Norway

E-mail: {Tor.Sorevik/Fredrik.Manne/assefaw}@ii.uib.no

Randi Moe

University of Bergen, Parallab

5020 Bergen, Norway

E-mail: Randi.Moe@ii.uib.no

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Applied parallel computing : new paradigms for HPC in industry and
academia ; 5th international workshop ; proceedings / PARA 2000,

Bergen, Norway, June 18 - 20, 2000. Tor Sørevik ... (ed.). - Berlin ;

Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ;

Paris ; Singapore ; Tokyo : Springer, 2001

(Lecture notes in computer science ; 1947)

ISBN 3-540-41729-X

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

ISSN 0302-9743

ISBN 3-540-41729-X Springer-Verlag 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-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York

a member of BertelsmannSpringer Science+Business Media GmbH

© Springer-Verlag Berlin Heidelberg 2001

Printed in Germany

Typesetting: Camera-ready by author, data conversion by PTP Berlin, Stefan Sossna

Printed on acid-free paper SPIN 10780911 06/3142 5 4 3 2 1 0

Lecture Notes in Computer Science

1947

Edited by G. Goos, J. Hartmanis and J. van Leeuwen

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Singapore

Tokyo

Preface

The papers in this volume were presented at PARA 2000, the Fifth International Workshop on Applied Parallel Computing. PARA 2000 was held in Bergen, Norway, June 18-21, 2000. The workshop was organized by Parallab and the Department of Informatics at the University of Bergen. The general theme for PARA 2000 was *New paradigms for HPC in industry and academia* focusing on:

- High-performance computing applications in academia and industry,
- The use of Java in high-performance computing,
- Grid and Meta computing,
- Directions in high-performance computing and networking,
- Education in Computational Science.

The workshop included 9 invited presentations and 39 contributed presentations. The PARA 2000 meeting began with a one-day tutorial on OpenMP programming led by Timothy Mattson. This was followed by a three-day workshop.

The first three PARA workshops were held at the Technical University of Denmark (DTU), Lyngby (1994, 1995, and 1996). Following PARA'96, an international steering committee for the PARA meetings was appointed and the committee decided that a workshop should take place every second year in one of the Nordic countries. The 1998 workshop was held at Umeå University, Sweden. One important aim of these workshops is to strengthen the ties between HPC centers, academia, and industry in the Nordic countries as well as worldwide. The University of Bergen organized the 2000 workshop and the next workshop in the year 2002 will take place at the Helsinki University of Technology, Espoo, Finland.

October 2000

Tor Sørenvik
Fredrik Manne
Randi Moe
Assefaw Hadish Gebremedhin

Organization

PARA 2000 was organized by the Department of Informatics at the University of Bergen and by Parallab.

Organizing Committee

Conference Chair:	Petter Bjørstad (University of Bergen, Norway)
Local Organization:	Assefaw Hadish Gebremedhin (University of Bergen, Norway)
	Fredrik Manne (University of Bergen, Norway)
	Randi Moe (University of Bergen, Norway)
	Synnøve Palmstrøm (University of Bergen, Norway)
	Tor Sørevik (University of Bergen, Norway)
	Aase Spieler (University of Bergen, Norway)

Steering Committee

Petter Bjørstad	University of Bergen (Norway) <i>PARA 2000 Chairman</i>
Jack Dongarra	University of Tennessee and Oak Ridge National Laboratory (USA)
Björn Engquist	PDC, Royal Institute of Technology, Stockholm (Sweden)
Kristjan Jonasson	University of Iceland, Reykjavik (Iceland)
Kari Laasonen	Center for Scientific Computing, Espoo (Finland)
Bo Kågström	Umeå University and HPC2N (Sweden), <i>PARA '98 Chairman</i>
Risto Nieminen	Helsinki University of Technology, Espoo (Finland)
Karstein Sørli	SINTEF, Dept of Industrial Mathematics, Trondheim (Norway)
Jerzy Waśniewski	Danish Computing Centre for Research and Education (UNI•C), Lyngby (Denmark), <i>PARA '94-'96 Chairman</i>

