PROCEEDINGS

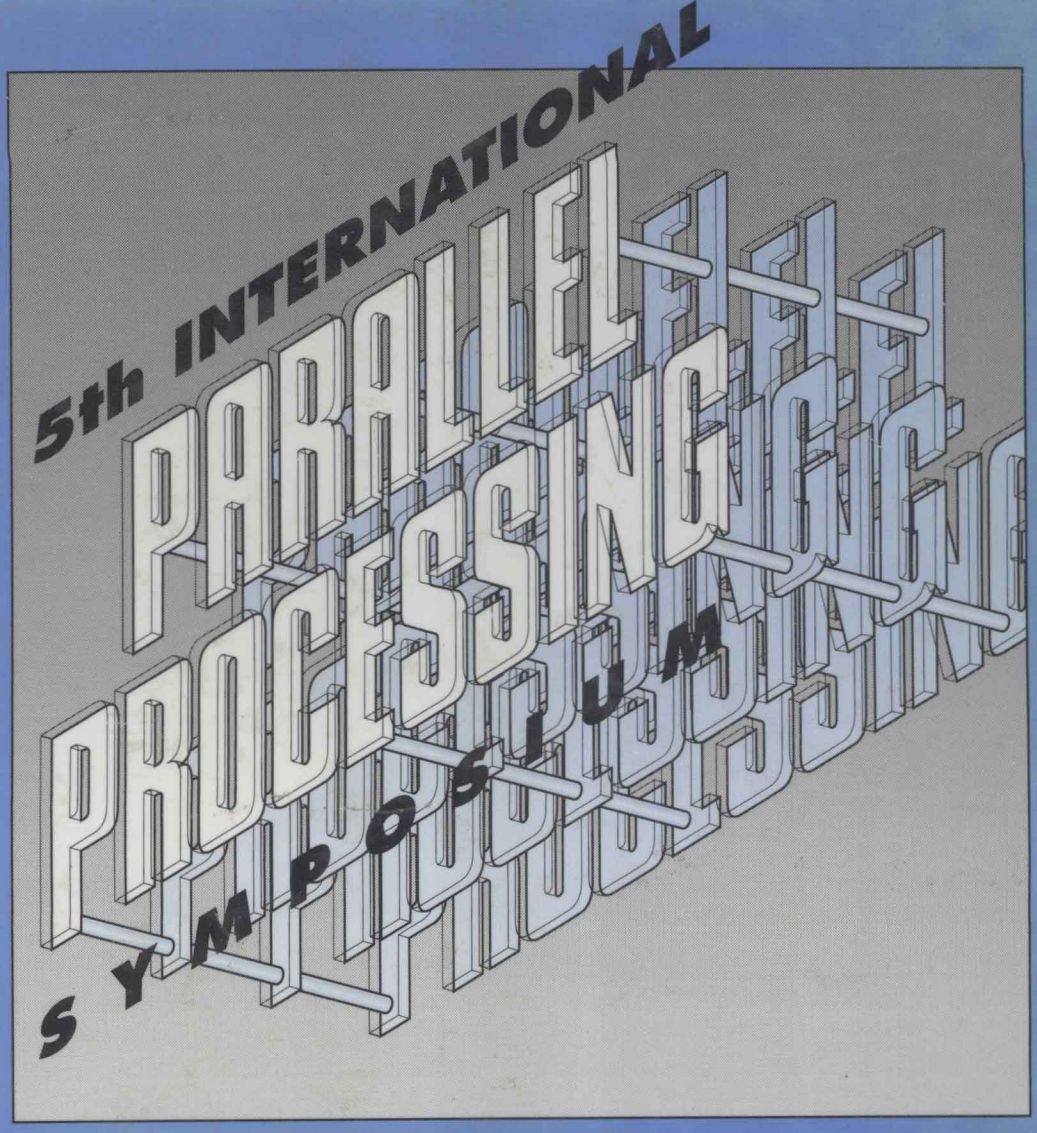

APRIL 30 - MAY 2, 1991 • ANAHEIM, CALIFORNIA

1951-1991
IEEE COMPUTER SOCIETY PRESS

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Proceedings

The Fifth International
Parallel Processing Symposium

V. K. Prasanna Kumar, Editor
University of Southern California

Proceedings of the fifth international symposium
sponsored by the
IEEE Computer Society
held in Anaheim, California,
April 30 - May 2, 1991

