

---

---

**PARALLEL, OBJECT-ORIENTED,  
AND ACTIVE  
KNOWLEDGE BASE SYSTEMS**

---

# The Kluwer International Series on ADVANCES IN DATABASE SYSTEMS

Series Editor  
**Ahmed K. Elmagarmid**

*Purdue University  
West Lafayette, IN 47907*

***Other books in the Series:***

DATABASE CONCURRENCY CONTROL: Methods, Performance, and Analysis  
*by Alexander Thomasian*

ISBN: 0-7923-9741-X

TIME-CONSTRAINED TRANSACTION MANAGEMENT: Real-Time Constraints in  
Database Transaction Systems

*by Nandit R. Soparkar, Henry F. Korth, Abraham Silberschatz*

ISBN: 0-7923-9752-5

SEARCHING MULTIMEDIA DATABASES BY CONTENT

*by Christos Faloutsos*

ISBN: 0-7923-9777-0

REPLICATION TECHNIQUES IN DISTRIBUTED SYSTEMS

*by Abdelsalam A. Helal, Abdelsalam A. Heddaya, Bharat B. Bhargava*

ISBN: 0-7923-9800-9

VIDEO DATABASE SYSTEMS: Issues, Products, and Applications

*by Ahmed K. Elmagarmid, Haitao Jiang, Abdelsalam A. Helal, Anupam Joshi, Magdy Ahmed*

ISBN: 0-7923-9872-6

DATABASE ISSUES IN GEOGRAPHIC INFORMATION SYSTEMS

*by Nabil R. Adam and Aryya Gangopadhyay*

ISBN: 0-7923-9924-2

INDEX DATA STRUCTURES IN OBJECT-ORIENTED DATABASES

*by Thomas A. Mueck and Martin L. Polaschek*

ISBN: 0-7923-9971-4

INDEXING TECHNIQUES FOR ADVANCED DATABASE SYSTEMS

*by Elisa Bertino, Beng Chin Ooi, Ron Sacks-Davis, Kian-Lee Tan, Justin Zobel, Boris Shidlovsky  
and Barbara Catania*

ISBN: 0-7923-9985-4

MINING VERY LARGE DATABASES WITH PARALLEL PROCESSING

*by Alex A. Freitas and Simon H. Lavington*

ISBN: 0-7923-8048-7

DATA MANAGEMENT FOR MOBILE COMPUTING

*by Evaggelia Pitoura and George Samaras*

ISBN: 0-7923-8053-3

---

---

**PARALLEL, OBJECT-ORIENTED,  
AND ACTIVE KNOWLEDGE BASE  
SYSTEMS**

*by*

**Ioannis Vlahavas**

*and*

**Nick Bassiliades**

*Department of Informatics  
Aristotle University of Thessaloniki  
Greece*



SPRINGER SCIENCE+BUSINESS MEDIA, LLC

ISBN 978-1-4419-5029-1

ISBN 978-1-4757-6134-4 (eBook)

DOI 10.1007/978-1-4757-6134-4

---

**Library of Congress Cataloging-in-Publication Data**

A C.I.P. Catalogue record for this book is available  
from the Library of Congress.

---

*The publisher offers discounts on this book when ordered in bulk quantities. For  
more information contact: Sales Department, Springer Science+Business Media, LLC*

**Copyright © 1998 by Springer Science+Business Media New York**

**Originally published by Kluwer Academic Publishers in 1998**

All rights reserved. No part of this publication may be reproduced, stored in a  
retrieval system or transmitted in any form or by any means, mechanical, photo-  
copying, recording, or otherwise, without the prior written permission of the  
publisher, Springer Science+Business Media, LLC

*Printed on acid-free paper.*

To my wife Chrysoula and my daughters Iro and Danae

*Ioannis Vlahavas*

To my fiancée Evdokia and my parents

*Nick Bassiliades*

# Contents

<b>List of Figures/Tables</b>	<b>ix</b>
<b>Foreword</b>	<b>xi</b>
<b>Preface</b>	<b>xiii</b>
<b>1. INTRODUCTION</b>	<b>1</b>
Basic Terminology	5
<b>PART I. KNOWLEDGE BASE SYSTEMS</b>	<b>11</b>
<b>2. DEDUCTIVE AND ACTIVE DATABASES</b>	<b>13</b>
Introduction: The Rule Spectrum	13
Deductive Databases	15
Active Databases	20
<b>3. INTEGRATION OF MULTIPLE RULE TYPES</b>	<b>27</b>
Unification of Production and Deductive Rule Semantics	27
Integration of ECA Rules in Declarative Rule Conditions	28
Integration of Declarative Rules in Active Databases	30
<b>4. AN ACTIVE OBJECT-ORIENTED KNOWLEDGE BASE SYSTEM</b>	<b>37</b>
The System Architecture	37
The Rule Language	39
Integration of Declarative Rules	41
Condition Compilation and Matching	46