Sponsoring Institutions

The City of Bergen

The Research Council of Norway (NFR)

Sun Microsystems AS, Norway

SGI Norge AS, Norway

The University of Bergen

Table of Contents

Speaker Information

Industrial Applications: Challenges in Modeling and Computing	1
<i>Jari Järvinen, Juha Haataja, Jari Hämäläinen</i>	
Data Grids for Next Generation Problems in Science and Engineering	17
<i>Carl Kesselman</i>	
High-Performance Computers: Yesterday, Today, and Tomorrow	18
<i>Erik Hagersten</i>	
The Evolution of OpenMP	19
<i>Timothy Mattson</i>	
JavaGrande – High Performance Computing with Java	20
<i>Michael Philippsen, Ronald F. Boisvert, Valdimir S. Getov, Roldan Pozo, José Moreira, Dennis Gannon, Geoffrey C. Fox</i>	
Ocean and Climate Prediction on Parallel Super Computers	37
<i>Geir Evensen</i>	
LAWRA: Linear Algebra with Recursive Algorithms	38
<i>Bjarne S. Andersen, Fred Gustavson, Alexander Karaivanov, Minka Marinova, Jerzy Waśniewski, Plamen Yalamov</i>	
Solving CFD Problems with Open Source Parallel Libraries	52
<i>Bill Gropp</i>	
High-Performance Library Software for <i>QR</i> Factorization	53
<i>Erik Elmroth, Fred Gustavson</i>	
Parallel Triangular Sylvester-Type Matrix Equation Solvers for SMP Systems Using Recursive Blocking	64
<i>Isak Jonsson, Bo Kågström</i>	
On the Efficiency of Scheduling Algorithms for Parallel Gaussian Elimination with Communication Delays	74
<i>Raimondas Čiegis, Vadimas Starikovičius, Jerzy Waśniewski</i>	
High Performance Cholesky Factorization via Blocking and Recursion That Uses Minimal Storage	82
<i>Fred Gustavson, Isak Jonsson</i>	

Parallel Two-Stage Reduction of a Regular Matrix Pair to Hessenberg-Triangular Form	92
<i>Björn Adlerborn, Krister Dackland, Bo Kågström</i>	
A Fast Minimal Storage Symmetric Indefinite Solver	103
<i>Fred Gustavson, Alexander Karaivanov, Minka Marinova, Jerzy Waśniewski, Plamen Yalamov</i>	
A Scalable Parallel Assembly for Irregular Meshes Based on a Block Distribution for a Parallel Block Direct Solver	113
<i>David Goudin, Jean Roman</i>	
MUMPS: A General Purpose Distributed Memory Sparse Solver	121
<i>Patrick R. Amestoy, Iain S. Duff, Jean-Yves L'Excellent, Jacko Koster</i>	
Runtime Adaptation of an Iterative Linear System Solution to Distributed Environments	131
<i>Masha Sosonkina</i>	
A Local Refinement Algorithm for Data Partitioning	140
<i>Jarmo Rantakokko</i>	
Feedback Guided Scheduling of Nested Loops	149
<i>T. L. Freeman, D. J. Hancock, J. M. Bull, R. W. Ford</i>	
A Comparison of Partitioning Schemes for Blockwise Parallel SAMR Algorithms	160
<i>Johan Steensland, Stefan Söderberg, Michael Thuné</i>	
Parallelizing an Adaptive Dynamical Grid Generator in a Climatological Trace Gas Transport Application	170
<i>Jörn Behrens</i>	
Optimal Parameter Values for a Parallel Structured Adaptive Mesh Refinement Algorithm	177
<i>Michael Thuné, Stefan Söderberg</i>	
Partition of Unstructured Finite Element Meshes by a Multilevel Approach	187
<i>Noureddine Bouhmala, Xing Cai</i>	
GRISK: An Internet Based Search for K-Optimal Lattice Rules	196
<i>Tor Sørøvik, Jan Frode Myklebust</i>	
Parallel and Distributed Document Overlap Detection on the Web	206
<i>Krisztián Monostori, Arkady Zaslavsky, Heinz Schmidt</i>	
A Parallel Implementation of a Job Shop Scheduling Heuristic	215
<i>U. Der, K. Steinhöfel</i>	

