

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53

Proceedings

Frontiers '96

The Sixth Symposium on the
Frontiers of Massively Parallel Computing

October 27–31, 1996
Annapolis, Maryland

Sponsored by
IEEE Computer Society

In cooperation with
NASA Goddard Space Flight Center
USRA/CESDIS



IEEE Computer Society Press
Los Alamitos, California

Washington • Brussels • Tokyo



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

Copyright © 1996 by The Institute of Electrical and Electronics Engineers, Inc.
 All rights reserved.

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. 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.

IEEE Computer Society Press Order Number PR07551
 IEEE Order Plan Catalog Number 96TB100062
 ISBN 0-8186-7551-9
 Microfiche ISBN 0-8186-7553-5
 ISSN 1088-4955

Additional copies may be ordered from:

IEEE Computer Society Press
 Customer Service Center
 10662 Los Vaqueros Circle
 P.O. Box 3014
 Los Alamitos, CA 90720-1314
 Tel: +1-714-821-8380
 Fax: +1-714-821-4641
 Email: cs.books@computer.org

IEEE Service Center
 445 Hoes Lane
 P.O. Box 1331
 Piscataway, NJ 08855-1331
 Tel: +1-908-981-1393
 Fax: +1-908-981-9667
 misc.custserv@computer.org

IEEE Computer Society
 13, Avenue de l'Aquilon
 B-1200 Brussels
 BELGIUM
 Tel: +32-2-770-2198
 Fax: +32-2-770-8505
 euro.ofc@computr.org

IEEE Computer Society
 Ooshima Building
 2-19-1 Minami-Aoyama
 Minato-ku, Tokyo 107
 JAPAN
 Tel: +81-3-3408-3118
 Fax: +81-3-3408-3553
 tokyo.ofc@computer.org

Editorial production by Penny Storms
 Cover by Kerry Bedford and Alex Torres
 Printed in the United States of America by KNI, Inc.



The Institute of Electrical and Electronics Engineers, Inc.

Contents

Message from the General Chair	ix
Message from the Program Chair	x
Conference Committee	xi
Referees	xiii

Session 1: Invited Speaker

From ASCI to Teraflops

John Hopson, Accelerated Strategic Computing Initiative (ASCI)

Session 2A: Scheduling 1

Gang Scheduling for Highly Efficient Distributed Multiprocessor Systems 4

H. Franke, P. Pattnaik, and L. Rudolph

Integrating Polling, Interrupts, and Thread Management..... 13

K. Langendoen, J. Romein, R. Bhoedjang, and H. Bal

A Practical Processor Design for Multithreading 23

M. Amamiya, T. Kawano, H. Tomiyasu, and S. Kusakabe

Session 2B: Routing

Analysis of Deadlock-Free Path-Based Wormhole Multicasting in
Meshes in Case of Contentions 34

E. Fleury and P. Fraigniaud

Efficient Multicast in Wormhole-Routed 2D Mesh/Torus Multicomputers:
A Network-Partitioning Approach..... 42

S-Y. Wang, Y-C. Tseng, and C-W. Ho

Turn Grouping for Efficient Multicast in Wormhole Mesh Networks 50

K-P. Fan and C-T. King

Session 3A: Applications and Algorithms

A³: A Simple and Asymptotically Accurate Model for Parallel Computation..... 60

A. Grama, V. Kumar, S. Ranka, and V. Singh

Fault Tolerant Matrix Operations Using Checksum and Reverse Computation..... 70

Y. Kim, J.S. Plank, and J.J. Dongarra

A Statistically-Based Multi-Algorithmic Approach for Load-Balancing
Sparse Matrix Computations 78

S. Nastea, T. El-Ghazawi, and O. Frieder

Session 3B: Petaflops Computing / Point Design Studies

Pursuing a Petaflop: Point Designs for 100 TF Computers Using PIM Technologies 88

