


T H I R D E D I T I O N

DATABASE SYSTEM CONCEPTS



ABRAHAM SILBERSCHATZ HENRY F. KORTH
S. SUDARSHAN

DATABASE SYSTEM CONCEPTS

Third Edition

Abraham Silberschatz

Bell Laboratories

Henry F. Korth

Bell Laboratories

S. Sudarshan

Indian Institute of Technology, Bombay

 **WCB
McGraw-Hill**

Boston, Massachusetts Burr Ridge, Illinois Dubuque, Iowa
Madison, Wisconsin New York, New York San Francisco, California St. Louis, Missouri

WCB/McGraw-Hill

A Division of the McGraw-Hill Companies



DATABASE SYSTEM CONCEPTS, THIRD EDITION, THIRD PRINTING

Copyright ©1999 by The McGraw-Hill Companies, Inc. All rights reserved. Previous editions © 1986 and 1991. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

This book is printed on acid-free paper.

5 6 7 8 9 0 DOC/DOC 1 0

ISBN 0-07-031086-6

The editor was Elizabeth A. Jones.
The production supervisor was Scott Hamilton.
The cover was designed by Christopher Brady.
R. R. Donnelley & Sons was printer and binder.

Library of Congress Cataloging-in-Publication Data

Silberschatz, Abraham.

Database system concepts / Abraham Silberschatz, Henry F. Korth, S. Sudarshan. – 3rd ed.

p. cm. – (McGraw-Hill series in computer science)

ISBN 0-07-031086-6 (acid-free paper)

Includes bibliographical references and index.

1. Database management. I. Korth, Henry F. II. Sudarshan, S.

III. Title. IV. Series

QA76.9.D3 S5637 1998

005.74—dc21

98-13188

DATABASE SYSTEM CONCEPTS

McGraw-Hill Series in Computer Science

Senior Consulting Editor

C.L. Liu, University of Illinois at Urbana-Champaign

Consulting Editor

Allen B. Tucker, Bowdoin College

Fundamentals of Computing and Programming

Computer Organization and Architecture

Computers in Society/Ethics

Systems and Languages

Theoretical Foundations

Software Engineering and Database

Artificial Intelligence

Networks, Parallel and Distributed Computing

Graphics and Visualization

The MIT Electrical and Computer Science Series

Software Engineering and Database

Cohen: *Ada as a Second Language*

Fairley: *Software Engineering Concepts*

Silberschatz, Korth, and Sudarshan: *Database System Concepts*

Musa, Iannino, and Okumoto: *Software Reliability*

Pressman: *Software Engineering: A Beginner's Guide*

Pressman: *Software Engineering: A Practitioner's Approach*

Wiederhold: *Database Design*

*In memory of my father Joseph Silberschatz,
and my grandparents Stepha and Aaron Rosenblum*

Avi Silberschatz

*In memory of my grandparents:
Giuseppa and Anthony Affatigato and Gertrude and Frank Korth*

Hank Korth

In memory of my grandfather, S. Mahalingam

S. Sudarshan

PREFACE

Database management has evolved from a specialized computer application to a central component of a modern computing environment. As such, knowledge about database systems has become an essential part of an education in computer science. Our purpose in this text is to present the fundamental concepts of database management. These concepts include aspects of database design, database languages, and database-system implementation.

This text is intended for a first course in databases at the junior or senior undergraduate, or first-year graduate, level. In addition to basic material for a first course, the text also contains advanced material that can be used for course supplements, or as introductory material for an advanced course.

We assume only a familiarity with basic data structures, computer organization, and a high-level (Pascal-like) programming language. Concepts are presented using intuitive descriptions, many of which are based on our running example of a bank enterprise. Important theoretical results are covered, but formal proofs are omitted. The bibliographic notes contain pointers to research papers in which results were first presented and proved, as well as references to material for further reading. In place of proofs, figures and examples are used to suggest why we should expect the result in question to be true.