Additional support from:
Department of Electrical Engineering-Systems,
University of Southern California
Alliant Computer Systems Corporation
Encore Computer Corporation

1951-1991

IEEE Computer Society Press
Los Alamitos, California

Washington • Brussels • Tokyo

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and are published as presented and without change, in the interests of timely dissemination. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society Press, or The Institute of Electrical and Electronics Engineers, Inc.

Published by

IEEE Computer Society Press
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264

IEEE Computer Society Press Order Number 2167
Library of Congress Number 90-86330
IEEE Catalog Number 91TH0363-2
ISBN 0-8186-9167-0 (case)
ISBN 0-8186-6167-4 (microfiche)

Additional copies can be ordered from:

IEEE Computer Society Press
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264

IEEE Computer Society
13, avenue de l'Aquilon
B-1200 Brussels
BELGIUM

IEEE Computer Society
Ooshima Building
2-19-1 Minami-Aoyama,
Minato-Ku
Tokyo 107, JAPAN

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331

Production Editor: Wally Hutchins

Printed in the United States by Braun-Brumfield, Inc.

Cover by Alex Torres

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Proceedings

The Fifth International
Parallel Processing Symposium

1951-1991
40 YEARS OF SERVICE

IEEE COMPUTER SOCIETY
A member society of the
Institute of Electrical and Electronics Engineers, Inc.

PROCEEDINGS

APRIL 30 - MAY 2, 1991 • ANAHEIM, CALIFORNIA

1951-1991
IEEE COMPUTER SOCIETY PRESS

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Foreword

The papers in this volume were presented at the 5th International Parallel Processing Symposium, held April 30 through May 2, 1991, in Anaheim, California. The symposium was sponsored by the IEEE Computer Society.

The program committee met on December 6, 1990, and selected these 112 papers from 243 manuscripts submitted in response to the Call for Papers. The selection was based on originality, technical quality, and relevance to the theme of the Symposium. The submissions were refereed, and many of them represent reports of continuing research. It is anticipated that most of these papers will appear in more polished and complete form in scientific journals.

The program committee wishes to thank all who submitted for consideration and those who evaluated the submissions.

V. K. Prasanna Kumar
Program Chair

Prith Banerjee
Jacob Barhen
Keith Bromley
Mary Eshaghian
Jose Fortes
Kai Hwang
S. S. Iyengar
H. T. Kung
Dennis Parkinson
A. Paulraj
K. Wojtek Przytula
Sartaj Sahni
H. J. Siegel
Abe Waksman

Message from the Chair

Welcome to the Fifth International Parallel Processing Symposium.

This year's event is the culmination of many hours of work from a dedicated group of individuals consisting of V. K. Prasanna Kumar, Program Committee Chairman-USC; Sally Jelinek, Publicity-Electronic Design Associated; Bill Pitts, Treasurer-Toshiba America; Susamma Barua, Exhibits and Local Arrangements-CSUF and George Westrom-Steering Committee Chair, Odetics, Inc. representing the Orange County IEEE Computer Society. As a result of growing interest from both the international community and from the United States, the name of the Symposium has been changed to reflect this broader participation. Because of the growth in attendance, the location was moved to the larger facilities of the Disneyland Convention Center. With the growing interest in research activity related to Parallel processing, the number of papers submitted has almost doubled to 243 papers. At the Symposium, 112 of these manuscripts will be presented in the technical and poster sessions. The Symposium theme continues to feature a blend of theoretical and applied parallel processing applications. This year the Symposium is co-located with the Federal Computer Conference. The Parallel Processing commercial exhibits will be combined with the FCC exhibits which will have almost 300 booths. In addition, three formal tutorials will be presented by Professor H. J. Siegal of Purdue University, Professors L. Bic and A. Nicolau of UCI, and Dr. Harold Szu of the Naval Surface Weapons Research Center.