Restructuring Irregular Computations for Distributed Systems Using Mobile Agents	223
<i>Rocco Aversa, Beniamino Di Martino, Nicola Mazzocca</i>	
An Information System for Long-Distance Cooperation in Medicine	233
<i>Harald Kosch, Renata Słota, László Böszörményi, Jacek Kitowski, Janusz Otfinowski</i>	
Hydra - Decentralized and Adaptative Approach to Distributed Computing	242
<i>Marian Bubak, Paweł Płaszczak</i>	
Object-Oriented Approach to Finite Element Modeling on Clusters	250
<i>Roman Wyrzykowski, Tomasz Olas, Norbert Szyci</i>	
An Object Oriented Framework for Parallel Multiple Optimizations	258
<i>Beidi S. Hamma</i>	
Experiments in Separating Computational Algorithm from Program Distribution and Communication	268
<i>R. B. Yehezkael, Y. Wiseman, H. G. Mendelbaum, I. L. Gordin</i>	
Performance Tuning on Parallel Systems: All Problems Solved?	279
<i>Holger Brunst, Wolfgang E. Nagel, Stephan Seidl</i>	
Performance Measurement Support for MPI Applications with PATOP ...	288
<i>Marian Bubak, Włodzimierz Funika, Bartosz Baliś, Roland Wismüller</i>	
A Parallel Volume Visualization Using Extended Space Leaping Method ..	296
<i>Sung-Up Jo, Chang-Sung Jeong</i>	
Hands-On Training for Undergraduates in High-Performance Computing Using Java	306
<i>Christian H. Bischof, H. Martin Bucker, Jörg Henrichs, Bruno Lang</i>	
A Parallel 3-D FFT Algorithm on Clusters of Vector SMPs	316
<i>Daisuke Takahashi</i>	
High-End Computing on SHV Workstations Connected with High Performance Network	324
<i>Lars Paul Huse, Håkon Bugge</i>	
From the Big Bang to Massive Data Flow: Parallel Computing in High Energy Physics Experiments	333
<i>C. Adler, J. Berger, D. Flierl, H. Helstrup, J. S. Lange, J. Lien, V. Lindenstruth, D. Röhrich, D. Schmischke, M. Schulz, B. Skaali, H. K. Sollveit, T. Steinbeck, R. Stock, C. Struck, K. Ullaland, A. Vestbø, A. Wiebalck</i>	

A Data Parallel Formulation of the Barnes-Hut Method for
N-Body Simulations 342
M. Amor, F. Argüello, J. López, O. Plata, E. L. Zapata

Par-T: A Parallel Relativistic Fully 3D Electromagnetic Particle-in-Cell
Code 350
Peter Messmer

Evaluation of MPI's One-Sided Communication Mechanism for
Short-Range Molecular Dynamics on the Origin2000..... 356
T. Matthey, J. P. Hansen

Ship Hull Hydrodynamic Analysis Using Distributed Shared Memory 366
João Silva, Paulo Guedes

Domain Decomposition Solvers for Large Scale Industrial Finite Element
Problems 373
Petter E. Bjørstad, Jacko Koster, Piotr Krzyżanowski

A High Parallel Procedure to Initialize the Output Weights of a Radial
Basis Function or BP Neural Network 384
Rossella Cancelliere

High-Performance Computing in Geomechanics by a Parallel Finite
Element Approach 391
Felicja Okulicka – Dlużewska

