Databases and Parallel Architectures

Proceedings

PARBASE-90

International Conference on Databases, Parallel Architectures, and Their Applications

Edited by N. Rishe, S. Navathe, and D. Tal

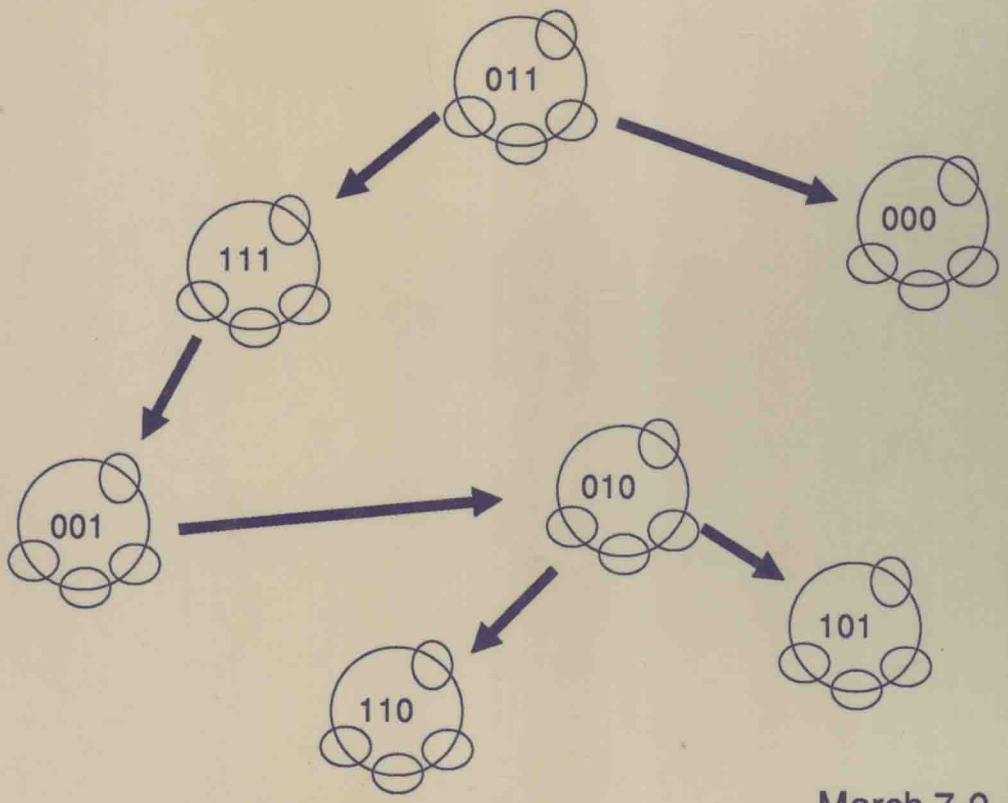

March 7-9, 1990
Miami Beach, Florida

Sponsored by
Florida International University
in cooperation with IEEE and Euromicro

IEEE Computer Society Press

Institute of Electrical and Electronics Engineers, Inc.

PARBASE-90

International Conference on Databases, Parallel Architectures, and Their Applications

March 7-9, 1990 Miami Beach, Florida

Edited by N. Rishe, S. Navathe, and D. Tal

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

Cover design by Wally Hutchins

Printed in United States of America

IEEE Computer Society Press Order Number 2035
Library of Congress Number 89-80524
IEEE Catalog Number 90CH2728-4
ISBN 0-8186-2035-8 (paper)
ISBN 0-8186-6035-X (microfiche)
SAN 264-620X

Additional copies may be ordered from:

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

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

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

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

PARBASE-90

International Conference on Databases, Parallel Architectures, and Their Applications

Databases and Parallel Architectures

Proceedings

PARBASE-90

International Conference on Databases, Parallel Architectures, and Their Applications