Special recognition is given to Dr. V. K. Prasanna Kumar for his excellent work as Program Chairman. Through Dr. Prasanna Kumar's leadership this year's Program Committee and Keynote Speakers consist of the leading scientists in the parallel processing field. Special thanks are given to Sally Jelinek for not only coordinating all the publicity and the advanced program, but also for the many additional administrative tasks she undertook. I would like to extend special thanks to Professor H. T. Kung of Carnegie Mellon University for his recommendations for this year's Symposium. Equally, special thanks go to our keynote speakers, Professor Kai Hwang of USC, Professor David Kuck of the University of Illinois, and Dr. Harold Stone of IBM, who bring to the Symposium a wealth of creative experiences in parallel processing which they will share with us.

Finally, I would like to extend our thanks to the many universities and industrial firms who support the Parallel Processing Symposium, particularly the USC-Dept of EE-Systems, Encore Computer Corporation, Alliant Computer Systems Corporation, Grumman Data Systems, Hughes Aircraft, Intel Corporation, JPL, Motorola Corporation, Rockwell International, Toshiba America Information Systems, UCI, CalTech, and CSUF.

Dr. Larry H. Canter
General Chairman

Acknowledgements

IPPS '91 could not have taken place without the contributions of many people, including:

Conference Chair

Larry Canter
Computer Systems Approach, Inc.

Program Committee

V. K. Prasana Kumar, Program Chair
University of Southern California

Prith Banerjee
University of Illinois

Jacob Barhen
JPL & Caltech

Keith Bromley
NOSC

Mary Eshaghian
Grumman Data Systems

Jose Fortes
Purdue University

Kai Hwang
USC

S. S. Iyengar
LSU

H. T. Kung
CMU

Dennis Parkinson
Queen Mary College & AMT

A. Paulraj
BEL & CDAC

K. Wojtek Przytula
Hughes

Sartaj Sahni
University of Florida

H. J. Siegel
Purdue University

Abe Waksman
AFOSR

Keynote Speakers

Professor Kai Hwang
Univ. of Southern Calif.

Professor David Kuck
University of Illinois

Dr. Harold Stone
IBM

Session Chairs

Hussein Alnuweiri
King Fahd University of Petroleum and Minerals

Joydeep Ghosh
Univ. of Texas at Austin

Wei K. (Kevin) Tsai
Univ. of Calif., Irvine

Alok Choudhary
Syracuse University

Dan Moldovan
Univ. of Southern Calif.

Michel Dubois
Univ. of Southern Calif.

Oscar Ibarra
Univ. of Calif., Santa Barbara

Milos Ercegovac
Univ. of Calif., Los Angeles

Dionisios Reisis
National Technical Univ., Athens

Herb Tesser
Grumman Data Systems

David Nassimi
New Jersey Institute of Technology

Greg Nash
Hughes Research Labs

Isaac Scherson
Univ. of Calif., Irvine

Marina Chen
Yale University

Ashfaq Khokhar
Univ. of Southern Calif.

David Okamura
Univ. of Southern Calif.

