

Dan Grigoras · Alex Nicolau
Bernard Toursel · Bertil Folliot (Eds.)

Advanced Environments, Tools, and Applications for Cluster Computing

NATO Advanced Research Workshop, IWCC 2001
Mangalia, Romania, September 2001
Revised Papers



Springer

Dan Grigoras Alex Nicolau
Bernard Tournel Bertil Folliot (Eds.)

Advanced Environments, Tools, and Applications for Cluster Computing

NATO Advanced Research Workshop, IWCC 2001
Mangalia, Romania, September 1-6, 2001
Revised Papers



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany

Juris Hartmanis, Cornell University, NY, USA

Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Dan Grigoras

University College Cork, Computer Science Department

College Road, Cork, Ireland

E-mail: d.grigoras@cs.ucc.ie

Alex Nicolau

University of California, Information and Computer Science Department

444 Computer Science Building, Irvine, CA 92697-3425, USA

E-mail: nicolau@ics.uci.edu

Bernard Tourssel

University Lille 1, Laboratoire d'Informatique Fondamentale de Lille (LIFL)

59655 Villeneuve d'Ascq, France

E-mail: bernard.tourssel@lifel.fr

Bertil Folliot

University Paris 6, LIP, 8, rue du Capitaine Scott, 75015 Paris, France

E-mail: bertil.folliot@lip6.fr

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Advanced environments, tools, and applications for cluster computing : NATO advanced research workshop ; revised papers / IWCC 2001, Mangalia, Romania, September 1 - 6, 2001. Dan Grigoras ... (ed.). - Berlin ; Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ; Tokyo : Springer, 2002 (Lecture notes in computer science ; Vol. 2326)

ISBN 3-540-43672-3

CR Subject Classification (1998): C.2, D.1.3, D.4, F.1.2, K.4.4, C.1, D.2

ISSN 0302-9743

ISBN 3-540-43672-3 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

<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2002

Printed in Germany

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

Printed on acid-free paper SPIN 10846725 06/3142 5 4 3 2 1 0

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Tokyo

Preface

Started by small group of well-known scientists with the aim of sharing knowledge, experiences, and results on all aspects of cluster computing, the initiative of a workshop on cluster computing received more attention after IFIP WG 10.3 and IEEE Romania Section accepted our request for sponsorship. Moreover, the application for a NATO ARW grant was successful, leading to a greater interest in the workshop. In this respect, we have to say that we chose Romania in order to attract scientists from Central and Eastern European countries and improve the cooperation in the region, in the field of cluster computing.

We had an extremely short time to organize the event, but many people joined us and enthusiastically contributed to the process. The success of the workshop is wholly due to the hard work of the organizing committee, members of the program committee, key speakers, speakers from industry, and authors of accepted papers. The workshop consisted of invited and regular paper presentations, followed by discussions, on many important current and emerging topics ranging from scheduling and load balancing to grids. The key speakers devoted their time and efforts to presenting the most interesting results of their research groups, and we all thank them for this. All papers were peer reviewed by two or three reviewers.

The proceedings published by Springer-Verlag include 8 invited papers and 24 regular papers, in that benefited from discussions during the workshop. We once again thank the authors for their efforts in preparing high-quality papers.

The workshop was generously sponsored by NATO Scientific Affairs Division, and co-sponsored by the IFIP Working Group 10.3, and IEEE Romania Sections. The “Gh. Asachi” Technical University, “Al. I. Cuza” University, and Black Sea University Foundation greatly contributed to the organization of the event. We also received generous support from the Romanian Ministry of Education and Research, Microsoft Romania, BRD-GSG, and Romaqua Group.

We would like to thank Mitica Craus for his help in processing this volume, and Delia Mitrea, Cristian Butincu, and Cristian Amarandei for their assistance with local organization matters.