Optimization and Extensibility 54

**PART II. PARALLEL DATABASE AND KNOWLEDGE BASE SYSTEMS 61**

**5. PARALLEL DATABASE SYSTEMS 63**

Introduction 63

Parallel Query Execution 65

Parallel Database System Architectures 71

Data Partitioning 74

Parallel Object-Oriented Database Systems 78

**6. PARALLEL KNOWLEDGE BASE SYSTEMS 83**

Introduction 83

Parallel (Database) Production Systems 84

Parallel Deductive Database Systems 97

Parallel Active Database Systems 101

**7. A PARALLEL OBJECT-ORIENTED KNOWLEDGE BASE SYSTEM 105**

The Parallel Object-Oriented Data Model 105

The Abstract Machine 108

Parallel Query Processing 110

Parallel Rule Evaluation 118

**8. CONCLUSIONS AND FUTURE DIRECTIONS 135**

**APPENDIX 139**

**References 141**

**Index 149**

## List of Figures/Tables

Figure 1. The rule spectrum and various rule integration schemata	14
Figure 2. The architecture of the DEVICE system	38
Figure 3. A sample complex event network	47
Figure 4. A negative event example	51
Figure 5. Hybrid memories of inter-object events	56
Figure 6. Sharing of conflict sets among the rule managers of DEVICE	59
Figure 7. Shared-everything architecture	71
Figure 8. Shared-nothing architecture	72
Figure 9. Hybrid architecture	73
Figure 10. Master-slave architecture	76
Figure 11. Parallel query execution	76
Figure 12. Non-uniform partitioning	77
Figure 13. A sample dependency graph	90
Figure 14. A cyclic read-write interference	90
Figure 15. A simple write-write interference	90
Figure 16. Instance hierarchy of PRACTIC objects	106
Figure 17. The hierarchical architecture of the abstract PRACTIC machine	108
Figure 18. A sample database schema	111
Figure 19. Horizontal partitioning of a class	113
Figure 20. Pipelined method execution	114
Figure 21. Hierarchical query execution	116
Figure 22. Hierarchical query execution with overlapping	117
Figure 23. Main loop of rule execution at each local rule manager	125
Figure 24. Auxiliary procedures for rule priority lock management	126
Figure 25. Auxiliary procedures for parallel rule processing	127
Table 1. ECA rule types handled by PRACTIC <sup>KB</sup>	131

# Foreword

It is a pleasure to me to commend this book. It brings together ideas about a number of different ways of organising computation that are often kept in separate compartments, namely Rule-based, Object-Oriented and Parallel computation. Each of them has seen great developments over the past two decades, but no one of them is a panacea, to be used to the exclusion of the others. This book compares their strengths and weaknesses and then makes a very clever synthesis of them to tackle the demanding problem of Knowledge-base management.

The history of database management systems (DBMS) has been a history of moving functionality from individual application programs in towards the shared kernel of the supporting system. First came indexing techniques for files, then query optimisation and then enforcement of integrity constraints. Constraints are themselves a form of high-level rule or knowledge, and this book argues that the maintenance and execution of other forms of rules (including Deductive rules, Production rules, Active rules) needs also to be done by a central shared facility, just as in a DBMS. Crucially, databases keep decisions on the physical details of data storage separate from the logical structure of data, and thus the KBMS kernel must also be free to change the physical details of rule storage and representation, so as to allow efficient, possibly parallel execution. This book explores ways to do this.

There is much to ponder on here, and we are still some years away from its widespread commercial application (although this book does discuss applications and is a step on the way towards it). As someone involved in database research for many years, I found this book a source of forward-thinking and stimulating ideas, and I am sure the reader will find this also!