Table of Contents

Foreword v
Message from the Chair vi
Acknowledgements vii

Algorithms

Efficient Implementations of a Class of $\pm 2^b$
Parallel Computations on a SIMD Hypercube 2
D. Nassimi and Y.-D. Tsai
Triangulation in a Plane and 3-D Convex Hull on
Mesh-Connected Arrays and Hypercubes 10
J. A. Holey and O. H. Ibarra
The Effects of Communication Latency Upon Synchronization and
Dynamic Load Balance on a Hypercube 18
D. C. Marinescu and J. R. Rice
Improved Graph Computations on the Reconfigurable Mesh 26
D. I. Reisis
Design Guidelines for Parallel Algorithms
Using Continuous Job Profiles 30
L. F. Wilson and M. J. Gonzalez
An Inherently Fault Tolerant Sorting Algorithm 37
I.-L. Yen, F. Bastani, and E. Leiss
Performance Analysis of Algorithms on a
Message Passing Multiprocessor 43
A. Basu, S. Srinivas, K. G. Kumar,
A. Paulraj, and L. M. Patnaik
Routing Algorithms in Interval and Circular-Arc Networks 51
M. A. Sridhar and S. Goyal
Deterministic Timing Schema for Parallel Programs 56
A. C. Shaw
Generalized Formulation and Hypercube Algorithms for
Relaxation Labeling 64
E. Leung and X. Li
General Asynchrony Is Not Expensive for PRAMs 70
J. C. Becker, C. U. Martel, and A. Park
Fast Parallel Algorithms for Solving Triangular Systems of
Linear Equations on the Hypercube 76
O. H. Ibarra and M. H. Kim
Optimal EREW Parallel Algorithms for Connectivity,
Ear Decomposition and st-Numbering of Planar Graphs 84
H. Gazit
Scalability of Parallel Sorting on Mesh Multicomputers 92
V. Singh, V. Kumar, G. Agha, and C. Tomlinson
An Algorithm for Generating Node Disjoint
Routes in Kautz Digraphs 102
G. J. M. Smit, P. J. M. Havinga, and P. G. Jansen
A Neural-Type Parallel Algorithm for Fast Matrix Inversion 108
M. M. Polycarpou and P. A. Ioannou
Parallel Multiple Search 114
Z. Wen

SIMD Data Communication Algorithms for
Multiply Twisted Hypercubes 120
S.-Q. Zheng
Two EREW Algorithms for Parentheses Matching 126
S. Prasad and N. Deo
A Cost-Optimal Parallel Algorithm for the Parentheses
Matching Problem on an EREW PRAM 132
C. C.-Y. Chen and S. K. Das
A Parallel Approximation Algorithm for Solving
One-Dimensional Bin Packing Problems 138
J. O. Berkey and P. Y. Wang
Fully Normal Algorithms for Incomplete Hypercubes 144
V. K. Prabhala and N. A. Sherwani
Refining Algorithm Mappings for Linear Systolic Arrays 151
R. Varadarajan and B. Ravichandran
Analysis of Parallel Mixed-Mode Simulation Algorithms 155
R. D. Chamberlain and M. A. Franklin
Implementation of an Oversize Neural Network on DAP-510 161
S. N. Gupta, M. Zubair, C. E. Grosch
Performance of Parallel Spanning Tree Algorithms on Linear
Arrays of Transputers and Unix Systems 165
S. K. Das and C.-Q. Yang

Applications

Trade-Offs In Mapping FFT Computations onto
Fixed Size Mesh Processor Array 170
J. W. Jang and W. Przytula
Parallel Vision Integration on the
AMT Distributed Array Processor 178
S. T. Toborg and K. Hwang
Functional Characterization of Sensor Integration in
Distributed Sensor Networks 186
L. Prasad, S. S. Iyengar,
R. L. Kashyap, and R. N. Madan
Connected Components with Split and Merge 194
J. J. Kistler and J. A. Webb
Optical Content Addressable Memories for
Data/Knowledge Base Processing 202
S. Akyokus and P. B. Berra
Reconfigurable Mesh Algorithms for Image Shrinking,
Expanding, Clustering, and Template Matching 208
J.-F. Jenq and S. Sahni
Parallel Implementation of Image Understanding Tasks on
Gated-Connection Networks 216
D. B. Shu, J. G. Nash, and K. Kim
Special Features of a VLIW Architecture 224
A. Abnous and N. Bagherzadeh
A Neural Algorithm for Variable Thresholding of Images 228
Z.-P. Lo and B. Bavarian
Design of an Optical Content-Addressable Parallel Processor with
Applications to Fast Searching and Information Retrieval 234
A. Louri

A Parallel Processing Approach to
Incremental Conceptual Clustering 240
P. Wohl and T. W. Christopher
Analysis of Neighborhood Interaction in
Kohonen Neural Networks 246
Z.-P. Lo, M. Fujita, and B. Bavarian
On a Binary Relation Inference Network 250
K. P. Lam and C. J. Su
MCMR: A Multiple Rule Firing Production System Model 256
S. Kuo, D. Moldovan, and S. Cha
Guided Scheduling Schemes for Image Understanding Tasks for
Shared and Distributed Memory Multiprocessors 260
A. N. Choudhary and R. Ponnusamy
Optical Database Machine Using Spatial Light Rebroadcasters 265
S. M. Chung

