

Lecture Notes in Computer Science

1797

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

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Singapore

Tokyo

Babak Falsafi Mario Lauria (Eds.)

Network-Based Parallel Computing

Communication, Architecture,
and Applications

4th International Workshop, CANPC 2000
Toulouse, France, January 8, 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

Babak Falsafi
Purdue University, School of Electrical & Computer Engineering
1258 Electrical Engineering Building, West Lafayette, Indiana 47907-1285, USA
E-mail: babak@ecn.purdue.edu

Mario Lauria
Ohio State University, Department of Computer and Information Science
Dreese Lab 395, 2015 Neil Avenue, Columbus, OH 43210-1277, USA
E-mail: lauria@cis.ohio-state.edu

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Network-based parallel computing : communication, architecture, and applications ; 4th international workshop ; proceedings / CANPC 2000, Toulouse, France, January 8, 2000. Babak Falsafi ; Mario Lauria (ed.). - Berlin ; Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ; Singapore ; Tokyo : Springer, 2000
(Lecture notes in computer science ; Vol. 1797)
ISBN 3-540-67879-4

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

ISSN 0302-9743

ISBN 3-540-67879-4 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 is a company in the BertelsmannSpringer publishing group.
© Springer-Verlag Berlin Heidelberg 2000
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Christian Grosche, Hamburg
Printed on acid-free paper SPIN: 10720115 06/3142 5 4 3 2 1 0

Preface

Clusters of workstations/PCs connected by off-the-shelf networks have become popular as platforms for cost-effective parallel computing. Technological advances in both hardware and software have made such a network-based parallel computing platform an affordable alternative to commercial supercomputers for an increasing number of scientific applications.

Continuing in the tradition of the three previously successful workshops, this fourth Workshop on Communication, Architecture and Applications for Network-based Parallel Computing (CANPC 2000) brought together researchers and practitioners working in architecture, system software, applications, and performance evaluation to discuss state-of-the-art solutions for network-based parallel computing. This year, the workshop was held in conjunction with the sixth International Symposium on High-Performance Computer Architecture (HPCA-6).

As in prior editions, the papers presented here are representative of a spectrum of research efforts from groups in academia and industry to further improve cluster computing's viability, performance, cost-effectiveness, and usability. Specifically, we have arranged the contributions in this edition into four groups: (1) program development and execution support, (2) network router architecture, (3) system support for communication abstractions, and (4) network software and interface architecture.

The first group contains three papers that focus on programming and usability of clusters. Unlike tightly-coupled parallel supercomputers, clusters are built with commodity networking technology and operating systems. As such, clusters require additional support in the areas of program deployment and transparent and robust network connectivity. The specific contributions include a remote execution facility to deploy programs in a cluster in a transparent, secured, and decentralized manner; and a transparent mechanism to virtualize network connectivity in dynamic cluster environments. This group also includes a paper on a visual parallel programming tool for the BSP distributed programming model.

The second and third groups of papers improve on crucial aspects of the cluster technology such as routing (papers on Up*/Down* routing, dynamic routing reconfiguration, deadlock-free routing), and system support for derived data types and collective communication. These efforts show how the cluster concept is influencing research on topics that for years have been associated with traditional large-scale supercomputers, like interconnect topologies, routing schemes, parallel libraries.

The fourth group addresses the high software and network interface communication overheads characteristic of clusters. One paper goes back to the root of the network computing approach, and performs an in-depth analysis of existing gigabit architectures (VIA, Gigabit Ethernet). Another evaluates various design

points for the emerging industry standard for interface architectures, VIA. The third paper in this section proposes a novel messaging implementation in which the system learns and predicts patterns of message sequences to reduce the processing overhead.

Our excellent program this year was only made possible with the help and great efforts of many people. First, we would like to thank all of the authors for submitting papers, the program committee for their timely and meticulous reviewing and selection of the papers, and the HPCA-6 organizing committee for their support of this workshop. Special thanks to Henri Bal for giving an excellent keynote address. Finally, we would like to thank the editorial staff of Springer-Verlag for agreeing to publish the final version of these proceedings.