Prof Peter M.D. Gray  
University of Aberdeen, Scotland  
October, 1997

## Preface

Modern data intensive real-world applications, such as data warehousing, data mining, information retrieval, expert database systems, network management, etc. strive for advanced data representation, querying and programming facilities, in order to capture the increasing demand for efficient, automated, tolerant, intelligent and really useful information systems. Such information systems can only be supported by application developing tools that provide for complex representation and efficient processing of knowledge.

Knowledge Base Systems are an integration of conventional database systems with Artificial Intelligence techniques. Knowledge Base Systems provide inference capabilities to the database system by encapsulating the knowledge of the application domain within the database. Furthermore, Knowledge Base Systems provide sharing, ease of maintenance, and reusability of this knowledge which is usually expressed in the form of high-level, declarative rules, such as production and deductive rules.

However, the enormous amount and complexity of data and knowledge to be processed by these systems imposes the need for increased performance and expressiveness from the Knowledge Base System. The problems associated with the large volumes of data are mainly due to the sequential data processing and the inevitable input/output bottleneck. In order to avoid this bottleneck parallel database systems have emerged to speed-up data intensive applications.

Furthermore, the synchronous, sequential execution of large numbers of rules leads to unnecessarily many inferencing cycles that slow down Knowledge Base Systems. Parallel rule-based systems try to speed-up rule processing by executing asynchronously the various phases of rule evaluation in multiprocessor environments. Finally, the decision about the applicability of a certain piece of knowledge to a certain information state requires a large amount of pattern matching and control synchronization that can be distributed in a multiprocessor environment.

On the other hand, the structure complexity of the data and data manipulating programs, along with the impedance mismatch between the programming languages

and the relational database management systems led to the advent of Object-Oriented Database systems, an intersection of object-oriented ideas and conventional databases. Object-Oriented Databases reduce the “semantic gap” between real world concepts and data representation models. This one-to-one mapping helps the development of complex applications, such as CAD/CAM, simulation, graphical user interfaces, etc. Object-Oriented Databases encapsulate within the database system both data and programs, with advantages such as program re-use, modularization, and ease of maintenance.

The object-oriented model offers a uniform, extensible and re-usable data and program representation that seems a promising solution for the integration of databases and knowledge-based systems. This book presents such an approach to Knowledge Base Systems: A Parallel Knowledge Base System that is built on top of a Parallel Active Object-Oriented Database System.

In the first part of the book, we discuss extensively the various attempts to integrate one or more rule types into databases in order to provide inferencing capabilities to the latter. The initial presentation of mostly sequential Knowledge Base Systems gives the reader a feel of the various problems and the proposed solutions for such systems. At the end of this first part, we present in detail one such system which integrates high-level, declarative rules into an active Object-Oriented Database. The resulting system is a flexible and extensible knowledge base system with multiple rule support.

In the second part of the book, we move into parallel Knowledge Base Systems by providing initial discussions of related research issues, such as parallel relational and object-oriented database systems. Many aspects of parallel rule execution are discussed including production, deductive, and active rules.

Finally, a complete parallel Knowledge Base System is presented. The system is based on the integration of a parallel Object-Oriented Database model with the multiple-rule integration scheme that is presented in the first part. The final system is implemented on a hierarchical multiprocessor architecture.

The book is intended as a reference text to the integration of database and knowledge base techniques for the researchers in the field of Knowledge Base Systems. It covers an extensive bibliography on the areas of rule integration in databases, namely active and deductive databases, as well as the unification of various rule types. Furthermore, the topics of parallel production, deductive, and active rule execution, both for databases and main-memory expert systems are reviewed.

Several chapters of the book (except probably of chapters 4 and 7 that describe a specific system) analyze in detail, using examples, various techniques for the above

topics. Therefore, the book can also be used as a textbook for an advanced course in Knowledge Base Systems. Finally, the book gives an in-depth insight to a specific parallel Knowledge Base System for the researchers that work in the fields of Active Databases, Knowledge Bases, and Object-Oriented Databases, on the one hand, and Parallel Databases, on the other.

We would like to thank Professor Ahmed K. Elmagarmid who encouraged us to write this book and read an earlier draft of it. We are also grateful to Professor Peter M.D. Gray who kindly agreed to write the foreword for this book.

Part of this work was carried out while Dr. Vlahavas was on sabbatical leave at Purdue University. Dr. Bassiliades was supported during his PhD studies by a scholarship from the Greek Foundation of State Scholarships (F.S.S. - I.K.Y.).