Edited by N. Rishe, S. Navathe, and D. Tal

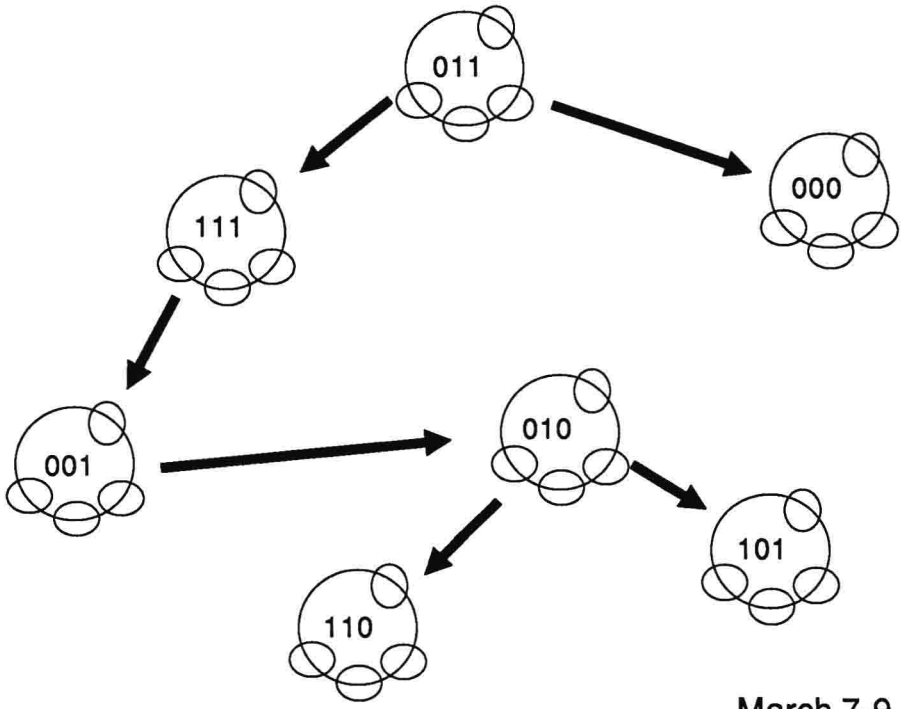

March 7-9, 1990
Miami Beach, Florida

Sponsored by
Florida International University
in cooperation with IEEE and Euromicro

IEEE Computer Society Press

Institute of Electrical and Electronics Engineers, Inc.

PARBASE-90: International Conference on Databases, Parallel Architectures, and their Applications, March 6-9, 1990, at Clarion Castle Hotel, Miami Beach, Florida, USA

Sponsored by Florida International University in cooperation with IEEE and Euromicro.

The conference provides a forum for exchange of information between scientists and engineers in the fields of theory, design and applications of databases and of parallel computer architectures. The technical program includes three keynote talks by D. DeGroot, C.A.R. Hoare, and G. Wiederhold; two panel sessions; 33 regular papers selected from 100 candidates, 50 short papers selected out of 150 candidates; poster papers; and tutorials.

ORGANIZING COMMITTEES

General Chair of the Conference
D. Tal, FIU

Steering Committee
N. Rishe (Chair), FIU
S. Navathe, University of Florida
J. Navlakha, FIU
D. Tal, FIU
N. Wiseman, Cambridge University (U.K.)

European Chair
N. Wiseman, Cambridge University (U.K.)

Japan Chair
H. Kunii, Ricoh Company, Japan

Local Arrangements Chair
M. Milani, FIU

Industrial Chairs
D. Barton, FIU
R. Ege, FIU

Publisher Relations Chair
B. Kraynek, FIU

Tutorial Chair
B. Furht, MODCOMP

Panel Chairs
D. Tabak, George Mason University
N. Prabhakaran, FIU

Publicity Chairs
A. Pelin, FIU
N. Prabhakaran, FIU

Conference Administration
G. Garrido, FIU, Conference Administrator
A. Dutt, FIU, Conference Assistant

Sponsoring Institution:
Florida International University —
M. Maidique, FIU President
J. Navlakha, FIU SCS Director

PROGRAM COMMITTEE

Chairpersons:
N. Rishe, FIU
S. Navathe, University of Florida

Program Committee members:
N. Adam, Rutgers University
M. Ahuja, Ohio State University
C.K. Baru, University of Michigan
B. Berra, Syracuse University
B. Bhargava, Purdue University
P.P.S. Chen, Louisiana State University
J. Comfort, FIU
K. Dittrich, University of Karlsruhe (FRG)
R. Elmasri, University of Houston
I. Erenyi, Hungarian Academy of Science
E. Fernandez, Florida Atlantic University
O. Frieder, Bell Communications Research
B. Furht, MODCOMP
S. Gadia, Iowa State University
E. Gudes, Ben-Gurion University (Israel)
H.V. Jagadish, AT&T Bell Laboratories
H. Kangassalo, University of Tampere (Finland)
Q. Li, FIU
P. Liu, University of Miami
R. Liuzzi, Rome Air Development, New York
A. Pelin, FIU
N. Prabhakaran, FIU
C. Pu, Columbia University
A. Sheth, Bell Communications Research
P. Shoval, Ben-Gurion University (Israel)
G.M. Silberman, Carnegie-Mellon University
A. Solberg, University of Trodheim (Norway)
D. Tabak, George Mason University
D. Tal, FIU
S. Urban, Arizona State University
R. Varadarajan, University of Florida
N. Wiseman, Cambridge University (U.K.)
N. Wolfe, University of Florida