Architectures

Design of a Clustered Multiprocessor for
Real-time Natural Language Understanding 270
R. F. DeMara and D. I. Moldovan
Prototype Implementation of a Highly Parallel
Dataflow Machine EM-4 278
S. Sakai, Y. Kodama, and Y. Yamaguchi
Proteus System Architecture and Organization 287
A. K. Somani, C. Wittenbrink, R. M. Haralick,
L. G. Shapiro, J.-N. Hwang, C.-H. Chen,
R. Johnson and K. Cooper
Development of a Mixed MIMD-SIMD Architecture
Using the AMT DAP 295
S. G. Skinner, P. M. Flanders, and A. G. Bale
Instruction Execution Trade-Offs for SIMD vs.
MIMD vs. Mixed Mode Parallelism 301
T. B. Berg and H. J. Siegel
An Effective Synchronization Network for
Large Multiprocessor Systems 309
W. T.-Y. Hsu and P.-C. Yew
Input/Output Operations for Hybrid
Data-Flow/Control-Flow Systems 318
P. Evripidou and J.-L. Gaudiot
A Linked List Cache Coherence Protocol:
Verifying the Bottom Layer 324
S. Gjessing, S. Krogdahl, and E. Munthe-Kaas
An Approach to Solve the Cache Thrashing Problem 330
Z. Fang, M. Lu, and H. Lin
Parallel Multi-Context Architecture with
High-Speed Synchronization Mechanism 336
K. Toda, K. Nishida, Y. Uchibori,
S. Sakai, and T. Shimada
Systolic Rank Updating and the Solution of
Non-Linear Equations 344
G. M. Megson

On the Design of Optimal Fault-Tolerant
Systolic Array Architectures 352
M. O. Esonu, A. J. Al-Khalili, and S. Hariri
From Algorithms to Parallel Architectures:
A Formal Approach 358
K. M. Elleithy and M. A. Bayoumi
Bounds on Performance of VLSI Processor Arrays 364
A. Nayak and N. Santoro
Systolic Array Architecture for Adaptive Eigenstructure
Decomposition of Correlation Matrices 371
S. Erlich and K. Yao
State Variable Model for a Class of Multiprocessor Systems 377
G. M. Chaudhry and J. S. Bedi
Parallel I/O Subsystems for Distributed-Memory Multicomputers 381
J. Ghosh and B. Agarwal
Amstrong II: A Loosely Coupled Multiprocessor with a
Reconfigurable Communications Architecture 385
P. M. Athanas and H. F. Silverman

Networks

Multidimensional Access Shared Memory
Parallel Processing Systems 390
I. D. Scherson
A Simple Algorithm to Route Arbitrary Permutations on
8-Input 5-Stage Shuffle/Exchange Network 398
K. Kim and C. S. Raghavendra
Characterization of Node Disjoint (Parallel) Path in Star Graphs 404
J.-S. Jwo, S. Lakshmivarahan, and S. K. Dhall
Efficient Randomized Routing on Clos Networks 410
A. Youssef
On Optimal Embeddings into Incomplete Hypercubes 416
A. Gupta, A. Boals, and N. Sherwani
Distributed Hierarchical Optimal Routing Using Aggregation/
Disaggregation and Decomposition/Composition Techniques 424
S.-W. Park and W. K. Tsai
A Note on Orthogonal Graphs 432
S. V. R. Madabhushi, S. Lakshmivarahan, and S. K. Dhall
Choosing the Interconnect of Distributed Memory Systems by
Cost and Blocking Behavior 438
W. K. Giloi and S. Montenegro
Optimum Embeddings of End-Around Meshes into
Pyramid Networks 445
A. Dingle and H. Barada
Link Augmented Binary (LAB) - Tree: Its Structure,
Routing and Fault Tolerance Properties 452
R. Mittal, B. N. Jain, and R. K. Patney
Congestion and Fault Tolerance of Binary
Tree Embeddings on Hypercube 458
K. Efe and K. Ramaiyer
Packed Exponential Connections—A Hierarchy of 2-D Meshes 464
W. W. Kirkman and D. Quammen
Cross-Cube: A New Fault Tolerant Hypercube-Based Network 471
E. Haq