The fundamental concepts and algorithms covered in the book are often based on those used in existing commercial or experimental database systems. Our aim is to present these concepts and algorithms in a general setting that is not tied to one particular database system.

In this third edition of *Database System Concepts*, we have retained the overall style of the first and second editions, while addressing the evolution of database management. Every chapter has been edited, and most have been modified extensively. We shall describe the changes in detail shortly.

Organization

The text is organized in eight major parts, plus two appendices:

- **Overview** (Chapters 1 and 2). Chapter 1 provides a general overview of the nature and purpose of database systems. We explain how the concept of a database system has developed, what the common features of database systems are, what a database system does for the user, and how a database system interfaces with operating systems. We also introduce an example database application: a banking enterprise consisting of multiple bank branches. This example is used as a running example throughout the book. This chapter is motivational, historical, and explanatory in nature. Chapter 2 presents the entity–relationship model. This model provides a high-level view of the issues in database design, and of the problems that we encounter in capturing the semantics of realistic applications within the constraints of a data model.
- **Relational model** (Chapters 3 through 5). Chapter 3 focuses on the relational data model, covering the relevant relational algebra and relational calculus. Chapter 4 focuses on the most influential of the user-oriented relational languages: SQL. Chapter 5 covers other relational languages. These three chapters describe data manipulation: queries, updates, insertions, and deletions. Algorithms and design issues are deferred to later chapters. Thus, these chapters are suitable for those individuals or lower-level classes who want to learn what a database system is, without getting into the details of what the internal algorithms and structure are.
- **Database constraints** (Chapters 6 and 7). Chapter 6 presents constraints from the standpoint of database integrity; Chapter 7 shows how constraints can be used in the design of a relational database. Functional dependencies and referential integrity are presented in Chapter 6, as are mechanisms for integrity maintenance, such as triggers and assertions. The theme of this chapter is the protection of the database from accidental damage. Chapter 7 introduces the theory of relational database design. Such topics as normalization and data dependencies are covered, with emphasis on the motivation for each normal form and on the intuitive meaning of each type of data dependency.
- **Object-based systems** (Chapters 8 and 9). Chapter 8 covers object-oriented databases. It introduces the concepts of object-oriented programming, and shows how these concepts form the basis for a data model. No prior knowledge of object-oriented languages is assumed. Chapter 9 covers object-relational databases, and shows how we can extend the relational data model to include object-oriented features, such as inheritance and complex types.
- **Data storage and retrieval** (Chapters 10 through 12). Chapter 10 deals with disk, file, and file-system structure, and with the mapping of relational and object data to a file system. A variety of data-access techniques are presented in Chapter 11, including hashing, B⁺-tree indices, and grid file indices. Chapter 12 addresses query-evaluation algorithms, and query optimization based on equivalence-preserving query transformations. These chapters are targeted

at people who want to gain an understanding of the internals of the storage and retrieval components of a database.