March 2002

Dan Grigoras
Alex Nicolau
Bernard Tournel
Bertil Folliot

Chairs

Alex NICOLAU, University of California at Irvine, USA
Bernard TOURSEL, Université des Sciences et Technologies de Lille, France
Bertil FOLLIOT, Université Pierre et Marie Curie, Paris, France
Dan GRIGORAS, Technical University of Iasi (on leave), Romania

Organizing Committee

Bertil FOLLIOT, Université Pierre et Marie Curie, Paris, France
Dan GRIGORAS, Technical University of Iasi, Romania
Alex NICOLAU, University of California at Irvine, USA
Isaac SCHERSON, University of California at Irvine, USA
Bernard TOURSEL, Université des Sciences et Technologies de Lille, France
Wolfgang GENTZSCH, SUN Microsystems, Grid Middleware, Palo Alto, USA

Program Committee

Florian Mircea BOIAN, University "Babes-Bolyai", Cluj-Napoca, Romania
Valentin CRISTEA, "Politehnica" University Bucharest, Romania
Michel COSNARD, INRIA Lorraine, France
Bertil FOLLIOT, Université Pierre et Marie Curie, Paris, France
Guang GAO, University of Delaware, USA
Wolfgang GENTZSCH, SUN Microsystems, Grid Middleware, Palo Alto, USA
Claude GIRAULT, Université Pierre et Marie Curie, Paris, France
Lucio GRANDINETTI, University of Calabria, Italy
Dan GRIGORAS, Technical University Iasi, Romania
Toader JUCAN, University "Al. I. Cuza" Iasi, Romania
Alex NICOLAU, University of California at Irvine, USA
Can OZTURAN, Bogazici University, Turkey
K. PINGALI, Cornell University, USA
L. RAUSCHWERGER, Texas A&M University, USA
Isaac SCHERSON, University of California at Irvine, USA
Bernard TOURSEL, Université des Sciences et Technologies de Lille, France
Marek TUDRUJ, Institute of Computer Science, Polish Academy of Sciences, Poland
Andrew WENDELBORN, University of Adelaide, Australia

Key Speakers

Alex NICOLAU, University of California at Irvine, USA

Bernard TOURSEL, Université des Sciences et Technologies de Lille, France

Isaac SCHERSON, University of California at Irvine, USA

Bertil FOLLLOT, Université Pierre et Marie Curie, Paris, France

Ionut LOPATAN, Microsoft Romania

Can OZTURAN, Bogazici University, Turkey

Marek TUDRUJ, Institute of Computer Science, Polish Academy of Sciences, Poland

Hai JIN, Huazhong University of Science and Technology, China

Lothar LIPPERT, GENIAS Software GmbH, Germany

Valentin CRISTEA, "Politehnica" University Bucharest, Romania

Dan GRIGORAS, Technical University Iasi, Romania

Sponsors and/or Supporting Organizations



North Atlantic Treaty Organization (NATO)



International Federation for Information Processing (IFIP) W.G. 10.3



IEEE Romanian Section



Technical University of Iasi



“Al.I.Cuza” University



Black Sea University Foundation



Microsoft Romania



Banca Romana pentru Dezvoltare (BRD) – GSG



Romaqua Group

Table of Contents

Invited Papers

Grid Computing: A New Technology for the Advanced Web	1
<i>Wolfgang Gentzsch</i>	
Beyond Flexibility and Reflection: The Virtual Virtual Machine Approach	16
<i>B. Folliot, I. Piumarta, L. Seinturier, C. Baillarguet, C. Khoury, A. Leger, F. Ogel</i>	
Programming Models for Cluster Computing	26
<i>Dan Grigoras</i>	
About Design and Efficiency of Distributed Programming: Some Algorithmic Aspects	36
<i>Bernard Tourset</i>	
A Collaborative Environment for High Performance Computing	47
<i>Valentin Cristea</i>	
Layer 4 Fault Tolerance: Reliability Techniques for Cluster System in Internet Services	60
<i>Guang Tan, Hai Jin, Liping Pang</i>	
Computational Issues in E-commerce Trading Exchanges	69
<i>Can Özturan</i>	
Embedded Cluster Computing through Dynamic Reconfigurability of Inter-processor Connections	77
<i>Marek Tudruj</i>	

