



Topics in Information Systems

# Query Processing in Database Systems

Edited by  
Won Kim  
David S. Reiner  
Don S. Batory



Springer-Verlag  
Berlin Heidelberg New York Tokyo

# Query Processing in Database Systems

Edited by  
Won Kim  
David S. Reiner  
Don S. Batory

With 127 Figures



Springer-Verlag  
Berlin Heidelberg New York Tokyo

## **Series Editors**

Dr. Michael L. Brodie  
Computer Corporation of America, Four Cambridge Center  
Cambridge, MA 02142/USA

Dr. John Mylopoulos  
Department of Computer Science, University of Toronto  
Toronto, Ontario M5S 1A4/Canada

Dr. Joachim W. Schmidt  
Fachbereich Informatik, Johann Wolfgang Goethe-Universität  
Dantestraße 9, D-6000 Frankfurt a. M. 11/FRG

## **Volume Editors**

Dr. Won Kim  
Microelectronics & Computer Technology Corporation  
9430 Research Boulevard, Austin, TX 78759/USA

Dr. David S. Reiner  
Computer Corporation of America, Four Cambridge Center  
Cambridge, MA 02142/USA

Dr. Don S. Batory  
Department of Computer Science, The University of Texas  
at Austin, Austin, TX 78712/USA

ISBN 3-540-13831-5 Springer-Verlag Berlin Heidelberg New York Tokyo  
ISBN 0-387-13831-5 Springer-Verlag New York Heidelberg Berlin Tokyo

Library of Congress Cataloging in Publication Data. Main entry under title: Query processing in database systems. (Topics in information systems) 1. Data base management--Addresses, essays, lectures. 2. Electronic data processing--Distributed processing--Addresses, essays, lectures. I. Kim, Won. II. Reiner, D. S. (David S.), 1952-. III. Batory, D. S. (Don S.), 1953-. IV. Series. QA76.9.D3Q47. 1984. 001.64. 84-20274

ISBN 0-387-13831-5 (U.S.)

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically those of translation, reprinting, re-use of illustrations, broadcasting, reproduction by photocopying machine or similar means, and storage in data banks. Under § 54 of the German Copyright Law where copies are made for other than private use a fee is payable to "Verwertungsgesellschaft Wort", Munich.

© Springer-Verlag Berlin Heidelberg 1985  
Printed in Germany

The use of registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Printing: Beltz Offsetdruck, Hemsbach/Bergstr.;  
Bookbinding: J. Schäffer OHG, Grünstadt  
2145/3140-543210

# Topics in Information Systems

Editors:

Michael L. Brodie

John Mylopoulos

Joachim W. Schmidt

# Topics in Information Systems

## Series Description

Dramatic advances in hardware technology have opened the door to a new generation of computer systems. At the same time, the growing demand for information systems of ever-increasing complexity and precision has stimulated the need in every area of Computer Science for more powerful higher-level concepts, techniques, and tools.

Future information systems will be expected to acquire, maintain, retrieve, manipulate, and present many different kinds of information. These systems will require user-friendly interfaces, powerful reasoning capabilities, and shared access to large information bases. Whereas the needed hardware technology appears to be within reach, the corresponding software technology for building these systems is not. The required dramatic improvements in software productivity will come from advanced application development environments based on powerful new techniques and languages.

The **concepts, techniques, and tools** necessary for the design, implementation, and use in future information systems are expected to result from the integration of those being developed and used in currently disjoint areas of Computer Science. Several areas bring their unique viewpoints and technologies to existing information processing practice. One key area is **Artificial Intelligence (AI)** which provides knowledge representation and reasoning capabilities for knowledge bases grounded on semantic theories of information for correct interpretation. An equally important area is **Databases** which provides means for building and maintaining large, shared databases based on computational theories of information for efficient processing. A third important area is **Programming Languages** which provides a powerful tool kit for the construction of large programs based on linguistic and methodological theories to ensure program correctness. To meet evolving information systems requirements, additional research viewpoints and technologies are or will be required from such areas as **Software Engineering, Computer Networks, Machine Architectures, and Office Automation.**

Although some integration of research results has already been achieved, a quantum leap in technological integration is needed to meet the demand for future information systems. This integration is one of the major challenges to Computer Science in the 1980s.

**Topics in Information Systems** is a series intended to report significant contributions on the integration of concepts, techniques, and tools that advance new technologies for information system construction. The series logo symbolizes the scope of topics to be covered and the basic theme of integration.

The logo will appear on each book to indicate the topics addressed.