Multi-level Hypercube Network 475
M. A. Aboelaze
A Resilient Decentralized Commit Protocol 481
S.-M. Yuan
Greedy Partitioning Strategy for Banyan-Hypercube Networks 487
A. Bellaachia and A. Youssef
High-Performance Parallel Pipelined Voting Networks 491
B. Parhami
Modeling and Performance Analysis for Processor-to-Processor Communications Unit Using a 100 Mb/s Optical Token Ring 495
T. Matsunaga, I. Oyaizu, E. Uozumi, T.Hoshiko, and Y. Kimura

Software

A Scheduling Algorithm for Parallelizable Dependent Tasks 500
K. P. Belkhale and P. Banerjee
Scheduling and Optimization for Multiprocessor Systems 507
C.-C. Hsu
A System for Algorithm-Architecture Mapping Based on Dependence Graph Matching and Hypergraphs 513
J. Li and L. H. Jamieson
A Message Segmentation Technique to Minimize Task Completion Time 519
S. Kim, S. S. Pande, D. P. Agrawal, and J. Mauney
Load Balancing on a Hypercube 525
J. Woo and S. Sahni
Scheduling Parallelizable Imprecise Computations on Multiprocessors 531
A. C. Yu and K.-J. Lin
Improved Algorithms for Load Balancing in Circuit-Switched Hypercubes 537
C. S. Raghavendra, S. Chalasani, and R. V. Boppana
Implementation of a Parallel Prolog Interpreter on Multiprocessors 543
L. V. Kale and B. Ramkumar
Design and Evaluation of a High Performance File System for Message Passing Parallel Computers 549
U. Nagaraj, U. S. Shukla, and A. Paulraj
Data Prefetching Strategies for Vector Cache Memories 555
J. W. C. Fu and J. H. Patel
Use of PARADISE: A Meta-Tool for Visualizing Parallel Systems 561
J. A. Kohl and T. L. Casavant
Chief: A Parallel Simultation Environment for Parallel Systems 568
J. D. Bruner, H. Cheong, A. Veidenbaum, and P.-C. Yew
Confining Imperative Languages for Parallel Processing 576
K. Dai and W. K. Giloi
A Geometrical Coding to Compile Affine Recurrence Equations on Regular Arrays 582
C. Mongenet, P. Clauss, and G. R. Perrin
A Load Sharing Interconnection Network for Hard Real-Time Systems 591
A. Avritzer, M. Gerla, and J. W. Carlyle

Loop Partitioning Unimodular Transformations for
Distributed Memory Multiprocessors 599
D. Kulkarni, K. G. Kumar, A. Basu, and A. Paulraj
Execution of Regular DO Loops on Asynchronous Multiprocessors 605
P. Ouyang
Process Migration in the GALAXY Distributed Operating System 611
P. K. Sinha, K. S. Park, X. Jia, K. Shimizu, and M. Maekawa
An Analysis of Recurrence Relations in Fortran
Do-loops for Vector Processing 619
C.-P. Chu and D. L. Carver
PRESHAK: A Generic Tool to Implement Application Specific
Message-Passing Communication Kernels for Concurrent Machines 626
A. K. Srivastava and S. C. Kshetramade
Parallelizing Tightly Nested Loops 630
K.-C. Kim and A. Nicolau
A Simulation of Demand-Driven Dataflow:
Translation from Lucid into MDC Language 634
G. K. Thiruvathukal and T. W. Christopher
PMACS: An Environment for Parallel Programming 638
B. Dehbonei, C. Laurent, N. Tawbi, and R. S. Kulkarni
An OR-Parallel and Restricted AND-Parallel,
Nonbacktracking Prolog Execution Model 642
I. Aybay and M. Baray