Papers

Parallelisation of Wave Propagation Algorithms for Odour Propagation in Multi-agent Systems	92
<i>Eugen Dedu, Stéphane Vialle, Claude Timsit</i>	
Parallel and Distributed Solutions for the Optimal Binary Search Tree Problem	103
<i>Mitică Craus</i>	
A Framework for Efficient Sparse LU Factorization in a Cluster Based Platform	118
<i>Laura Grigori</i>	

Task Allocation Using Processor Load Prediction on Multiprocessors Cluster	126
<i>Thierry Monteil, Patricia Pascal</i>	
Heterogeneous Dynamic Load Balancing	136
<i>Tiberiu Rotaru, Hans-Heinrich Nägeli</i>	
Multi-application Scheduling in Networks of Workstations and Clusters of Processors	145
<i>D. Kebbal, E.-G. Talbi, J.-M. Geib</i>	
The Parallel Algorithm of Conjugate Gradient Method	156
<i>Andrzej Jordan, Robert Piotr Bycul</i>	
A Parallel Similarity Search in High Dimensional Metric Space Using M-Tree	166
<i>Adil Alpkocak, Taner Danisman, Tuba Ulker</i>	
A User-Level Interface for Clustering Mathematical Software Kernels	172
<i>Dana Petcu, Horia Popa, Dorin Țepeneu</i>	
Implementing Multithreaded Protocols for Release Consistency on Top of the Generic DSM-PM ² Platform	179
<i>Gabriel Antoniu, Luc Bougé</i>	
Communication Balancing for Wireless Nodes in a LAN-WAN Cluster	189
<i>Cosmina Ivan</i>	
An Architecture and Task Scheduling Algorithm for Systems Based on Dynamically Reconfigurable Shared Memory Clusters	197
<i>Marek Tudruj, Lukasz Masko</i>	
Distributed Rendering Engine	207
<i>Nicolae Tapus, Emil Slusanschi, Tudor Popescu</i>	
A Distributed Application for Monte Carlo Simulations	216
<i>Nicolae Tapus, Mihai Burcea, Vlad Staicu</i>	
A Model for Accessing Resources of the Distributed File Systems	224
<i>Sabin C. Buraga</i>	
Coarse-Grain Parallelization of Test Vectors Generation on Multiprocessor Systems	231
<i>Felicia Ionescu, Mihail Ionescu, Cristina Coconu, Valentin Stoica</i>	
OpenMP/MPI Programming in a Multi-cluster System Based on Shared Memory/Message Passing Communication	241
<i>Adam Smyk, Marek Tudruj</i>	

Expressing Parallelism in Java Applications Distributed on Clusters	249
<i>Violeta Felea, Nathalie Devesa, Bernard Tournel, Pierre Lecouffe</i>	
Advanced Collaboration Techniques between Java Objects Distributed on Clusters	259
<i>Florian Mircea Boian, Corina Ferdean</i>	
Java Tools for Measurement of the Machine Loads	271
<i>A. Bouchi, R. Olejnik, Bernard Tournel</i>	
The Parallel FDTD Algorithm and Its Application to the Analysis of Complex Structures	279
<i>Boguslaw Butrylo</i>	
A Theoretical Application of Feedback Guided Dynamic Loop Scheduling	287
<i>Tatiana Tabirca, Len Freeman, Sabin Tabirca</i>	
DIM – A Distributed Image Processing System, Based on ISO IEC 12087 Image Processing Standard	293
<i>Paulina Mitrea</i>	
DisMedJava – A Distributed Application for Medical Image Processing . . .	308
<i>Cristian Buțincu, Dan Grigoras</i>	
Author Index	321

Grid Computing: A New Technology for the Advanced Web

Wolfgang Gentzsch

Sun Microsystems, Inc., Palo Alto, USA

Wolfgang.Gentzsch@Sun.COM

Abstract. The aim of our Grid Computing team at Sun is to take Cluster and Grid computing onto the next level of widest acceptance. For this purpose, we present a three-tier Grid architecture which relates to an intuitive model, including Cluster, Campus and Global Grids, which allows users to embrace Grid Computing starting with their current compute clusters, and moving to the next level of Grid implementation in an evolutionary way. In addition, we present a complete Grid software stack for all three Grid stages, which contains modules developed at Sun and by our partners.