	Artificial Intelligence	Databases	Programming Languages
concepts			
techniques			
tools			

The first book of the series, "On Conceptual Modelling: Perspectives from Artificial Intelligence, Databases and Programming Languages", Michael L. Brodie, John Mylopoulos, and Joachim W. Schmidt (Eds.), February 1984, which deals with concepts in the three areas, has the logo:

	Artificial Intelligence	Databases	Programming Languages
concepts	•	•	•
techniques			
tools			

The second book, "Query Processing in Database Systems", Won Kim, David S. Reiner, and Donald S. Batory (Eds.), March 1985, which deals with Database and Programming Language concepts, AI and Database techniques, and Database system tools, has the logo:

	Artificial Intelligence	Databases	Programming Languages
concepts		•	•
techniques	•	•	
tools		•	

The third book, "Office Automation", Dionysios C. Tschritzis (Ed.), March 1985, which will deal with the design and implementation of Office Systems, has the logo:

	Artificial Intelligence	Databases	Programming Languages
concepts		•	
techniques	•	•	
tools	•	•	•

Future books in the series will provide timely accounts of ongoing research efforts to reshape technologies intended for information system development.

March, 1985

Michael L. Brodie  
John Mylopoulos  
Joachim W. Schmidt

## Introduction

This book is an anthology of the results of research and development in database query processing during the past decade. The relational model of data provided tremendous impetus for research into query processing. Since a relational query does not specify access paths to the stored data, the database management system (DBMS) must provide an intelligent query-processing subsystem which will evaluate a number of potentially efficient strategies for processing the query and select the one that optimizes a given performance measure. The degree of sophistication of this subsystem, often called the optimizer, critically affects the performance of the DBMS.

Research into query processing thus started has taken off in several directions during the past decade. The emergence of research into distributed databases has enormously complicated the tasks of the optimizer. In a distributed environment, the database may be partitioned into horizontal or vertical fragments of relations. Replicas of the fragments may be stored in different sites of a network and even migrate to other sites. The measure of performance of a query in a distributed system must include the communication cost between sites. To minimize communication costs for queries involving multiple relations across multiple sites, optimizers may also have to consider semi-join techniques.

An extra dimension of difficulty arises when the database is distributed over a network of heterogeneous computers running DBMSs which support different data models. Typically, one translates the conceptual models (schemas) of the underlying databases to a common conceptual model and then defines a global conceptual model using the view mechanisms. User queries against this global model are translated into queries against the conceptual models of the underlying systems.

Much effort has been directed toward understanding the problems that arise when updates are issued against views. Many current DBMSs allow users to define views on the physical database and grant query privileges on the views to other users. However, it is often difficult, even impossible, for the effect of an update issued against the view to propagate to the physical database.

Current DBMSs have been designed largely to meet business data processing requirements, such as inventory control, airline reservations, banking, and so forth. However, they do not provide the functions and performance demanded by data management in computer-aided design of VLSI chips, database access by expert

systems, and statistical data processing. Research into the query characteristics and storage requirements of these applications is currently of great interest.

Optimizers in current relational systems only attempt to minimize the cost of processing single queries. To significantly improve the performance of these systems, multiple queries which reference the same relations may have to be simultaneously evaluated, or common subexpressions in a set of queries isolated and the results reused. Research along these lines is only beginning.

Recognizing the inherent mismatch between the conventional von Neumann processor architecture and the nonnumeric processing requirements of DBMSs, many researchers have proposed alternative architectures for nonnumeric data processing. They have proposed various algorithms for processing queries under these architectures.

Optimizers find optimal plans for evaluating queries based on available access paths and storage structures. Physical database design addresses the problems of selecting an optimal set of access paths and storage structures to support a class of applications. This was an active area of research even before relational query processing became fashionable. Physical database design must take into account such factors as frequency and types of access, database characteristics (size, distribution of data values, etc.), and query processing techniques. Further, it needs an accurate cost model of the existing access structures.

The book begins with a survey, by Jarke, Koch, and Schmidt, of important research results in query processing. Then it guides the reader through most of the important research topics in query processing, organized around seven sections: query processing in distributed DBMSs, query processing for multiple data models, database update through views, database access for special applications, techniques for optimizing the processing of multiple queries, query processing in database machines, and physical database design.

We have three papers on distributed query processing. Lohman, et al., provide an overview of the architecture of the R\* distributed relational DBMS, and discuss the query processing strategy explored for the system. To preserve autonomy among sites, R\* compiles a query at each site having a relation referenced in the query. The R\* optimizer considers CPU, I/O, and communications costs. Yu analyzes the costs and benefits of various semi-join strategies, and discusses estimation of intermediate result sizes. He also considers tree versus cyclic queries, and multiple copies of relations and relation fragments. Kambayashi summarizes techniques for converting cyclic queries into tree queries, which may then be processed by well-known semi-join based techniques.

