REAL-TIME DATABASE SYSTEMS

Architecture and Techniques

Edited by
Kam-Yiu Lam
Tei-Wei Kuo

Kluwer Academic Publishers

REAL-TIME DATABASE SYSTEMS

Architecture and Techniques

Edited by

Kam-Yiu Lam
City University of Hong Kong

Tei-Wei Kuo
National Taiwan University

KLUWER ACADEMIC PUBLISHERS
Boston / Dordrecht / London

Distributors for North, Central and South America:
Kluwer Academic Publishers
101 Philip Drive
Assinippi Park
Norwell, Massachusetts 02061 USA
Telephone (781) 871-6600
Fax (781) 681-9045
E-Mail <kluwer@wkap.com>

Distributors for all other countries:
Kluwer Academic Publishers Group
Distribution Centre
Post Office Box 322
3300 AH Dordrecht, THE NETHERLANDS
Telephone 31 78 6392 392
Fax 31 78 6392 254
E-Mail <services@wkap.nl>


Electronic Services <http://www.wkap.nl>

Library of Congress Cataloging-in-Publication Data

Real-time database systems : architecture and techniques / edited by Kam-Yiu Lam, Tei-Wei Kuo.
p. cm. -- (The Kluwer international series in engineering and computer science ; SECS 593)
Includes bibliographical references and index.
ISBN 0-7923-7218-2 (alk. paper)
1. Database design. 2. Real-time data processing. I. Lam, Kam-Yiu, 1964- II. Kuo, Tei-Wei. III. Series.

QA76.9.D26 R435 2000
004'.3--dc21 00-046214

Printed on acid-free paper.

Printed in the United States of America

REAL-TIME DATABASE SYSTEMS

Architecture and Techniques

The Kluwer International Series in Engineering and Computer Science

REAL-TIME SYSTEMS

Consulting Editor: **John A. Stankovic,** *University of Virginia*

Other books in this series:

DEADLINE SCHEDULING FOR REAL-TIME SYSTEMS: ***EDF and Related Algorithms,*** John A. Stankovic, Marco Spuri, Krithi Ramamritham, Giorgio C. Buttazzo, ISBN: 0-7923-8269-2

HARD REAL-TIME COMPUTING SYSTEMS: ***Predictable Scheduling Algorithms and Applications***, by Giorgio C. Buttazzo, ISBN: 0-7923-9994-3

REAL-TIME DATABASE AND INFORMATION RESEARCH ADVANCES, by Azer Bestavros, Victor Wolfe, ISBN: 0-7923-8011-8

REAL-TIME SYSTEMS: ***Design Principles for Distributed Embedded Applications,*** by Hermann Kopetz, ISBN: 0-7923-9894-7

REAL-TIME DATABASE SYSTEMS: ***Issues and Applications,*** edited by Azer Bestavros, Kwei-Jay Lin and Sang Hyuk Son, ISBN: 0-7923-9897-1

FAULT-TOLERANT REAL-TIME SYSTEMS: ***The Problem of Replica Determinism,*** by Stefan Poledna, ISBN: 0-7923-9657-X

RESPONSIVE COMPUTER SYSTEMS: ***Steps Toward Fault-Tolerant Real-Time Systems,*** by Donald Fussell and Miroslaw Malek, ISBN: 0-7923-9563-8

IMPRECISE AND APPROXIMATE COMPUTATION, by Swaminathan Natarajan, ISBN: 0-7923-9579-4

FOUNDATIONS OF DEPENDABLE COMPUTING: ***System Implementation***, edited by Gary M. Koob and Clifford G. Lau, ISBN: 0-7923-9486-0

FOUNDATIONS OF DEPENDABLE COMPUTING: ***Paradigms for Dependable Applications,*** edited by Gary M. Koob and Clifford G. Lau, ISBN: 0-7923-9485-2

FOUNDATIONS OF DEPENDABLE COMPUTING: ***Models and Frameworks for Dependable Systems,*** edited by G. M. Koob and C. G. Lau, ISBN: 0-7923-9484-4

