

# An Introduction to Database Systems

VOLUME I

Fourth Edition

C. J. DATE

# An Introduction to Database Systems

VOLUME I

Fourth Edition

C. J. DATE



**ADDISON-WESLEY PUBLISHING COMPANY**

Reading, Massachusetts • Menlo Park, California  
Don Mills, Ontario • Wokingham, England • Amsterdam  
Sydney • Singapore • Tokyo • Mexico City • Bogotá  
Santiago • San Juan

This book is in the  
**Addison-Wesley Systems Programming Series**

Consulting Editors: IBM Editorial Board  
Sponsoring Editor: James T. DeWolf  
Project Editor: Marion E. Howe  
Packaging Service: Superscript Associates  
Text Designer: Quadrata, Inc.  
Cover: Marshall Henrichs  
Manufacturing Supervisor: Hugh Crawford

**Library of Congress Cataloging in Publication Data**

Date, C. J.

An introduction to database systems.

Includes bibliographies and index.

	1. Data base management	1. Title.
QA76.9.D3D37	1986	001.64 85-1422
ISBN 0-201-14201-5 (v. 1)		

Copyright © 1986 by Addison-Wesley Publishing Company, Inc. 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, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America. Published simultaneously in Canada.

ABCDEFGHIJ-HA-898765

*For Lindy*

## THE SYSTEMS PROGRAMMING SERIES

- \* The Program Development Process  
Part I—The Individual Programmer Joel D. Aron
- \* The Program Development Process  
Part II—The Programming Team Joel D. Aron
- \* Mathematical Foundations of  
Programming Frank Beckman
- \* Structured Programming: Theory  
and Practice Richard C. Linger  
Harlan D. Mills  
Benard I. Witt
- \* Coded Character Sets: History and  
Development Charles E. Mackenzie
- \* The Structure and Design of Program-  
ming Languages John E. Nicholls
- \* The Environment for Systems Programs Frederic G. Withington
- \* Communications Architecture for  
Distributed Systems R. J. Cypser
- \* An Introduction to Database Systems,  
Fourth Edition C. J. Date
- \* An Introduction to Database Systems:  
Vol. II C. J. Date
- \* Database Security and Integrity Eduardo B. Fernandez  
R. C. Summers  
C. Wood
- \* Fundamentals of Interactive  
Computer Graphics James Foley
- \* Compiler Design Theory Andries Van Dam  
Philip M. Lewis II  
Daniel J. Rosenkrantz  
Richard E. Stearns
- \* Sorting and Sort Systems Harold Lorin
- \* Operating Systems Harold Lorin  
Harvey M. Deitel  
William Burge
- \* Recursive Programming Techniques
- \* Modeling and Analysis: An Introduction  
to System Performance Evaluation  
Methodology Hisashi Kobayashi
- \* Conceptual Structures: Information  
Processing in Mind and Machines John F. Sowa
- \* Data Structures: Form  
and Function Harry F. Smith
- \* Published

# Foreword

The field of systems programming primarily grew out of the efforts of many programmers and managers whose creative energy went into producing practical, utilitarian systems programs needed by the rapidly growing computer industry. Programming was practiced as an art where each programmer invented unique solutions to problems with little guidance beyond that provided by immediate associates. In 1968, the late Ascher Opler, then at IBM, recognized that it was necessary to bring programming knowledge together in a form that would be accessible to all systems programmers. Surveying the state of the art, he decided that enough useful material existed to justify a significant codification effort. On his recommendation, IBM decided to sponsor The Systems Programming Series as a long-term project to collect, organize, and publish those principles and techniques that would have lasting value throughout the industry.

The Series consists of an open-ended collection of text-reference books. The contents of each book represent the individual author's view of the subject area and do not necessarily reflect the views of the IBM Corporation. Each is organized for course use but is detailed enough for reference. Further, the Series is organized in three levels: broad introductory material in the foundation volumes, more specialized material in the software volumes, and very specialized theory in the computer science volumes. As such, the Series meets the needs of the novice, the experienced programmer, and the computer scientist.

Taken together, the Series is a record of the state of the art in systems programming that can form the technological base for the systems programming discipline.

*The Editorial Board*

## IBM EDITORIAL BOARD

Marilyn Bohl  
Richard P. Case, Chairman  
Willy W. Chiu  
Gerhard Chroust  
Robert H. Glaser  
Charles I. Gold  
Edward A. Green  
Paul S. Herwitz

James P. Morrissey  
John E. Nicholls  
Jean E. Sammet  
David Sayre  
Robert G. Webster

William B. Gruener (Addison-Wesley)  
Mark S. Dalton (Addison-Wesley)  
M. C. Varley (Addison-Wesley)

## About the Author.

C. J. Date is an independent author, lecturer, and consultant, specializing in relational database systems. He is currently based in Saratoga, California.

Mr. Date joined IBM in 1967 in England, where he worked on the integration of database functions into PL/I. Later (1974) he moved to the IBM Development Center in California, where he was responsible for the design of a database language known as UDL (Unified Database Language). More recently, he was involved in technical planning and externals design for the IBM relational products SQL/DS (announced in 1981 for DOS and in 1983 for VM) and DB2 (announced in 1983 for MVS). He left IBM in May 1983.

