


Tutorial: DISTRIBUTED DATABASE MANAGEMENT

JAMES A. LARSON • SAEED RAHIMI



IEEE CATALOG NUMBER EHO222-0
LIBRARY OF CONGRESS NUMBER 84-48797
IEEE COMPUTER SOCIETY ORDER NUMBER 575
ISBN 0-8186-0575-8



IEEE COMPUTER SOCIETY





THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

COMPUTER
SOCIETY
PRESS



Tutorial: DISTRIBUTED DATABASE MANAGEMENT

JAMES A. LARSON • SAEED RAHIMI



IEEE CATALOG NUMBER EH0222-0
LIBRARY OF CONGRESS NUMBER 84-48797
IEEE COMPUTER SOCIETY ORDER NUMBER 575
ISBN 0-8186-0575-8

 IEEE COMPUTER SOCIETY

IEEE
COMPUTER
SOCIETY
PRESS 

 1984
A CENTURY OF ELECTRICAL ENGINEERING

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Published by IEEE Computer Society Press
1109 Spring Street
Suite 300
Silver Spring, MD 20910

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 29 Congress Street, Salem, MA 01970. Instructors are permitted to photocopy isolated articles for noncommercial classroom use without fee. For other copying, reprint or republication permission, write to Director, Publishing Services, IEEE, 345 E. 47 St., New York, NY 10017. All rights reserved. Copyright © 1985 by The Institute of Electrical and Electronics Engineers, Inc.

IEEE Catalog Number EHO222-0
Library of Congress Number 84-48797
IEEE Computer Society Order Number 575
ISBN 0-8186-0575-8 (Paper)
ISBN 0-8186-4575-X (Microfiche)

Order from: IEEE Computer Society
Post Office Box 80452
Worldway Postal Center
Los Angeles, CA 90080

IEEE Service Center
445 Hoes Lane
Piscataway, NJ 08854



THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC

Preface

In order to design, implement, and maintain a successful distributed database management system, a basic understanding of the components that constitute such a system is needed. It is important to understand how each component operates and how the components relate to each other. If designers, implementers, and users don't understand these principles, the distributed database management system may fail to meet all of its objectives.

The primary goals of this book are to provide a written description of the basic components of distributed database management systems and to examine how those components relate to each other. After reading the overviews and papers in this book, the reader should understand what command translators, semantic integrity constraint enforcers, runtime support processors, query decomposers, distributed execution monitors, and communications processors do; how these components relate to each other; and how each of these components works. Chapter 1 describes the requirements of a distributed database management system and outlines a reference architecture of processors that satisfies these requirements.

Another goal of this book is to provide two collections of papers and reports. The first collection (in Chapters 2 through 8) provides descriptions of each of the major components of distributed database management systems. Chapter 2 discusses database command translations, which are used to support data model independence and nonhomogeneous database management systems. Chapter 3 discusses semantic integrity constraints, why they are important, and how they can be enforced. Chapter 4 reviews how to partition a database into fragments to be stored at different sites in the distributed database management system and how to decompose a database request into subrequests that will be processed at each site. The software that decomposes database requests is the key to location transparency. Chapters 5 and 6 describe the distributed execution monitor which is responsible for the coordination of the execution of subrequests at the various sites. Chapter 7 describes the communication subsystem used to send requests and data between the sites. Chapter 8 considers issues in the design and administration of distributed database management systems. The papers at the end of each chapter form a core of basic reference material. Readers who are interested in greater detail may use the suggestions for additional reading found at the end of each chapter.

The second collection of papers and reports (Chapter 9) presents descriptions of existing distributed database management systems. These papers describe both the production systems used today and the experimental prototypes of the distributed database management systems that will be available tomorrow.

This book and the papers included in it are used as reference material in professional development seminars conducted by the authors and in an advanced course in database management systems at the University of Minnesota.