THE TESTABILITY OF DISTRIBUTED REAL-TIME SYSTEMS, Werner Schütz; ISBN: 0-7923-9386-4

A PRACTITIONER'S HANDBOOK FOR REAL-TIME ANALYSIS: ***Guide to Rate Monotonic Analysis for Real-Time Systems,*** Carnegie Mellon University (M. Klein, T. Ralya, B. Pollak, R. Obenza, M. G. Harbour); ISBN: 0-7923-9361-9

FORMAL TECHNIQUES IN REAL-TIME FAULT-TOLERANT SYSTEMS, J. Vytopil; ISBN: 0-7923-9332-5

SYNCHRONOUS PROGRAMMING OF REACTIVE SYSTEMS, N. Halbwachs; ISBN: 0-7923-9311-2

REAL-TIME SYSTEMS ENGINEERING AND APPLICATIONS, M. Schiebe, S. Pferrer; ISBN: 0-7923-9196-9

SYNCHRONIZATION IN REAL-TIME SYSTEMS: ***A Priority Inheritance Approach,*** R. Rajkumar; ISBN: 0-7923-9211-6

CONSTRUCTING PREDICTABLE REAL TIME SYSTEMS, W. A. Halang, A. D. Stoyenko; ISBN: 0-7923-9202-7

FOUNDATIONS OF REAL-TIME COMPUTING: ***Formal Specifications and Methods,*** A. M. van Tilborg, G. M. Koob; ISBN: 0-7923-9167-5

FOUNDATIONS OF REAL-TIME COMPUTING: ***Scheduling and Resource Management,*** A. M. van Tilborg, G. M. Koob; ISBN: 0-7923-9166-7

REAL-TIME UNIX SYSTEMS: ***Design and Application Guide,*** B. Furht, D. Grostick, D. Gluch, G. Rabbat, J. Parker, M. McRoberts, ISBN: 0-7923-9099-7

Acknowledgments

Dedicated to my wife I-Ju Chen and my parents Cheng-I Kuo and Chin-Yun Hsu for their love and support.

Tei-Wei Kuo

We would like to thank Prof. Stankovic for his valuable comments on the organization and the contents of this book. We also want to express our gratitude to the Distributed Real-Time Systems (DRTS) Research Group of the University of Skövde and the Distributed Real-Time Multimedia Databases of the University of Virginia for providing us valuable information regarding their real-time database prototypes which are included in the chapter for future directions. Special thanks go to all of the contributing authors, particularly Prof. Krithi Ramamritham and Prof. Sten Andler, for their support and valuable advice during the development of the book. At last, we would like to thank Mr. Reynold Cheng for helping us in integrating and formating all of the chapters.

Kam-Yiu Lam and Tei-Wei Kuo

Preface

In recent years, a lot of research work has been devoted to the design of database systems for real-time applications. A real-time database system (RTDBS) is usually defined as a database system where transactions are associated with deadlines on their completion times. In addition, some of the data items in a real-time database are associated with temporal constraints on their validity. Example applications include systems for stock trading, navigation, telephone management, and computer integrated manufacturing. In order to commit a real-time transaction, the transaction has to be completed before its deadline, and all of its accessed data items must be valid up to its commit time. Otherwise, the usefulness of the results will be seriously affected. For many cases, any deadline violation of a hard real-time transaction may result in disasters.

There are two major technology components in building a RTDBS: real-time systems and database systems. Unfortunately, these two technologies might have conflicting goals. For example, a database system is often designed to maximize the system throughput, while real-time techniques usually aim at meeting deadlines and improving the predictability of the system performance. Frequently, when conventional database technologies are applied to a RTDBS, many real-time requirements of the systems might be seriously affected and cannot be guaranteed. In the past decade, a lot of research effort has been devoted to the study of RTDBS, and a large number of research papers were published in this area [9]. Important conferences and special issues of international journals were organized and devoted to real-time database systems, such as the 1996 and 1997 International Workshop on Real-time Database Systems [2, 3], the 1995, 1997, and 1999 International Workshop on Active and Real-Time Database Systems [4, 5, 6], Information Systems [7], Real-Time Systems Journal [11], ACM SIGMOD Records [14, 15], IEEE Transactions on Knowledge and Data Engineering [8], and Journal of Systems and Software [10]. A recent issue of the Real-Time Systems Journal was also devoted to RTDBS [12].