Author Index 399

Industrial Applications: Challenges in Modeling and Computing

Jari Järvinen¹, Juha Haataja¹, and Jari Hämäläinen²

¹ CSC – Scientific Computing Ltd., P.O. Box 405, FIN-02101 Espoo, Finland
{Jari.Jarvinen, Juha.Haataja}@csc.fi

² Valmet Corp., P.O. Box 587, FIN-40101 Jyväskylä, Finland
Jari.P.Hamalainen@valmet.com

Abstract. Complex industrial applications can be examined using modeling and computing even when theoretical and experimental approaches fail. Modeling and computing are rapidly developing tools for engineering and science, but applying these tools in practice is a challenge. This article focuses on industrial multi-physics and multi-scale problems in the production of silicon crystals and the design of paper machines.

Keywords: Modeling, computing, multi-physics, multi-scale, industry, parallel computing, numerical methods, paper machines, silicon crystals.

1 Introduction

Scientific and engineering problems can be examined using modeling and computing [5,17], but this approach poses several challenges. We focus on industrial multi-physics applications, such as the growth of silicon crystals [16] and the modeling of paper machines [12,6].

The growth of silicon crystals, shown in Fig. 1, contains coupled physical phenomena: electromagnetism, fluid dynamics, heat transfer, and structural mechanics. Thorough modeling of this application requires multi-scale modeling spanning several length and time scales [3,16].

Modeling of paper machines is a different kind of challenge. The headbox of a paper machine, presented in Fig. 2, contains flows in a complex geometry. These flows are three-dimensional, turbulent and multi-phase. A comprehensive mathematical model leads to demanding computational tasks.

Computation of a multi-physics problem is a challenge. Coupled problems lead to systems of nonlinear equations, which have to be solved iteratively. Solving these coupled systems can be a huge computational effort. We'll present some ideas on numerical solution strategies in multi-physics problems. In addition, we'll discuss how to combine computational fluid dynamics and optimization.




Fig. 1. Growth of a silicon crystal.



Fig. 2. The headbox of a paper machine.

2 Possibilities and Challenges in Industrial Applications

Industrial applications are difficult to model and compute: both the physical phenomena and the structures and geometries are typically very complex. These modeling tasks involve significant challenges for scientific computing, and high performance computing (HPC) is undoubtedly needed. Thus, HPC is an essential part of industrial research and development (R&D).

There are a multitude of modeling problems presented by the modern industry. Challenges are presented, for example, by phenomena in the human body or by the physics of miniaturized electronics components.

Microelectromechanical systems (MEMS) [20] pose a multi-physics problem, which contains electrostatics, structural analysis, fluid dynamics and heat transfer in miniaturized geometries. In fact, modeling MEMS devices may require us to combine molecular dynamics and continuous models.

Multi-physics problems are solved increasingly often. As an example, the model may contain fluid-structure interactions. Other possibilities are free surface problems or phase-change problems. These problems are typically solved in 2D today, but in 3D in the future.

Another demanding task is the modeling the human body, which is of interest to the pharmaceutical industry. We may be interested in simulating the blood flow in 3D bifurcations, designing new drugs, or modeling how electromagnetic fields behave in living tissues.

The modeling of turbulence in fluid dynamics poses significant challenges. Today, the Reynolds stress model is applied routinely in industrial problems, and large eddy simulation (LES) is becoming feasible. However, direct numerical simulation (DNS) of the Navies-Stokes equations will be impossible for several years to come.

Another challenge are multi-phase problems. This is somewhat similar to turbulence modeling: there exist smaller and smaller phenomena arising from direct simulation, and it is less and less feasible to perform averaging. Furthermore, multi-scale problems present us with a combination of different scales (micro, meso and macro scales).

We may need to perform several simulations to solve a single problem. For example, one may be interested in finding the optimal geometry in a fluid dynamics application. This is an example of shape optimization in computational fluid dynamics (CFD). For each optimization run, one may have to solve a CFD problem tens, hundreds or thousands of times. Also, there are requirements for performing real-time interactive simulations, for example in a virtual reality environment, or a virtual learning system (Fig. 3).