- **Transaction management** (Chapters 13 through 15). Chapter 13 focuses on the fundamentals of a transaction-processing system, including transaction atomicity, consistency, isolation, and durability, as well as the notion of serializability. It is suitable for those individuals or classes that need an introduction to the issues of transaction management, yet do not require the details of concurrency-control and recovery protocols. In Chapter 14, we focus on concurrency control, and present several techniques for ensuring serializability, including locking, timestamping, and optimistic (validation) techniques. Deadlock issues are also covered in that chapter. Chapter 15 covers the primary techniques for ensuring correct transaction execution despite system crashes and disk failures. These techniques include logs, shadow pages, checkpoints, and database dumps.
- **Parallel and distributed systems** (Chapters 16 through 18). Chapter 16 covers computer-system architecture, and describes the influence of the underlying computer system on the database system. We discuss centralized systems, client-server systems, parallel and distributed architectures, and network types. In Chapter 17, on parallel databases, we explore a variety of parallelizing techniques, including I/O parallelism, interquery and intraquery parallelism, and interoperation and intraoperation parallelism. We also discuss cost estimation, query optimization, and parallel-system design. Chapter 18 revisits issues of database design, transaction management, and query evaluation and optimization, in the context of a distributed database system.
- **Advanced topics** (Chapters 19 through 21). Chapter 19 covers numerous special topics, including security and integrity, standardization, performance benchmarks and performance tuning, time in databases, user interfaces, and active databases. Chapter 20 deals with advanced transaction processing. We discuss transaction-processing monitors, high-performance transaction systems, real-time transaction systems, and transactional workflows. In Chapter 21, we introduce numerous new applications for database systems. First, we discuss decision-support systems, including data-analysis, data-mining, and data-warehousing applications. We then examine spatial and geographic databases, multimedia databases, and mobile and personal databases. Finally, we investigate information-retrieval systems for textual data, and distributed information systems, including the World Wide Web.
- **Appendices.** Although most new database applications use either the relational model or the object-oriented model, the network and hierarchical data models are still in use. As we did in the second edition, we have covered systems based on the network and hierarchical models in Appendices A and B, respectively. However, due to a decline in interest in teaching these older models, the full text of these appendices has been moved to the internet, with only a summary appearing in the printed text. The full version of the appendices is available on the web (<http://www.bell-labs.com/topic/books/db-book>) or via anonymous FTP from <ftp.research.bell-labs.com> in directory `dist/db-book`.

The Third Edition

Many comments and suggestions were forwarded to us concerning the first and second editions. These inputs — coupled with our own observations while teaching at the University of Texas, IIT Bombay, and IBM, and with our analysis of the directions in which database technology is evolving — have prodded us to produce this third edition, and have guided us in its production. Our basic procedure was to rewrite the material in each chapter, bringing the older material up to date, adding discussions on recent developments in database technology, improving the exercises, and adding new references. We have also restructured the organization of parts of the book. For the benefit of those readers familiar with the second edition, we explain the main changes here.

- **Entity–Relationship model.** We have improved our coverage of the entity–relationship (E-R) model. Chapter 2 in the third edition is similar to the previous Chapter 2; however, we have expanded coverage of issues in E-R database design. Design issues are discussed throughout the chapter, and are addressed in particular in two new sections, 2.2 and 2.8. Extended E-R features are discussed in more detail than in the second edition, as is reduction of an E-R schema to tables (Sections 2.7 and 2.9, respectively).
- **Relational databases.** In the second edition, we expanded our coverage of the relational model. In the third edition, we again devote Chapter 3 to the relational model and to the formal relational languages: the relational algebra and the relational calculus. In Section 3.5, we have added discussions of the generalized projection, outer-join operations, and aggregation. Chapter 4 now covers SQL exclusively. Our coverage of SQL has been significantly expanded to include features of the SQL-92 standard, in addition to the existing coverage based on the SQL-89. Some SQL implementations may support only SQL-89 and not SQL-92; we explicitly identify those features of SQL-92 that are not supported in SQL-89. We now cover QBE, and Quel, in Chapter 5. The research language Datalog, which had been presented in the second edition in Chapter 14, is included now in Chapter 5, and receives a more detailed discussion. Chapter 6 covers integrity constraints; Chapter 7 covers database-design issues and normal forms. These chapters were Chapters 5 and 6, respectively, in the second edition.
- **Object-based databases.** Expansions to our coverage of the object-oriented data model appear in Chapters 8, 9, and 10. In Chapter 8, we augment the material from Chapter 13 in the second edition with a discussion of object-oriented programming languages. Chapter 9 is a new chapter in which we present object-relational data models. These models extend the relational data model by providing a richer type system, including object orientation, and add constructs to relational query languages, such as SQL, to deal with the added data types. Chapter 10 now includes a section on the storage structures for object-oriented databases (Section 10.9).

