

Lecture Notes in Computer Science

853

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

Advisory Board: W. Brauer D. Gries J. Stoer



Kevin Bolding Lawrence Snyder (Eds.)

Parallel Computer Routing and Communication

First International Workshop, PCRCW '94
Seattle, Washington, USA, May 16-18, 1994
Proceedings

Springer-Verlag

Berlin Heidelberg New York
London Paris Tokyo
Hong Kong Barcelona
Budapest

Series Editors

Gerhard Goos

Universität Karlsruhe

Postfach 69 80, Vincenz-Priessnitz-Straße 1, D-76131 Karlsruhe, Germany

Juris Hartmanis

Department of Computer Science, Cornell University

4130 Upson Hall, Ithaca, NY 14853, USA

Jan van Leeuwen

Department of Computer Science, Utrecht University

Padualaan 14, 3584 CH Utrecht, The Netherlands

Volume Editors

Kevin Bolding

Lawrence Snyder

Department of Computer Science and Engineering, University of Washington

FR-35, Seattle, WA 98195, USA

CR Subject Classification (1991): C.2, C.1.2, B.4.3, C.4, F.2.2

ISBN 3-540-58429-3 Springer-Verlag Berlin Heidelberg New York

CIP data applied for

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 1994

Printed in Germany

Typesetting: Camera-ready by author

SPIN: 10478946 45/3140-543210 - Printed on acid-free paper

Preface

Due to advances in both hardware and systems design over the last decade, the power of parallel computers has leaped forward dramatically. The backbone of a massively parallel computer is the interconnection network that ties all of the processors together into one coherent machine. Although both interest in networks and routing, as well as the technology itself, has grown very quickly in recent years, there has been little cohesiveness to the routing/interconnection network community. For this reason, the first Parallel Computer Routing and Communication Workshop (PCRCW) was formed.

PCRCW '94, held at the University of Washington in Seattle on May 16, 17, and 18, 1994 had three primary goals:

1. To bring researchers from the routing/network community together.
2. To provide a platform for the presentation of the latest work in interconnection networks.
3. To produce a proceedings to disseminate the latest results to the entire research community.

The first goal was accomplished at the workshop, with over 60 attendees, approximately 40% of whom were from industry, with the remaining 60% from academia. The panels, banquet, and forum provided many opportunities for people to meet others in the community. The second goal was satisfied with the invited keynote address by Duncan Lawrie and 23 other talks covering the papers presented in this volume. Finally, the third goal is satisfied with the publication of these proceedings.

This volume includes 23 regular papers, covering areas from details of hardware design to proofs of theoretical issues. There are also many papers dealing with the performance of various adaptive routing schemes, new network topologies, network interfaces and fault-tolerance issues.

We would like to thank the many researchers and attendees who contributed to this successful workshop. The members of the program committee deserve extra thanks for helping in the organization of the workshop and the selection of papers to include. Clifford Lau and the US Office of Naval Research are especially appreciated for their support of this workshop. Finally, we wish to thank Kay Beck for handling the local arrangements and organizational details of the workshop in a nearly glitch-free manner.

Seattle, July 1994

Kevin Bolding
Lawrence Snyder

Program Committee

Lawrence Snyder, University of Washington (chair)

Kevin Bolding, University of Washington

Andrew Chien, University of Illinois

William Dally, Massachusetts Institute of Technology

Smaragda Konstantinidou, IBM T. J. Watson Research Center

Burton Smith, Tera Computer

Contents

Do Faster Routers Imply Faster Communication?	1
Vijay Karamcheti and Andrew A. Chien	
<i>University of Illinois</i>	
Fast Arbitration in Dilated Routers	16
Matthew E. Becker and Thomas F. Knight, Jr.	
<i>Massachusetts Institute of Technology</i>	
Performance Analysis of a Minimal Adaptive Router	31
Thu Duc Nguyen and Lawrence Snyder	
<i>University of Washington</i>	
Performance Evaluation of Adaptive Routing Algorithms for k-ary n-cubes	45
José Duato and Pedro López	
<i>Universidad Politécnica de Valencia</i>	
The Performance of Adaptive Routers on Worst Case Permutations	60
Donald D. Chinn	
<i>University of Washington</i>	
On the Effect of Queues Sizes and Channel Scheduling Policies in the Segment Router	72
Smaragda Konstantinidou	
<i>IBM T. J. Watson Research Center</i>	
The Offset Cube: An Optoelectronic Interconnection Network	86
D. Scott Wills, W. Stephen Lacy and Jose Cruz-Rivera	
<i>Georgia Institute of Technology</i>	
Three-Dimensional Network Topologies	101
John Nguyen	
<i>University of Michigan</i>	
John Pezaris, Gill Pratt and Steve Ward	
<i>Massachusetts Institute of Technology</i>	
Support for Multiple Classes of Traffic in Multicomputer Routers	116
Jennifer Rexford and Kang G. Shin	
<i>University of Michigan</i>	
Multidestination Message Passing Mechanism Conforming to Base Wormhole Routing Scheme	131
Dhabaleswar K. Panda, Sanjay Singal and Pradeep Prabhakaran	
<i>Ohio State University</i>	

Multi-address Encoding for Multicast	146
Chi-Ming Chiang and Lionel M. Ni	
<i>Michigan State University</i>	
Routing Algorithms for IBM SP1	161
Bülent Abali	
<i>IBM T. J. Watson Research Center</i>	
Cevdet Aykanat	
<i>Bilkent University</i>	
Congestion-Free Routing on the CM-5 Data Router	176
Steve Heller	
<i>Thinking Machines Corporation</i>	
ROMM Routing: A Class of Efficient Minimal Routing Algorithms	185
Ted Nesson	
<i>Harvard University</i>	
Lennart Johnsson	
<i>Thinking Machines Corporation</i>	
Packaging and Multiplexing of Hierarchical Scalable Expanders	200
Frederic T. Chong, Eric A. Brewer, F. Thomson Leighton and	
Thomas F. Knight, Jr.	
<i>Massachusetts Institute of Technology</i>	
Guaranteeing Idempotence for Tightly-Coupled, Fault-Tolerant Networks	215
Ian Eslick, André DeHon, and Thomas Knight, Jr.	
<i>Massachusetts Institute of Technology</i>	
Design of a Router for Fault-Tolerant Networks	226
Kevin Bolding and William Yost	
<i>University of Washington</i>	
The Reliable Router: A Reliable and High-Performance Communication	
Substrate for Parallel Computers	241
William J. Dally, Larry R. Dennison, David Harris, Kinhong Kan and	
Thucydides Xanthopoulos	
<i>Massachusetts Institute of Technology</i>	
Network Interface Support for User-Level Buffer Management	256
Cezary Dubnicki, Kai Li and Malena Mesarina	
<i>Princeton University</i>	

CRANIUM: An Interface for Message Passing on Adaptive Packet Routing Networks	266
Neil R. McKenzie, Kevin Bolding, Carl Ebeling and Lawrence Snyder <i>University of Washington</i>	
Optimized Routing in the Cray T3D	281
Steve Scott and Greg Thorson <i>Cray Research, Inc.</i>	
R2: A Damped Adaptive Router Design	295
A. Davis <i>University of Utah</i> R. Hodgson, I. Robinson, L. Cherkasova, V. Kotov and T. Rokicki <i>Hewlett Packard Laboratories</i>	
Arctic Routing Chip	310
G. Andrew Boughton <i>Massachusetts Institute of Technology</i>	