1 Grid Components: Networks, Computers, Software

Mankind is right in the middle of another evolutionary technological transition which once more will change the way we do things. And, you guessed right, it has to do with the Internet. It's called "The Grid", which means the infrastructure for the Advanced Web, for computing, collaboration and communication.

The Internet itself has dramatically changed over the last three decades. While, in the late Sixties, it was built mainly to provide scientists with an infrastructure for faster communication via electronic mail, it has rapidly grown and improved since then, mainly because of three driving elements: networks, computers, and software.

In the mid Nineties, George Gilder predicted the "Network Abundance". Every nine months, total network bandwidth doubles. For many years, every day, thousands of miles of fiberoptic cables are laid down, ready for deployment in private and business applications. While many of us still suffer from 56 bps (Bits/sec) telephone modem speed, researchers and enterprises have already access to networks with a bandwidth of 10 million bps, and some even up to one billion bps. Soon, we will see 100 billion bps and more. Thus, network bandwidth will grow by a factor of 5000 over the next 10 years. Just this July, the US National Science Foundation approved the \$53 million DTF Distributed TeraScale Facility project, a network with 40 billion bps, connecting research centers in San Diego (SDSC), Pasadena (Caltech), Urbana-Champaign (NCSA), and Chicago (ARNL). It seems that there is no limit.

Another building block of the Internet are the computers. Today, for example, there are over one hundred million PCs in homes and at work, plus some 10 million powerful compute servers, from midrange to high-end, used at ISP Internet Service Providers, or for high-performance scientific, engineering and commercial appli-

cations. Their performance doubles every 18 months, which was observed and predicted by former Intel Chairman Gordon Moore in 1965. Thus, computer performance will grow by a factor of 100 over the next 10 years, then breaking the Petaflops Performance Barrier. This will make today's handheld electronic devices soon very powerful nodes in the Internet.

The third and most complex Internet building block is software, either for running the networks and the computers and their intercommunication - then called the *middleware*, or for solving our day-to-day problems and running our business, called *application* software. The ever increasing benefit, resulting in the combination of the networked computers, the software to run them, and the people who use them, is called Metcalfe's Law, after Bob Metcalfe, who developed Ethernet at Xerox Parc, in 1973: "The usefulness of a network equals the square of the number of users."

These three laws of Gilder, Moore, and Metcalfe, respectively, and the technological evolution they describe, are currently converging into and enabling the Advanced Web, on top of the Internet infrastructure. Past Internet and World Wide Web mainly enabled information provision, retrieval and exchange, and some e-commerce.

The new Advanced Web adds a wide variety of opportunities, based on computing, collaboration and communication, for individuals, groups, research and engineering teams, and for the whole community. It will provide great services in our private, community, and business environments. Universal connectivity gives users immediate and easy access to any kind of information and service they want, helps them in solving problems and in making personal and business decisions, and allows them to easily offer their own services to anybody. The new Advanced Web changes the way we live and work.

Enter *The Grid*. The term has been derived from the "Power Grid" infrastructure which provides electricity to every wall socket. In our context, The Grid describes the technology infrastructure for the Advanced Web, for computing, collaboration and communication.

2 Distributed Computing and the Grid

In the early Nineties, research groups started exploiting distributed computing resources over the Internet: scientists collected and utilized hundreds of workstations for parallel applications like molecular design and computer graphics rendering. Other research teams glued large supercomputers together into one virtual metacomputer, distributing subsets of a meta-application to specific vector, parallel and graphics computers, over wide-area networks, e.g. the computer simulation of multi-physics applications like the interaction of a fluid with a rotating propeller blade. Additionally the scope of many of these research projects was to understand and demonstrate the actual potential of the networking, computing and software infrastructure and to develop it further.