James A. Larson
Saeed K. Rahimi

Table of Contents

Preface	iii
Chapter 1: Introduction	1
Interconnecting Heterogeneous Database Management Systems	13
<i>V.D. Gligor and G.L. Luckenbaugh (Computer, January 1984, pages 33-43)</i>	
A Model for Describing Distributed Database Management Systems Architecture	24
<i>P. Decitre (Proceedings of the 16th Annual Electronics and Aerospace Conference and Exposition, September 1983, pages 367-376)</i>	
Concurrency Control Issues in Distributed Heterogeneous Database Management Systems	34
<i>V.D. Gligor and R. Popescu-Zeletin (Distributed Data Sharing Systems, F.A. Schreiber and W. Litwin (Editors), 1985, pages 43-56)</i>	
Chapter 2: Transforming Database Commands	47
Bridging the Gap between Network and Relational Database Management Systems	53
<i>J.A. Larson (Computer, September 1983, pages 82-92)</i>	
Local Query Translation and Optimization in a Distributed System	64
<i>E. Onuegbe, S. Rahimi, and A.R. Hevner (Proceedings of the AFIPS National Computer Conference, Volume 52, 1983, pages 229-239)</i>	
Chapter 3: Semantic Integrity Constraints	75
On Efficient Monitoring of Database Assertions in Distributed Databases	77
<i>D.Z. Badal (Proceedings of the Fourth Berkeley Conference on Distributed Data Management and Computer Networks, August 1979, pages 125-137)</i>	
Chapter 4: Decomposing Requests	91
Comparative Models of the File Assignment Problem	95
<i>L.W. Dowdy and D.V. Foster (Computing Surveys, Volume 14, June 1982, pages 287-313)</i>	
A Comprehensive Approach to Fragmentation and Allocation of Data in Distributed Databases	122
<i>S.B. Navathe and S. Ceri</i>	
Optimal Query Processing for Distributed Database Systems	135
<i>W.W. Chu and P. Hurley (IEEE Transactions on Computers, Volume C-31, September 1982, pages 835-850)</i>	
Methods for Data Retrieval in Distributed Systems	151
<i>A.R. Hevner (Proceedings of the Second Symposium on Reliability in Distributed Software and Database Systems, 1982, pages 1-10)</i>	
Transaction Optimization in a Distributed Database Testbed System	161
<i>P.A. Dwyer and A.R. Hevner (Proceedings of COMPSAC, 1983, pages 564-570)</i>	
Chapter 5: Concurrency and Replication Control	169
A Survey of Techniques for Synchronization and Recovery in Decentralized Computer Systems ...	175
<i>W.H. Kohler (Computing Surveys, Volume 13, Number 2, June 1981, pages 149-183)</i>	
Concurrency Control in Distributed Database Systems	210
<i>P.A. Bernstein and N. Goodman (Computing Surveys, Volume 13, Number 2, June 1981, pages 185-221)</i>	
A Sophisticate's Introduction to Distributed Database Concurrency Control	247
<i>P.A. Bernstein and N. Goodman (Proceedings of the 8th VLDB Conference, September 1982, pages 62-76)</i>	