- **Query processing.** Our treatment of query processing has been greatly expanded from the second edition. In Chapter 12, we now present detailed explanations of what different ways we have for implementing various relational operations, and of how to estimate their execution costs. We have also expanded our coverage of query transformations that preserve equivalence of the query results, and have incorporated new material on query optimization. Chapter 17 now has expanded coverage of parallel query processing.
- **Transaction processing.** As we did for the second edition, we have reorganized and slightly expanded our coverage of transaction processing (Chapters 13 through 15). Some of the advanced material on transaction processing from Chapter 12 in the second edition is now organized in a new chapter on advanced transaction processing: Chapter 20 (see the note on advanced topics that follows). Chapter 13 introduces issues in all aspects of transaction management, with details deferred to later chapters. This organization allows instructors to choose between just introducing transaction-processing concepts (by covering only Chapter 13), or offering detailed coverage (based on Chapters 13 through 15).
- **Computer architectures and parallel systems.** Chapter 16 is a new chapter that covers computer-system architecture and the influence of the underlying computer system on the database system. We discuss centralized systems, client-server systems, and network types. We also present parallel and distributed architectures; we cover these systems in detail in Chapters 17 and 18, respectively. Chapter 17 is also a new chapter; it covers parallel databases. We explore a variety of parallelizing techniques, including I/O parallelism, interquery and intraquery parallelism, and interoperation and intraoperation parallelism. We also discuss cost evaluation, query optimization, and parallel-system design. New material on disk organization based on the RAID architectures is included in Chapter 10.
- **Advanced topics.** Although we have modified and updated the entire text, we concentrate our presentation of material pertaining to ongoing database research and new database applications in three new chapters.

We present numerous special topics in Chapter 19. Chapter 16 in the second edition on security and integrity, now appears as Section 19.1. The remaining sections are new material pertaining to standardization projects, performance benchmarks, performance tuning, time in databases, user interfaces, and active databases.

Some of the material on transaction processing in Chapter 20 appeared previously in Chapter 12 in the second edition; Section 20.1 on transaction processing monitors, Section 20.2 on high-performance transaction systems, Section 20.4 on real-time transaction systems, and Section 20.6 on transactional workflows are all new.

In Chapter 21, we introduce several new applications for database systems; the material in this chapter is all new. First, we discuss decision-support systems, including data-analysis, data-mining, and data-warehousing applications. We then discuss spatial and geographic databases, multimedia

databases, and mobile and personal databases. Finally, we discuss information-retrieval systems for textual data, and distributed information systems, including the World Wide Web.

Instructor's Note

The book contains both basic and advanced material, which might not be covered in a single semester. It is possible to design courses using various subsets of the chapters. We outline possibilities here:

- Chapter 5 can be omitted if students will not be using QBE, Quel, or Datalog as part of the course.
- Chapter 7 contains a series of normal forms, in decreasing order of practical importance. Later sections (Section 7.4 on) may be omitted, if desired.
- If object-orientation is to be covered in a separate advanced course, Chapters 8 and 9, and Section 10.9, can be omitted. Alternatively, they could constitute the foundation of an advanced course in object databases.
- Chapters 11 and 12 contain some material that may be more suitable for an advanced course. You might choose to omit some or all of Sections 11.6, 11.9, 12.7, 12.8, and 12.10.
- Both our coverage of transaction processing (Chapters 13 through 15) and our coverage of database-system architecture (Chapters 16 through 18) consist of an overview chapter (Chapters 13 and 16, respectively), followed by chapters with details. You might choose to use Chapters 13 and 16, while omitting Chapters 14, 15, 17, and 18, if you defer these latter chapters to an advanced course.
- The sections of the final three chapters (Chapters 19 through 21) are largely independent. Based on instructor or student interest, a custom-tailored set of subsections may be chosen as end-of-semester enrichment material. Those chapters as a whole are suitable for an advanced topics course.