General Chairman's Message

PARBASE-90 is a two-track International Conference on Databases, Parallel Architectures, and their Applications, taking place in Miami Beach, Florida. The conference strives to encourage exchange of information between scientists and engineers in the fields of theory, design and applications of databases and parallel computer architectures. This conference, although the first one of its kind, has attracted vast attention, as evidenced by the number of papers submitted from five continents.

The factor that contributes most to the success of the conference is the quality and diversity of the program. The scientific and technical program includes three keynote addresses: "Speculating and Throttling in Parallel Architecture," by D. DeGroot, Texas Instruments, "Future Architectures for Information Processing Systems," by G. Wiederhold, Stanford University, and "A Theory of Conjunction and Concurrency," by C.A.R. Hoare, University of Oxford; two panel sessions: "Parallel Architecture and Databases," chaired by D. Tabak, George Mason University, and "Databases," chaired by N. Prabhakaran, FIU, and moderated by S. Navathe, University of Florida; eleven regular paper sessions, seven short paper sessions, and poster paper sessions; and five tutorials intended for a broader audience.

Coordinating a conference of this magnitude is a difficult endeavor. Thus, we are most grateful for the support of many devoted persons who made this event possible. We thank all the program committee members, their colleagues, and external referees who served as additional reviewers and spent so much time and energy. In particular, we wish to express our gratitude to N. Rishe and S. Navathe (Program Chairpersons) and N. Wiseman (European Chair) for their outstanding efforts and contribution to the preparation of this conference. Special thanks go to M. Milani (Local Arrangements Chair), D. Barton and R. Ege (Industrial Chairs), B. Kraynek (Publisher Relations), D. Tabak and N. Prabhakaran (Panel Chairs), B. Furht (Tutorial Chair), A. Pelin (Publicity Chair), and H. Kunii (Japan Chair), and the Steering Committee chaired by N. Rishe.

We are most grateful for the financial help and support of Florida International University and its President M. Maidique, and the continuous help and advice of J. Navlakha, Director, FIU School of Computer Science. Our gratitude goes to all our colleagues from the School of Computer Science, FIU. Also, we thank G. Garrido (Conference Administrator) and A. Dutt (Conference Assistant) for their invaluable administrative assistance.

Last but not least, we owe thanks to all those who submitted papers for consideration and the delegates participating in PARBASE-90.

To conclude, we expect the PARBASE-90 conference to promote novel technology, open new research areas, and benefit all the participants. We hope that PARBASE-90 participants will find the conference enjoyable, stimulating, and a source of inspiration for new ideas.

Doron Tal
School of Computer Science
FIU, Miami, FL 33199
E-mail: tald@servax.bitnet

Program Chairmen's Message

We are pleased to present in these proceedings the full-length papers, short papers, selected formal poster papers, keynote papers, and panel position papers which are to be presented at PARBASE-90 — The First International Conference on Databases, Parallel Architectures and their Applications. The conference is sponsored by Florida International University. We were happy to receive the cooperation of IEEE and Euromicro for this conference.

We had an impressive collection of 100 full-length papers and 83 short papers submitted to this conference. All papers were refereed anonymously, without disclosing the names or affiliations of the authors. Each full paper was reviewed by three program reviewers — typically, two of whom were program committee members. We took the liberty of using some of the authors as referees for other authors' papers. We also called upon many experts in the respective fields to assist us with the refereeing process. The reviews evaluated primarily on originality, quality and relevance to the conference topics. The program committee met in Miami on November 11-12, 1989, and selected 33 papers for inclusion in these proceedings as full papers. All the remaining 150 papers were then considered as potential short papers; and out of these, 50 were selected for inclusion in the proceedings as short papers.

