

Edited by G. Goos and J. Hartmanis

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



Pierre America (Ed.)

Parallel Database Systems

PRISMA Workshop

Noordwijk, The Netherlands, September 24-26, 1990

Proceedings

Springer-Verlag

Berlin Heidelberg New York

London Paris Tokyo

Hong Kong Barcelona

Budapest

Series Editors

Gerhard Goos
GMD Forschungsstelle
Universität Karlsruhe
Vincenz-Priessnitz-Straße 1
W-7500 Karlsruhe, FRG

Juris Hartmanis
Department of Computer Science
Cornell University
Upson Hall
Ithaca, NY 14853, USA

Volume Editor

Pierre America
Philips Research Laboratories
P.O. Box 80.000, 5600 JA Eindhoven, The Netherlands

CR Subject Classification (1991): H.2.4, C.1.2, D.1.3

ISBN 3-540-54132-2 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-54132-2 Springer-Verlag New York Berlin Heidelberg

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 other ways, and storage in data banks. Duplication of this publication or parts thereof is only permitted under the provisions of the German Copyright Law of September 9, 1965, in its current version, and a copyright fee must always be paid. Violations fall under the prosecution act of the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1991
Printed in Germany

Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr.
2145/3140-543210 - Printed on acid-free paper

Preface

This volume constitutes the proceedings of the Workshop on Parallel Database Systems organized by the PRISMA (Parallel Inference and Storage Machine) Project, and held in Noordwijk on the Dutch North Sea coast, September 24–26, 1990. It contains contributions from invited speakers as well as from the organizing PRISMA Project. The invited speakers, all recognized experts in their fields, present their views on several subjects that are closely related to the title of the workshop. The PRISMA papers, taken together, give an in-depth overview of the PRISMA system. This system is based on a parallel machine, where the individual processors each have their own local memory and communicate with each other over a packet-switched network. On this machine a parallel object-oriented programming language, POOL-X, has been implemented, which provides dedicated support for database systems as well as general facilities for parallel programming. This POOL-X system then serves as a platform for a complete relational main-memory database management system, which uses the parallelism of the machine to speed up significantly the execution of database queries. The presentation of the PRISMA system, together with the invited papers, gives a broad overview of the state of the art in parallel database systems.

The workshop marked the end of the PRISMA Project, which started in October 1986. The project was a cooperation of Philips Research Laboratories, Eindhoven, the Netherlands, with the Centre for Mathematics and Computer Science (CWI) in Amsterdam and the Universities of Amsterdam, Twente, Leiden, and Utrecht. It was sponsored by the Dutch 'Stimuleringsprojectteam Informatica-onderzoek' (SPIN).

The contributions were selected by the programme committee, consisting of

Pierre America (Philips Research Laboratories, chair)
Peter Apers (University of Twente)
Bob Hertzberger (University of Amsterdam)
Martin Kersten (CWI, Amsterdam)

These proceedings were edited by Pierre America.

I would like to take this opportunity to thank the many persons who have contributed to PRISMA, both within the project and through interaction from related research. They are too numerous to mention in this preface, but many of them are authors of articles in these proceedings or are referenced therein. One person, however, deserves special mention here: Fred Robert, who diligently took care of logistic arrangements throughout PRISMA and this workshop.

February 1991

H. H. Eggenhuisen
PRISMA Project Leader

Contents

Invited Presentations

Languages for Parallel Programming	3
Henri Bal (<i>Free University of Amsterdam</i>)	
Relational Algebra Operations	24
Kjell Bratbergsengen (<i>University of Trondheim</i>)	
An Overview of Parallel Strategies for Transitive Closure on Algebraic Machines	44
Filippo Cacace (<i>Politecnico di Milano</i>), Stefano Ceri (<i>Università di Modena</i>), Maurice Houtsma (<i>University of Twente</i>)	
Data Management with Massive Memory: A Summary	63
Hector Garcia-Molina, Robert Abbott, Christopher Clifton, Carl Staelin, Kenneth Salem (<i>Princeton University</i>)	
An Approach to Implement Dynamically Defined Complex Objects	71
Theo Härder (<i>University of Kaiserslautern</i>)	
Experience with MIMD Message-Passing Systems: Towards General Purpose Parallel Computing	99
Anthony Hey (<i>University of Southampton</i>)	
Parallel Associative Processing for Knowledge Bases	112
Simon Lavington (<i>University of Essex</i>)	
Performance Estimates of a Join	124
Suresh Patel (<i>ICL, Manchester</i>)	
The EDS Parallel Relational Database System	149
Paul Watson, Paul Townsend (<i>ICL, Manchester</i>)	

PRISMA Contributions

PRISMA, a Platform for Experiments with Parallelism	169
Peter Apers (<i>University of Twente</i>), Bob Hertzberger (<i>University of Amsterdam</i>), Ben Hulshof, Hans Oerlemans (<i>Philips Research, Eindhoven</i>), Martin Kersten (<i>CWI, Amsterdam</i>)	
The PRISMA Experiment	181
Peter Apers (<i>University of Twente</i>), Bob Hertzberger (<i>University of Amsterdam</i>), Ben Hulshof (<i>Philips Research, Eindhoven</i>)	
Back-End Aspects of a Portable POOL-X Implementation	193
Marcel Beemster (<i>University of Amsterdam</i>)	
Logging and Recovery in PRISMA	229
Carel van den Berg, Martin Kersten (<i>CWI, Amsterdam</i>)	
Parallel Handling of Integrity Constraints	242
Paul Grefen, Jan Flokstra, Peter Apers (<i>University of Twente</i>)	

Performance Analysis of a Dynamic Query Processing Scheme	259
Martin Kersten, Satish Shair-Ali, Carel van den Berg (<i>CWI, Amsterdam</i>)	
Evaluation of a Communication Architecture by Means of Simulation	275
Henk Muller (<i>University of Amsterdam</i>)	
A Parallel Implementation of the EQUIP Expert System	294
Jan Willem Spee (<i>PTT Research, Groningen</i>)	
POOL-X and its Implementation	309
Juul van der Spek (<i>Philips Research, Eindhoven</i>)	
Automatic Recovery of a Parallel Stable File System	345
Sun Chengzheng, Bob Hertzberger (<i>University of Amsterdam</i>)	
Duplicates and Translation of Nested SQL Queries into XRA	354
Nanda Verbrugge (<i>CWI, Amsterdam</i>)	
The POOMA Architecture	365
Marnix Vlot (<i>Philips Research, Eindhoven</i>)	
The POOMA Operating System	396
Rogier Wester, Ben Hulshof (<i>Philips Research, Eindhoven</i>)	
Parallel Query Execution in PRISMA/DB	424
Annita Wilschut, Peter Apers, Jan Flokstra (<i>University of Twente</i>)	