Model course syllabi, based on the text, can be found on the World Wide Web home page of the book (see following section).

Supplements and Mailing List

For information about the teaching supplements that complement this book, send electronic mail to wcbcomp@mcgraw-hill.com. In the United States, you may call 800-338-3987. The McGraw-Hill web page for Computer Science texts is <http://www.mhhe.com/engcs/compsci>.

We also now provide an environment in which users can communicate among themselves and with us. We have created a mailing list consisting of users of our book with the electronic mail address db-book@research.bell-labs.com. If you wish to be on the list, please send a message to db-book@research.bell-labs.com, include your name, affiliation, title, and electronic mail address.

Home Page

A World Wide Web home page for the book is available with the URL:

<http://www.bell-labs.com/topic/books/db-book>

The home page contains information about the book, such as up-to-date errata, model course syllabi, and information about teaching supplements.

Errata

We have endeavored to eliminate typos, bugs, and the like from the text. But, as in new releases of software, bugs probably remain. We would appreciate it if you would notify us of any errors or omissions in the book. An updated errata page will be accessible from the book's WWW home page. Also, if you would like to suggest improvements or to contribute exercises, we would be glad to hear from you. Any correspondence should be sent to Avi Silberschatz, Bell Laboratories, Lucent Technologies Inc., 700 Mountain Avenue, Murray Hill, NJ 07974, USA. Internet electronic mail should be addressed to db-book@research.bell-labs.com.

Acknowledgments

This edition is based on the two previous editions, so we thank once again the many people who helped us with the first and second editions, including Don Batory, Haran Boral, Robert Brazile, Sara Strandtman, Won Kim, Anil Nigam, Bruce Porter, Carol Kroll, Jim Peterson, Fletcher Mattox, Ron Hitchens, Alberto Mendelzon, Alan Fekete, Hyoung-Joo Kim, Keith Marzullo, Mark Roth, Greg Speegle, and Henry Korth (father of Henry F.)

Greg Speegle and Dawn Bezviner helped us to prepare the instructor's manual for the first edition. Their work served as the basis for the new instructor's manual for the third edition, which was prepared with the help of K. V. Raghavan.

This edition has benefited from the many useful comments provided to us by the numerous students who have used the second edition in our classes at the University of Texas and at IBM, or have used drafts of the third edition at IIT Bombay. In addition, numerous people have written or spoken to us about the book, and have offered suggestions and comments. Although we cannot mention all these people here, we especially thank R. B. Abhyankar, Paul Bourgeois, Michael Carey, J. Edwards, Christos Faloutsos, Homma Farian, Shashi Gadia, Jim Gray, Le Gruenwald, Yannis Ioannidis, Gary Lindstrom, Dave Maier, Hector Garcia-Molina, Ami Motro, Cyril Orji, K. V. Raghavan, Marek Rusinkiewicz, S. Seshadri, Shashi Shekhar, Amit Sheth, Nandit Soparkar, and Marianne Winslett.

Lyn Dupré copyedited the book. Sara Strandtman edited our text into the \LaTeX format, and helped us prepare the revised text for this edition. Sara also edited the instructor's manual and transparencies.

The new cover—an evolution of the covers of the first two editions—was drawn by Chris Brady of McGraw-Hill. Marilyn Turnamian created an early draft of the cover design. The idea of using ships as part of the cover concept was

originally suggested to us by Bruce Stephan. We would also like to thank our editor, Eric Munson, for his wisdom and patience.

Finally, Sudarshan would like to acknowledge his brother, Mohan, and mother, Indira, for their support. Hank would like to acknowledge his wife, Joan, and his children, Abby and Joe, for their love and understanding. Avi would like to acknowledge his mother for her love and support.