Early study in RTDBS was focused on the design of efficient transaction scheduling algorithms and concurrency control protocols [1]. The goals are to satisfy the time constraints of transactions and the temporal properties of data items [13]. Other design issues essential to the performance of RTDBS are buffer management, I/O scheduling, etc. Although a large number of research papers have been published, there is a lack of a comprehensive system reference book which covers the important technologies in this area. Furthermore, there still exists many misunderstandings about this area, e.g., the relationship with active database systems and the performance issues of a RTDBS. It is the purpose of this book to provide an in-depth review on the current technologies

and algorithms in the design of RTDBS. Its target readers are graduate students, researchers, and practitioners in real-time systems and database systems.

The book is divided into six parts, which covers all the important topics in the area. Part I discusses the important basic concepts of RTDBS and the characteristics of their potential applications. Part II covers one of the most important issues in RTDBS: real-time concurrency control, which is essential in maintaining database consistency and providing a consistent view to transactions. Part III is on other run-time system management issues, such as buffer management, recovery and system failure handling, disk scheduling, security, and overload management. An important characteristics of many real-time database applications is the active property, which is important for generating timely responses to critical events in the external environment. The active issues in RTDBS are covered in Part IV. Since many RTDBS, such as traffic management systems and tele-communication management systems, are often distributed in nature, Part V addresses important issues related to the processing of transactions in a distributed real-time database system. They are commitment, distributed concurrency control, and replicated data management. The last part of the book, Part VI, describes a RTDB prototype and the future directions for RTDB research.

Kam-Yiu Lam and
Tei-Wei Kuo

References

[1] R. Abbott and H. Garcia-Molina, "Scheduling Real-time Transactions: A Performance Evaluation," in *Proceedings of 14th International Conference on Very Large Data Bases*, September 1988.

[2] *The First International Workshop on Real-Time Databases: Issues and Applications*, Newport Beach, California, March 1996.

[3] A. Bestarvos and V. Fay-Wolfe (editors), *Real-Time Database and Information Systems: Research Advances*, Kluwer Academic Publishers, Massachusetts, 1997.

[4] M. Berndtsson and J. Hansson (editors), *Active and Real-Time Database Systems (ARTDB-95), Proceedings of the First International Workshop on Active and Real-Time Database Systems*, Skövde, Sweden, 9-11 June 1995, 284pp, ISBN 3-540-19983-7, Springer Verlag (London) Ltd.

[5] S.F. Andler and J. Hansson (eds), *Proceedings of the Second International Workshop on Active, Real-Time, and Temporal Database Systems (ARTDB-97)*, Como/Milan, Italy, Springer-Verlag Lecture Notes, 1998.

[6] A.H. Frigeri, W.A. Halang, S.H. Son (eds), *Active and Real Time Data Base Systems, in Real Time Programming 1999, Proceedings of the 24th IFAC/IFIP Workshop Real Time Programming (WRTP'99) and 3rd International Workshop Active and Real-Time Database Systems (ARTDB-99)*, Schloss Dagstuhl, May - June 1999.

[7] *Information Systems*, vol. 21, no. 1, 1996.

[8] *IEEE Transactions on Knowledge and Data Engineering*, vol. 7, no. 4, 1995.

[9] G. Ozsoyoglu and R. Snodgrass, "Temporal and Real-Time Databases: A Survey", *IEEE Transactions on Knowledge and Data Engineering*, vol. 7, no. 4, pp. 513-532, 1995.

[10] *Journal of Systems and Software*, vol. 42, no. 3, 1998.

[11] *Real-Time Systems Journal*, vol. 4, no. 3, September 1992.