P.M. Kogge, S.C. Bass, J.B. Brockman, D.Z. Chen, and E. Sha

Hybrid Technology Multithreaded Architecture 98

G. Gao, K.K. Likharev, P.C. Messina, and T.L. Sterling

The Illinois Aggressive Coma Multiprocessor Project (I-ACOMA)..... 106

J. Torrellas and D. Padua

1
2 **Panel Session—How Do We Break the Barrier to the Software Frontier?**

3 *Panel Chair: Rick Stevens, Argonne National Laboratory*

4
5 **Session 4: Invited Speaker**

6
7 **Session 5A: Scheduling 2**

8 Largest-Job-First-Scan-All Scheduling Policy for 2D Mesh-Connected Systems..... 118

9 *S.M. Yoo and H.Y. Youn*

10 Scheduling for Large-Scale Parallel Video Servers..... 126

11 *M-Y. Wu and W. Shu*

12 Effect of Variation in Compile Time Costs on Scheduling Tasks on
13 Distributed Memory Systems 134

14 *S. Darbha and S. Pande*

15
16 **Session 5B: SIMD**

17 Processor Autonomy and Its Effect on Parallel Program Execution..... 144

18 *D.M. Hawver and G.B. Adams III*

19 Particle-Mesh Techniques on the MasPar 154

20 *P. MacNeice, C. Mobarry, and K. Olson*

21 MIMD Programs on SIMD Architectures 162

22 *M-Y. Wu and W. Shu*

23
24 **Session 6A: I/O Techniques**

25 Intelligent, Adaptive File System Policy Selection..... 172

26 *T.M. Madhyastha and D.A. Reed*

27 An Abstract-Device Interface for Implementing Portable Parallel-I/O Interfaces..... 180

28 *R. Thakur, W. Gropp, and E. Lusk*

29 PMPIO - A Portable Implementation of MPI-IO 188

30 *S.A. Fineberg, P. Wong, B. Nitzberg, and C. Kuszmaul*

31 Disk Resident Arrays: An Array-Oriented I/O Library for Out-Of-Core
32 Computations..... 196

33 *J. Nieplocha and I. Foster*

34
35 **Session 6B: Memory Management**

36 Hardware-Controlled Prefetching in Directory-Based Cache Coherent Systems 206

37 *W. Hu and P. Xia*

38 Preliminary Insights on Shared Memory PIC Code Performance on the
39 Convex Exemplar SPP1000..... 214

40 *P. MacNeice, C.M. Mobarry, J. Crawford, and T.L. Sterling*

41 Scalability of Dynamic Storage Allocation Algorithms 223

42 *A. Iyengar*

43 An Interprocedural Framework for Determining Efficient Data
44 Redistributions in Distributed Memory Machines 233

45 *S.K.S. Gupta and S. Krishnamurthy*

1
2 **Panel Session—Petaflops Alternative Paths**

3 *Panel Chair: Paul Messina, California Institute of Technology*

4
5 **Session 7: Invited Speaker**

6 Independence Day

7 *Steven Wallach, HP-Convex*

8
9 **Session 8A: Synchronization**

10 A Fair Fast Distributed Concurrent-Reader Exclusive-Writer Synchronization 246

11 *T.J. Johnson and H. Yoon*

12 Lock Improvement Technique for Release Consistency in Distributed
13 Shared Memory Systems 255

14 *S.S. Fu and N-F. Tzeng*

15 A Quasi-Barrier Technique to Improve Performance of an Irregular Application 263

16 *H.V. Shah and J.A.B. Fortes*

17
18 **Session 8B: Networks**

19 Performance Analysis and Fault Tolerance of Randomized Routing on
20 Clos Networks 272

21 *M. Bhatia and A. Youssef*

22 Performing BMMC Permutations in Two Passes through the Expanded
23 Delta Network and MasPar MP-2 282

24 *L.F. Wisniewski, T.H. Cormen, and T. Sundquist*

25 Macro-Star Networks: Efficient Low-Degree Alternatives to Star Graphs
26 for Large-Scale Parallel Architectures 290

27 *C-H. Yeh and E. Varvarigos*

28
29 **Session 9A: Performance Analysis**

30 Modeling and Identifying Bottlenecks in EOSDIS 300

31 *J. Demmel, M.Y. Ivory, and S.L. Smith*

32 Tools-Supported HPF and MPI Parallelization of the NAS Parallel Benchmarks 309

33 *C. Cl  men  on, K.M. Decker, V.R. Deshpande, A. Endo,*

34 *J. Fritscher, P.A.R. Lorenzo, N. Masuda, A. M  ller,*

35 *R. R  hl, W. Sawyer, B.J.N. Wylie, and F. Zimmerman*

36 A Comparison of Workload Traces from Two Production Parallel Machines 319

37 *K. Windisch, V. Lo, D. Feitelson, R. Moore, and B. Nitzberg*

38 Morphological Image Processing on Three Parallel Machines 327

39 *M.D. Theys, R.M. Born, M.D. Allemang, and H.J. Siegel*

40
41
42 **Session 9B: Petaflops Computing / Point Design Studies**

43 MORPH: A System Architecture for Robust High Performance Using
44 Customization (An NSF 100 TeraOps Point Design Study) 336

45 *A.A. Chien and R.K. Gupta*

46 Architecture, Algorithms and Applications for Future Generation
47 Supercomputers 346

48 *V. Kumar, A. Sameh, A. Grama, and G. Karypis*

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.