3 The Three Methods of R&D

We can use three different approaches in industrial R&D: experimental, theoretical and computational (Fig. 4). All these approaches are essential and needed, and they are usually applied together to solve a problem.




Fig. 3. Industrial R&D utilizes modern computer technology: interacting with a simulation in a virtual reality environment.

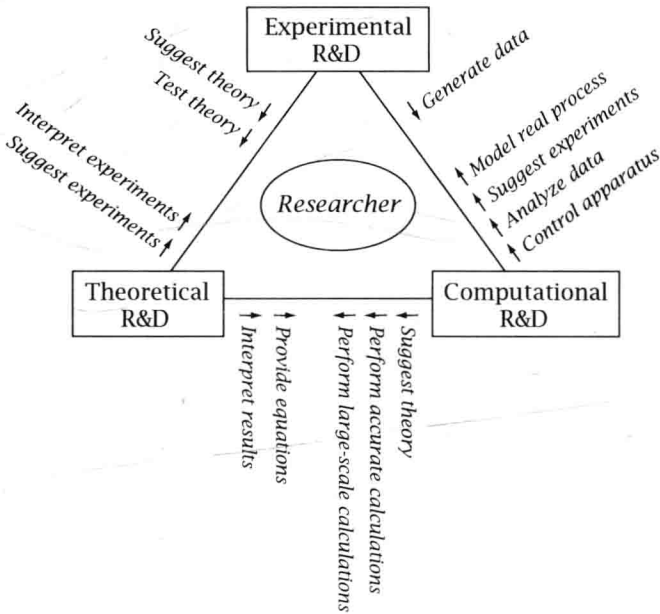


Fig. 4. The three methods of R&D.

The experimental approach is typically straightforward — basically, we measure what happens in the physical process we are interested in. Experiments are often reliable, at least to a known extent, which depends on the physical setup and the human interface of the experiment. Also, experiments are usually quantitative. However, experiments can be laborious or even impossible to carry out in complex cases. We mention the lack of commercial techniques to measure the turbulence of dense fiber-water suspensions. Experiments can also be very expensive and slow to carry out. For example, one may have to build a prototype to do the experiment.

The theoretical approach is efficient in simple cases. It is usually qualitative and suggestive — a good way to check the numerical or experimental results. However, the theoretical approach is insufficient in complex cases.

Computational modeling is repeatable, which is an asset: one does not need to build a new prototype to make a new experiment. However, there are a great many steps where mistakes can occur: mathematical modeling, selection of the numerical methods, and implementation of the methods. In addition, unanticipated hardware and software problems may occur during the computation.

4 The Challenges of Industrial HPC

Traditionally, industrial R&D has involved experiments and the internal know-how of a company. Nowadays R&D has shifted to a more innovative direction: combining the use of computation, experiments, and theory. Because all these approaches demand a different set of skills, this usually involves a collaboration of several experts.

The computational approach requires people who know how to solve mathematical models on a computer. Also, programming skills are required: real-life problems are not usually solvable using off-the-shelf software packages. Thus, R&D departments have to compete for talented people in other commercial sectors: www services, networking, etc.

The modeling and computing facilities of industrial companies are usually modest. Due to the fast development of this area, there is a continuing lack of know-how of numerical methods, programming, and computer technology (efficient workstations, clusters etc.). If one builds a computational research center from scratch, the costs easily exceed any budget. Therefore, industry is interested in co-operation with sites which have the required expertise, in addition to software and hardware resources. Usually this co-operation involves university laboratories.

However, also in the universities there is a lack of efficient training programs on HPC. At least in Finland, there may be just a single course on HPC topics, or none. Therefore, there is urgent need for teaching modeling and computing tools to the students at the university level. These courses should encourage students to use modeling techniques, numerical methods, commercial software packages, and programming languages.