System Level Concurrency Control for Distributed Database Systems	262
<i>D.J. Rosenkrantz, R.E. Stearns, and P.M. Lewis, II (ACM Transactions on Database Systems, Volume 3, Number 2, June 1978, pages 178-198)</i>	
A Majority Consensus Approach to Concurrency Control for Multiple Copy Databases	283
<i>R.H. Thomas (ACM Transactions on Database Systems, Volume 4, Number 2, June 1979, pages 180-209)</i>	
On Optimistic Methods for Concurrency Control	308
<i>H.T. Kung and J.T. Robinson (ACM Transactions on Database Systems, Volume 6, Number 2, June 1981, pages 213-226)</i>	
Distributed Deadlock Detection Algorithm	320
<i>R. Obermarck (ACM Transactions on Database Systems, Volume 7, Number 2, June 1982, pages 187-208)</i>	
Detection of Mutual Inconsistency in Distributed Systems	342
<i>D.S. Parker, Jr., G.J. Popek, G. Rudisyon, A. Sloughton, B.J. Walker, F. Walton, L.M. Chow, D. Edwards, S. Kiser, and C. Kline (IEEE Transactions on Software Engineering, Volume SE-9, Number 3, May 1983, pages 240-246)</i>	
A Robust Algorithm for Updating Duplicate Databases	360
<i>C.A. Ellis (Proceedings of the 2nd Berkeley Workshop on Distributed Data Management and Computer Networks, 1977, pages 146-158)</i>	
Chapter 6: Distributed Execution Monitor	363
Computation and Communication in R ² : A Distributed Database Manager	367
<i>B.G. Lindsay, L.M. Haas, C. Mohan, P.F. Wilms, and R.A. Yost (ACM Transactions on Computer Systems, Volume 2, Number 3, February 1984, pages 24-38)</i>	
SER: A System for Distributed Execution Based on Decentralized Control Techniques	380
<i>G. Sergeant and L. Trelle (Proceedings of the Fifth International Conference on Computer Communications: Computer Communications: Increasing Benefits for Society, October 1980, pages 225-230)</i>	
The Notions of Consistency and Predicate Locks in a Database System	386
<i>K.P. Eswaran, J.N. Gray, R.A. Lorie, and J.L. Traiger (Communications of the ACM, Volume 19, Number 11, November 1976, pages 624-633)</i>	
The Recovery Manager of the System R Database Manager	396
<i>J. Gray, P. McJones, M. Blasgen, B. Lindsay, R. Lorie, T. Price, F. Putzola, and J. Traiger (Computing Surveys, Volume 13, Number 2, June 1981, pages 223-242)</i>	
A Quorum-Based Commit Protocol	416
<i>D. Skeen (Proceedings of the 6th Berkeley Workshop on Distributed Data Management and Computer Networks, February 1982, pages 69-80)</i>	
Chapter 7: Communications Subsystem	429
Network Protocols	433
<i>A.S. Tanenbaum (Computing Surveys, Volume 13, Number 4, December 1981, pages 453-489)</i>	
Architecture Considerations for Local Computer Networks	470
<i>K.J. Thurber and H.A. Freeman (First International Conference on Distributed Computing System Proceedings, October 1979, pages 131-142)</i>	
Chapter 8: Design of Distributed DBMSs	483
Distribution of Data Administration Authority	489
<i>F.B. Wilson and J.A. Larson (Proceedings of the ACM Pacific Northwest Conference, 1975)</i>	
Administering a Distributed Data Base Management System	493
<i>H.M. Walker (ACM SIGMOD Record, Volume 12, Number 3, April 1982, pages 86-99)</i>	

Chapter 9: Case Studies	507
Introduction to a System for Distributed Databases (SDD-1)	517
<i>J.B. Rothnie, Jr., P.A. Bernstein, S. Fox, N. Goodman, M. Hammer, T.A. Landers, C. Reeve, D.W. Shipman, and E. Wong (ACM Transactions on Database Systems, Volume 5, Number 1, March 1980, pages 1-17)</i>	
An Introduction to Distributed Query Compilation in R*	534
<i>D. Daniels, P. Selinger, L. Haas, B. Lindsay, C. Mohan, A. Walker, and P. Wilms (IBM Research Report RJ3497 (41354), June 1982)</i>	
An Overview of MULTIBASE	556
<i>T. Landers and R.L. Rosenberg (Distributed Data Bases, H-J. Schneider (Editor), 1982, pages 153-183)</i>	
An Overview of the Mermaid System—A Frontend to Heterogeneous Databases	580
<i>M. Templeton, D. Brill, A. Hwang, I. Kameny, and E. Lund (Proceedings of EASCON 83, 1983, pages 387-402)</i>	
Overview of an Ada Compatible Distributed Database Manager	596
<i>A. Chan, U. Dayal, S. Fox, N. Goodman, D.R. Ries, and D. Skeen (Proceedings of SIGMOD 1983, May 1983, pages 228-237)</i>	
Five-Schema Architecture Extends DBMS to Distributed Applications	606
<i>C. Devor, R. Elmasri, J.A. Larson, S. Rahimi, and J.P. Richardson (Electronic Design, March 18, 1982, pages ss27-ss32)</i>	
The SIRIUS-DELTA Architecture: A Framework for Co-Operating Database Systems	612
<i>C. Esculier (Computer Networks, Volume 8, 1984, pages 43-48)</i>	
A Distributed Data Base Version of INGRES	618
<i>M. Stonebraker and E. Neuhold (1977 Berkeley Workshop on Distributed Data Management and Computer Networks, May 1977, pages 19-36)</i>	
EMPACT: A Distributed Database Application	648
<i>A. Norman and M. Anderton (Proceedings AFIPS National Computer Conference, Vol. 52, pages 203-217)</i>	
Glossary	663
Subject Index	667