There are two papers on multimodel query optimization. Dayal discusses integration of conceptual schemas of existing distributed databases that may be stored under heterogeneous DBMSs, and presents optimization of queries against the integrated schema. Logical integration is achieved through generalization (ISA) hierarchies. He also considers optimization of nonprocedural queries against underlying DBMSs that support procedural interfaces. Rosenthal and Reiner describe an architecture for supporting relational queries to data stored under both network and relational models.



They express intermodel mappings as system-generated relational views, and use view substitution to perform translations automatically.

There is only one paper on view update. Furtado and Casanova discuss two basic approaches to view update problems. One is to treat a view as an abstract data type containing the view definition and the set of allowed view updates. The other is to define general procedures for translating view updates to database updates based on conceptual schema dependencies.

We have three papers which describe data access problems that some important special classes of applications impose. Lorie, et al. discuss extensions to System R to support database organization and access requirements of VLSI CAD applications. These extensions support complex objects, clusters of related records that belong to different relations as single units for purposes of retrieval. Vassiliou, Clifford, and Jarke identify the database access requirements of expert systems, present four architectural strategies for coupling these systems to databases, and investigate criteria for choosing among the strategies. Ozsoyoglu and Ozsoyoglu describe Summary Table by Example (STBE), a QBE-like screen-oriented query language for statistical databases. They also compare expressive powers of the language with SQL and Aggregation by Example, and then outline techniques for evaluating queries expressed in STBE.

To stimulate research in the relatively unexplored area of global optimization, we include three papers. Jarke presents techniques for recognizing and exploiting common subexpressions in different representations (including PROLOG) of relational queries. He also shows that the programming language abstraction "selector" is a useful concept in considering access path support for multiple queries. Kim examines various issues that arise in simultaneously processing a set of SQL queries, and offers a first-cut strategy for grouping queries and sequencing the execution of the query groups. Queries are grouped based on considerations of data dependence and estimated access costs of individual queries. Kambayashi and Ghosh introduce the reader to the notion of the consecutive retrieval (CR) property. A set of queries referencing a common file can be processed in a single scan of the file, if the queries possess the CR property. The authors also discuss techniques for processing a set of  $N$  queries that does not have the CR property in  $n$  ( $1 < n \ll N$ ) scans of the file.

There are two papers on database machines, one on a commercial system and one on an ongoing large-scale research project. Ubell provides an overview of the architecture of Britton-Lee's Intelligent Database Machine (IDM). He outlines the structure of the IDM DBMS software and discusses its performance advantages. Shaw gives a brief description of the organization of the NONVON relational database machine prototype and considers query processing algorithms on the machine.

The book ends with three papers in the section on physical database design. Batory reviews recent advances in the modeling of the storage structures in DBMSs and shows how these advances can lead to a technology that may automate the development of physical database software, software which handles placement and retrieval of data on secondary storage. March and Carlis survey some classical modeling and optimization techniques that have been used for the record structuring and access path design

problems. Whang, Wiederhold and Sagalowicz present an overview of the theory of separability, which reduces the problem of assigning access structures to a database to the collective subproblems of optimizing the access structures of individual relations (record types), independently of other relations (record types). They show application of the theory to the physical design of relational and network databases.

March 1985

Won Kim  
David Reiner  
Don Batory

---

# On Conceptual Modelling:

**Perspectives from Artificial Intelligence, Databases, and Programming Languages**

Editors: **M. L. Brodie, J. Mylopoulos, J. W. Schmidt**

1984. 25 figures. XI, 510 pages. (Topics in Information Systems). ISBN 3-540-90842-0

**Contents:** Artificial Intelligence, Database, and Programming Language Overviews. – Perspectives from Artificial Intelligence. – Perspectives from Databases. – Perspectives from Programming Languages. – Concluding Remarks from Three Perspectives. – References. – Authors and Symposium Participants. – Index.

Conceptual modelling relates to all areas of computer science, but especially to artificial intelligence, databases, and programming languages. Here is the first published collection of state-of-the-art research papers in these domains. Its purpose is to consider conceptual modelling as a topic in its own right rather than as an aspect of data modelling, and to present and compare research on knowledge representation, semantic data models, and data abstraction in this context.