Panel
Moderator: D. Parkinson

How Do We Make Parallel Processing a Reality?:
Bridging the Gap Between Theory and Practice 648
Panelists: A. Brenner, R. Freund, R. S. Gaines,
R. Kelly, L. Lome, R. McAndrew, A. Nicolau, J. Patel,
T. Probert, J. Reif, J. L. C. Sanz, H. J. Siegel, and J. Webb

Author Index 654

ALGORITHMS

Efficient Implementations of a Class of $\pm 2^b$ Parallel Computations on a SIMD Hypercube

David Nassimi

CIS Department
New Jersey Institute of Technology
Newark, New Jersey 07102

Yuh-Dong Tsai

Dept. of CS, Math, Stat
University of Baltimore
Baltimore, Maryland 21201

Abstract

We identify an important class of parallel computations, called $\pm 2^b$ − *descend*, with an efficient implementation on a hypercube. Given the input $A[0 : N-1]$, a computation in this class consists of $\log N$ iterations. Iteration b, $b = \log N - 1, \ldots, 0$, computes the new value of each $A[i]$ as a function of $A[i]$, $A[i+2^b]$ and $A[i-2^b]$. We obtain a general algorithm for implementing any computation in this class in $O(\log N)$ time on a SIMD hypercube. Our general descend algorithm results in an efficient $O(\log N)$ implementation of Batcher's odd-even merge algorithm on a hypercube. The best previously known implementation of odd-even merge on a SIMD hypercube requires $O(\log^2 N)$ time.

1 Introduction

A Boolean n-cube, also called n-dimensional hypercube, is a parallel computer with $N = 2^n$ processors numbered $0, 1, \ldots, N-1$. A pair of processors i, j are directly connected iff the binary representaions of i and j differ in exactly one bit. Let $i = i_{n-1} i_{n-2} \ldots i_0$ for $i \in [0 : N-1]$, so

$$i \oplus 2^b = i_{n-1} \ldots i_{b+1} \overline{i_b} i_{b-1} \ldots i_0.$$

Processor i is connected to $\log N$ other processors: $i \oplus 2^b$, $0 \le b \le n-1$.

The hypercube interconnection has n *dimensions*. Dimension b is the connections between $N/2$ pairs of processors $i, i \oplus 2^b$. The $N/2$ processors with $i_b = 0$ and the $N/2$ processors with $i_b = 1$ each form a *subcube*, having the structure of an (n-1)-cube.

The processors can communicate only via the processor-to-processor interconnection. Two communication models of hypercube have been considered in the literature. In a SIMD model, all communications in a single step are restricted to a single dimension. In a MIMD model, this restriction does not exist.

There is a wide class of parallel algorithms which require processor i to communicate with processors $i \pm 2^b \bmod N$, $0 \le i \le N-1$, where b is an ineger determined by each stage of the algorithm.

Consider the complexity of performing a 2^b-permutation on a hypercube. (A 2^b-permutation of a sequence $a[0 : N-1]$ sends the element initially at position i to position $i+2^b \bmod N$. The inverse is called -2^b-permutation.) If we assume that element $a[i]$ is mapped onto processor i, then it is easy to see that at least $\log N - b$ routing steps are needed to perform a 2^b-permutation. (There are known $O(\log N - b)$ algorithms for performing a 2^b-permutation [8] on a hypercube.)

Considering the more general case where the elements may be mapped onto processors according to any one-to-one mapping π, the following bounds are known for performing a 2^b permutation on a hypercube:

- For a SIMD model, the lower bound of $\log N - b$ still holds [12].
- For a MIMD model, a 2^b permutation can be performed in $O(1)$ steps using an appropriate one-to-one mapping [5], [12].

The lower bound for a SIMD model suggests that if a parallel algorithm requires $f(N)$ communication steps of type $\pm 2^b$, it may require $O(f(N) \cdot \log N)$ steps on a hypercube.

A well-known example of a parallel algorithm requiring $\pm 2^b$ communications is Batcher's odd-even merge. To merge two sorted sequences of length $N/2$, this algorithm runs through $\log N$ iterations,

TH0363-2/91/0000/0002$01.00 © 1991 IEEE