[12] *Real-Time Systems Journal*, vol. 19, no. 3, Nov. 2000.

[13] K. Ramamritham, "Real-time Databases", *International Journal of Distributed and Parallel Databases*, vol. 1, no. 2, 1993.

[14] S. Son, "Real-Time Databases," *ACM SIGMOD Record*, vol. 17, no. 1, 1988.

[15] *ACM SIGMOD Record*, vol. 23, no. 3, 1994.

Contributing Authors

Brad Adelberg
Department of Computer Science, Northwestern University, Evanston, Illinois, USA

Sten F. Andler
Department of Computer Science, University of Skövde, Skövde, Sweden

Reynold Cheng
Department of Computer Science and Information Systems, University of Hong Kong, Hong Kong

Anindya Datta
Centre for Electronic Commerce, Georgia Institute of Technology, Atlanta, Georgia, USA

Joakim Eriksson
Department of Computer Science, University of Skövde, Skövde, Sweden

Binto George
Department of Computer Science, Western Illinois University, Macomb, Illinois, USA.

Ramesh Gupta
Goldencom Technologies, Inc., Fremont, California, USA.

Jörgen Hansson
Department of Computer and Information Science, Linköping University, Linköping, Sweden

Jayant R. Haritsa
Database Systems Lab, Supercomputer Education and Research Centre, Indian Institute of Science, Bangalore, India.

Ben Kao
Department of Computer Science and Information Systems, University of Hong Kong, Hong Kong

Tei-Wei Kuo
Department of Computer Science and Information Engineering, National Taiwan University, Taipei, Taiwan

Kam-Yiu Lam
Department of Computer Science, City University of Hong Kong, Hong Kong

Doug Locke
TimeSys Corporation, USA

Jonas Mellin
Department of Computer Science, University of Skövde, Skövde, Sweden

Sarit Mukherjee
Panasonic Labs, Princeton, New Jersey, USA

Krithi Ramamritham
Department of Computer Science, University of Massachusetts, Amherst, Massachusetts, USA.

S. Seshadri
Lucent Bell Labs, Murray Hill, New Jersey, USA.

Rajendran M. Sivasankaran
GTE Laboratories, Waltham Massachusetts, USA.

Sang H. Son
Department of Computer Science, University of Virginia, Charlottesville, Virginia, USA

John A. Stankovic
Department of Computer Science, University of Virginia, Charlottesville, Virginia, USA

Özgür Ulusoy
Department of Computer Engineering, Bilkent University, Ankara, Turkey

List of Figures

List of Tables

PART I
OVERVIEW, MISCONCEPTIONS, AND ISSUES

Chapter 1

REAL-TIME DATABASE SYSTEMS: AN OVERVIEW OF SYSTEM CHARACTERISTICS AND ISSUES

Tei-Wei Kuo
Department of Computer Science and Information Engineering
National Taiwan University
ktw@csie.ntu.edu.tw

Kam-Yiu Lam
Department of Computer Science
City University of Hong Kong, Kowloon, Hong Kong
cskylam@cityu.edu.hk

1. INTRODUCTION

In the past two decades, the research in *real-time database systems (RTDBS)* has received a lot of attention [7, 8, 9, 10]. It consists of two different important areas in computer science: *real-time systems* and *database systems*. Similar to conventional real-time systems, transactions in RTDBS are usually associated with time constraints, e.g., deadlines. On the other hand, RTDBS must maintain a database for useful information, support the manipulation of the database, and process transactions. Some example applications of RTDBS are integrated manufacturing systems, programmed stock trading systems, air traffic control systems, and network management system. Typically, these application systems need predictable response time, and they often have to process various kinds of queried in a timely fashion.

Differently from traditional database systems, RTDBS must not only maintain database integrity but also meet the urgency of transaction executions. Different performance metrics are adopted. Real-time transactions are, in general, classified into three types: No *hard real-time transaction* should have its deadline missed, and its deadline must be guaranteed by the system. On the other hand, any deadline violations of *soft real-time transactions* may only