A. S.

H. F. K.

S. S.

CONTENTS

Preface	xv
1 Introduction	1
1.1 Purpose of Database Systems	1
1.2 View of Data	4
1.3 Data Models	7
1.4 Database Languages	12
1.5 Transaction Management	13
1.6 Storage Management	14
1.7 Database Administrator	15
1.8 Database Users	15
1.9 Overall System Structure	16
1.10 Summary	19
Exercises	20
Bibliographical Notes	20
2 Entity–Relationship Model	23
2.1 Basic Concepts	23
2.2 Design Issues	28
2.3 Mapping Constraints	30
2.4 Keys	34
2.5 Entity–Relationship Diagram	36
2.6 Weak Entity Sets	37
2.7 Extended E-R Features	41
2.8 Design of an E-R Database Schema	47
2.9 Reduction of an E-R Schema to Tables	52

2.10	Summary	58
	Exercises	59
	Bibliographical Notes	62
3	Relational Model	63
3.1	Structure of Relational Databases	63
3.2	The Relational Algebra	71
3.3	The Tuple Relational Calculus	86
3.4	The Domain Relational Calculus	90
3.5	Extended Relational-Algebra Operations	94
3.6	Modification of the Database	100
3.7	Views	102
3.8	Summary	106
	Exercises	107
	Bibliographical Notes	110
4	SQL	111
4.1	Background	111
4.2	Basic Structure	113
4.3	Set Operations	120
4.4	Aggregate Functions	122
4.5	Null Values	124
4.6	Nested Subqueries	125
4.7	Derived Relations	129
4.8	Views	130
4.9	Modification of the Database	131
4.10	Joined Relations	136
4.11	Data-Definition Language	140
4.12	Embedded SQL	145
4.13	Other SQL Features	148
4.14	Summary	148
	Exercises	149
	Bibliographical Notes	152
5	Other Relational Languages	153
5.1	Query-by-Example	153
5.2	Quel	165
5.3	Datalog	174
5.4	Summary	188
	Exercises	188
	Bibliographical Notes	190
6	Integrity Constraints	193
6.1	Domain Constraints	193
6.2	Referential Integrity	195
6.3	Assertions	200
6.4	Triggers	201
6.5	Functional Dependencies	202
6.6	Summary	210
	Exercises	211
	Bibliographical Notes	213

7	Relational Database Design	215
7.1	Pitfalls in Relational-Database Design	215
7.2	Decomposition	217
7.3	Normalization Using Functional Dependencies	221
7.4	Normalization Using Multivalued Dependencies	231
7.5	Normalization Using Join Dependencies	239
7.6	Domain-Key Normal Form	242
7.7	Alternative Approaches to Database Design	244
7.8	Summary	246
	Exercises	247
	Bibliographical Notes	250
8	Object-Oriented Databases	251
8.1	New Database Applications	251
8.2	The Object-Oriented Data Model	253
8.3	Object-Oriented Languages	262
8.4	Persistent Programming Languages	263
8.5	Persistent C++ Systems	267
8.6	Summary	271
	Exercises	272
	Bibliographical Notes	272
9	Object-Relational Databases	275
9.1	Nested Relations	275
9.2	Complex Types and Object Orientation	278
9.3	Querying with Complex Types	283
9.4	Creation of Complex Values and Objects	287
9.5	Comparison of Object-Oriented and Object-Relational Databases	288
9.6	Summary	289
	Exercises	289
	Bibliographical Notes	290
10	Storage and File Structure	293
10.1	Overview of Physical Storage Media	293
10.2	Magnetic Disks	296
10.3	RAID	301
10.4	Tertiary Storage	307
10.5	Storage Access	309
10.6	File Organization	312
10.7	Organization of Records in Files	318
10.8	Data-Dictionary Storage	322
10.9	Storage Structures for Object-Oriented Databases	324
10.10	Summary	332
	Exercises	333
	Bibliographical Notes	336
11	Indexing and Hashing	339
11.1	Basic Concepts	339
11.2	Ordered Indices	340