April 2000

Babak Falsafi and Mario Lauria
Program Co-chairs

CANPC 2000 Program Committee

Babak Falsafi, *Purdue University (co-chair)*

Mario Lauria, *The Ohio State Univeristy (co-chair)*

Henri Bal, *Vrije University*

Angelos Bilas, *University of Toronto*

Alok Choudhary, *Northwestern University*

Robert Clay, *Sandia National Labratories*

Jose Duato, *Universidad Politécnicna de Valencia*

Rudolf Eigenmann, *Purdue University*

Bob Horst, *3Ware Inc.*

Giulio Iannello, *Università di Napoli*

Liviu Iftode, *Rutgers University*

Alain Kägi, *Intel MRL*

Vijay Karamcheti, *New York University*

Steve Lumetta, *University of Illinois*

Dhabaleswar Panda, *The Ohio State University*

Phil Papadopoulos, *University of California*

Ioannis Schoinas, *Intel SAL*

Anand Sivasubramaniam, *The Penn State University*

Peter Steenkiste, *Carnegie Mellon University*

Robert Stets, *Compaq WRL*

Craig Stunkel, *IBM T.J. Watson*

Jaspal Subhlok, *University of Houston*

Rich Wolski, *University of Tennessee*

Table of Contents

Program Development and Execution Support

REXEC: A Decentralized, Secure Remote Execution Environment for Clusters	1
<i>Brent N. Chun, David E. Culler (University of California)</i>	
A Visual BSP Programming Environment for Distributed Computing	15
<i>Jeremy Martin (Oxford Supercomputing Center), Alex Wilson (Embedded Solutions, Ltd.)</i>	
Transparent Network Connectivity in Dynamic Cluster Environments	30
<i>Xiaodong Fu, Hua Wang, Vijay Karamcheti (New York University)</i>	

Network Router Architecture

A New Methodology to Computer Deadlock-Free Routing Tables for Irregular Networks	45
<i>José Carlos Sancho, Antonio Robles, José Duato (Universidad Politécnica de Valencia)</i>	
On the Performance of Up*/Down* Routing	61
<i>Francisco J. Alfaro (University of Murcia), Aurelio Bermúdez, Rafael Casado (University of Castilla-La Mancha), José Duato (Universidad Politécnica de Valencia), Francisco J. Quiles, José L. Sánchez (University of Castilla-La Mancha)</i>	
Extending Dynamic Reconfiguration to NOWs with Adaptive Routing	73
<i>Aurelio Bermúdez (University of Castilla-La Mancha), Francisco J. Alfaro (University of Murcia), Rafael Casado (University of Castilla-La Mancha), José Duato (Universidad Politécnica de Valencia), Francisco J. Quiles, José L. Sánchez (University of Castilla-La Mancha)</i>	

System Support for Communication Abstractions

MPI Derived Data Types Support in VIRTUS	84
<i>Rosario Cristaldi, Giulio Iannello (Università di Napoli)</i>	
Fast Collective Communication Algorithms for Reflective Memory Network Clusters	100
<i>Vijay Moorthy, Dhableswar K. Panda, P. Sadayappan (The Ohio State University)</i>	

Broadcast/Multicast over Myrinet Using NIC-Assisted Multidestination Messages 115
Darius Buntinas, Dhabaleswar K. Panda (The Ohio State University), José Duato (Universidad Politécnica de Valencia), P. Sadayappan (The Ohio State University)

Network Software and Interface Architecture

Simulation Studies of Gigabit Ethernet Versus Myrinet Using Real Application Cores 130
Helen Chen, Pete Wyckoff (Sandia National Laboratories)

Comparison and Evaluation of Design Choices for Implementing the Virtual Interface Architecture (VIA) 145
Mohammad Banikazemi (The Ohio State University), Bulent Abali (IBM TJ Watson), Dhabaleswar K. Panda (The Ohio State University)

Efficient Communication Using Message Prediction for Cluster Multiprocessors 162
Ahmad Afsahi, Nikitas J. Dimopoulos (University of Victoria)

Author Index 179