We have not made any deliberate effort to separate the papers into parallel architectures and databases. It was evident from the submissions that it is not possible to divide these two areas along tight boundaries. We hope that this conference will bring together researchers from both disciplines and improve the cross-fertilization of these two technologies. It was interesting to find at the end that the sessions fell into two parallel tracks — one mostly dealing with database issues and the other mostly dealing with architecture issues.

In addition, the conference has two poster sessions, keynote addresses by D. DeGroot, C.A.R. Hoare, and G. Wiederhold, and two panels on database technology and parallel architecture technology.

We wish to take this opportunity to thank the various officials of Florida International University for making this conference possible. We are particularly pleased to thank D. Tal, the Conference Chairman, for his extensive help to the program committee and for the numerous effort of the organization of this conference. Our special thanks go to all the authors, program committee members and external referees who really shaped the content of the program of this conference.

We hope that with the encouraging support of the attendees at this conference we will be able to have the next one in this series soon.

Naphtali Rishe
School of Computer Science
FIU, Miami, Fla. 33199
E-mail: rishen@servax.bitnet

Shamkant B. Navathe
Department of Computer Science
University of Florida
Gainesville, FL 32611
E-mail: sham@ufl.edu

Table of Contents

Short Paper Session SB2: Parallel Algorithms and OS Scheduling

Short Paper Session SC2: Architectures and Algorithms for Parallel and Distributed Databases II

Formal Poster Papers

Keynote Address

D. DeGroot

Throttling and Speculating on Parallel Architectures

(Abstract)

Doug DeGroot
Texas Instruments
Computer Science Center
degroot@csc.ti.com

Two of the lesser studied aspects of parallel processing are those of throttling and speculating. Loosely speaking, throttling may be viewed as a way of keeping parallelism under control, while speculating is a way of letting it get a bit out of control. It is frequently possible to significantly increase parallelism through less throttling or more speculation, and in certain application domains, we may want to do some of both. (In fact, there is at least one parallel computation model that *requires* both if any speedup is likely to occur.) Unfortunately, however, the greater parallelism due to a lessening of throttling or to speculating with the parallelism may lead to either significantly increased execution speedup or to significant execution slowdown, and surprisingly, it is only sometimes possible to predict ahead of time which will occur.

Throttling of parallelism is of importance to dynamic parallelism models in which the number and sizes of possible parallel code segments (tasks, processes, threads, etc.) are unknown at compile time, and which, if left uncontrolled, may lead to explosive parallelism with resulting slowdown rather than speedup. The primary goal of throttling is preventing a dynamically unfolding parallel execution from creating so many parallel code segments that either: 1) the system runs out of resources with which to manage the parallel segments, 2) the system begins to thrash as a result of the increased process management and communication requirements, or 3) the system simply wastes CPU time creating extra code segments which end up being executed sequentially anyway. This aspect of throttling then is responsible for ensuring that "too much" parallelism does not occur; it allows all processors to become busy, sets up a pool of "spare" work, and then "shuts down" the creation of parallelism until more work is needed to replace previously completed work, at which point it turns creation back on. It is a bit related to the job of the medium-term scheduler of a multiprogramming system whose job is to control the degree of multiprogramming for similar reasons.

Another goal of throttling is properly controlling the reduction of granularity in self-decomposing parallel procedures, although this aspect of throttling is not yet well understood. Here, the throttler is responsible for ensuring that tasks of too small a size do not decompose into two or more tasks of even smaller size, resulting in what has been called "embarrassing parallelism." This aspect of throttling is difficult to examine at the kernel-benchmark level and must be examined within the overall applications-level environment.

Several software approaches to throttling have been explored recently, at both the system level and the application level; most report more failures in solving the problem than successes. Certainly it can be easily demonstrated that programmer control of throttling is impractical. A significantly smaller number of hardware approaches to throttling have been examined; these too have met with occasional failure. Consequently, it is generally agreed that only appropriate combinations of both hardware and system-level software control appear suitable for effecting throttling.

Speculative parallelism involves the parallel execution of code segments which are not yet known to be required. It can be exploited in two basic forms: parallel search and parallel test. Parallel search techniques have been studied for years, but usually not within the speculative framework. "All solutions" parallel searches, as typified by many parallel relational or logic database query schemes, for example, are usually non-speculative. On the other hand, when only one solution to a problem or query is

CH2728-4/90/0000/0002$01.00 © 1990 IEEE