The contributions consist of overviews and reports, each chapter having been written and edited for readers in all three areas. Also included are transcripts of symposium discussions which took place among the contributors during a workshop on conceptual modelling at Intervale; these interdisciplinary discussions of each paper clarify many aspects which might otherwise remain obscure to non-specialists. Key features of the book include introductions to pertinent concepts, and the integration of recent results; focus on twelve research projects, involving specific applications such as database design; and challenging suggestions for further research, especially in the concluding comments by leading experts in the three main fields of inquiry.



Springer-Verlag  
Berlin  
Heidelberg  
New York  
Tokyo

---

# Relational Database Systems

## Analysis and Comparison

Editors: **J. W. Schmidt, M. L. Brodie**

1983. XV, 618 pages. ISBN 3-540-12032-7

**Contents:** Introduction. – Features of Relational Database Systems. – Analysis of Relational Database Management Systems. – Feature Summaries and Comparisons. – References.

The book is the most comprehensive and detailed analysis of existing relational database management systems to date. It presents a generic characterization of an RDBMS (independently of specific RDBMSs) in terms of:

- relational database constituents
- relational functional capabilities
- definition, generation and administration facilities
- interfaces and DBMS architecture
- operational aspects

These features are then used as a common basis to analyze fourteen of the most important existing RDBMSs. The fourteen systems analyses are then compared, in a tabular format, with respect to the features, and system feature summaries are presented.

The book is introduced by a foreword written by Dr. E. F. Codd, the inventor of the relational approach to databases.

The book is intended to assist the reader in developing a detailed understanding of the Relational Data Model, Relational Database Management Systems, and the state-of-the-art in relational DBMS technology. It provides a comprehensive check list of features with which to evaluate RDBMSs or DBMSs in general.



Springer-Verlag  
Berlin  
Heidelberg  
New York  
Tokyo

---

# Table of Contents

## Section I. Introduction to Query Processing

### Introduction to Query Processing

*Matthias Jarke, Jürgen Koch, Joachim Schmidt* . . . . . 3

## Section II. Query Processing in Distributed Database Management Systems

### Query Processing in R\*

*Guy Lohman, C. Mohan, Laura Haas, Dean Daniels, Bruce Lindsay, Paricia Selinger, Paul Wilms* . . . . . 31

### Distributed Database Query Processing

*Clement Yu* . . . . . 48

### Processing Cyclic Queries

*Yahiko Kambayashi* . . . . . 62

## Section III. Query Processing for Multiple Data Models

### Query Processing in a Multidatabase System

*Umeshwar Dayal* . . . . . 81

### Querying Relational Views of Networks

*Arnon Rosenthal, David S. Reiner* . . . . . 109

## Section IV. Database Updates through Views

Updating Relational Views <i>Anthony Furtado, Marco Casanova</i> . . . . .	127
-------------------------------------------------------------------------------	-----

## Section V. Database Access for Special Applications

Supporting Complex Objects in a Relational System for Engineering Databases <i>Raymond Lorie, Won Kim, Dan McNabb, Wil Plouffe, Andreas Meier</i> . . . . .	145
Database Access Requirements of Knowledge-Based Systems <i>Yannis Vassiliou, Jim Clifford, Matthias Jarke</i> . . . . .	156
A Query Language for Statistical Databases <i>Z. Meral Ozsoyoglu, Gultekin Ozsoyoglu</i> . . . . .	171

## Section VI. Techniques for Optimizing the Processing of Multiple Queries

Common Subexpression Isolation in Multiple Query Optimization <i>Matthias Jarke</i> . . . . .	191
Global Optimization of Relational Queries: A First Step <i>Won Kim</i> . . . . .	206
Query Processing Using the Consecutive Retrieval Property <i>Yahiko Kambayashi, Sakti Ghosh</i> . . . . .	217

## Section VII. Query Processing in Database Machines

The Intelligent Database Machine (IDM) <i>Michael Ubell</i> . . . . .	237
Relational Query Processing on the NON-VON Supercomputer <i>David Shaw</i> . . . . .	248

## Section VIII. Physical Database Design

Progress Toward Automating the Development of Database System Software <i>Don Batory</i> . . . . .	261
Physical Database Design: Techniques for Improved Database Performance <i>Salvatore March, John Carlis</i> . . . . .	297
The Property of Separability and Its Application to Physical Database Design <i>Kyu-Young Whang, Gio Wiederhold, Daniel Sagalowicz</i> . . . . .	297
 <i>References</i> . . . . .	 319
 <i>List of Authors</i> . . . . .	 347
 <i>Subject Index</i> . . . . .	 351

Section I

Introduction to  
Query Processing