Mr. Date has been active in the field of database for over fifteen years. He was one of the first people anywhere to recognize the significance of Codd's pioneering work on the relational model. He has lectured widely on technical subjects—principally on database topics, and especially on relational database—throughout the United States and also in Europe, Australia, Latin America, and the Far East. In addition to the present book, his publications include *An Introduction to Database Systems: Volume II* (1982), which covers a range of more advanced aspects of the subject; *Database: A Primer* (1983), which treats database systems (from micros to mainframes) specifically from the nonspecialist's point of view; and *A Guide to DB2* (1984), a description of the IBM relational product DB2 and its companion products QMF and DXT. He has also produced numerous technical papers and articles and has made a number of contributions to database theory. He is currently conducting a series of professional seminars on database technology, both in the United States and elsewhere.

# Preface to the Fourth Edition

The fourth edition of this book is so different from all previous editions that it is to all intents and purposes a brand new book. Even the structure has changed, as will be explained below. But of course the overall objective is still the same as it always was—namely, to provide the basis for a solid education in the fundamentals of database technology. The general style and level of exposition are also still very much the same; the book is still intended primarily as a textbook rather than as a work of reference, and it still retains its introductory and tutorial flavor.

Readers are assumed to be professionally interested in some aspect of data processing. They may, for example, be systems analysts or designers, application programmers, systems programmers, students following a university or similar course in computer science, or teachers of such a course. Readers are expected to have a reasonable appreciation of the capabilities of a modern computer system, with particular reference to the file-handling features of such a system; they should also have some knowledge of at least one high-level programming language. Since these prerequisites are not particularly demanding, however, I am hopeful that the book will prove suitable as an introductory text for anyone concerned with using or implementing a database system, or for anyone who simply wishes to broaden a general knowledge of the computer science field.

The book is divided into six major parts:

1. Basic Concepts
2. A Sample Relational System
3. Relational Database Management
4. The Database Environment
5. Nonrelational Systems
6. Some Current Developments

Each part in turn is subdivided into a number of chapters:

- Part 1 (three chapters) provides a broad introduction to the concepts of database systems in general, and in particular distinguishes between relational and nonrelational systems. At the time of writing it is still the case that the majority of installed systems are nonrelational. However, this state of affairs is unlikely to persist for very much longer—relational products are being developed and installed at an ever-increasing rate; indeed, it is a fact that almost all modern products are relational. Furthermore, almost all current research and development is based on relational ideas. The emphasis of this book is therefore heavily on the relational approach.
- Part 2 (six chapters) consists of a tutorial treatment of the major concepts of relational systems, using the IBM system DB2 as a vehicle for illustrating those concepts.
- Part 3 (eight chapters), which is the longest part of the book (indeed, it is almost a book within a book) consists of a detailed examination of the ideas, both theoretical and pragmatic, underlying relational systems such as DB2. Part 3 also includes a detailed discussion of another relational system, namely INGRES, which is significantly different from DB2 in a number of ways.
- Part 4 (three chapters) then goes on to discuss certain other aspects of database systems (relational or otherwise)—recovery, concurrency, security, integrity, and so on.
- Next, Part 5 (three chapters) introduces the major nonrelational approaches (inverted list, hierarchic, and network), and describes a representative set of nonrelational systems: DATACOM/DB, IMS, and IDMS.
- Finally, Part 6 (two chapters) discusses some current research directions in the database field.

*Note:* The six parts have deliberately been written to be somewhat independent of one another; Part 3 (at least Chapters 11–14) can be studied before Part 2 if the reader wishes, and Part 5 can be studied before Parts 2 and 3 or after Part 6. However, my own recommendation is to read the material in sequence as written.

Most chapters are followed by a set of exercises and (usually) answers also; often the answers give additional information about the subject of the exercise. Most chapters are also followed by an extensive list of references, most of them annotated. References are identified in the text by numbers in square brackets. For example, [11.1] refers to the first item in the list of references at the end of Chapter 11, namely, a paper by E. F. Codd published in *Communications of the ACM* (CACM), Vol. 13, No. 6 (June 1970).

For readers who may be familiar with the third edition, the major differences from that edition are summarized below.

*Part 1:* The old introductory chapter (Chapter 1) has been expanded into two chapters (1 and 2) and made a little more gentle. The old Chapter 2 on storage structures (now Chapter 3) has been completely rewritten and greatly expanded. The old Chapter 3 on different data models has been dropped; however, all the material of that chapter still survives (indeed, in expanded form) at other points in the book.

*Part 2:* Chapters 4–9 replace Chapters 5–10 from the third edition. The material on the System R prototype has been replaced throughout by material on the DB2 product.

*Part 3:* This part is largely new, or at least considerably improved over the third edition. Chapter 10 (new) describes the INGRES system in some detail. Chapters 11–14 present a significantly expanded and improved treatment of the relational model (replacing the old Chapters 4 and 11–13); in particular, the integrity aspects of the model are covered in depth. Chapter 15 (new) discusses a number of practical aspects of relational systems. Chapter 16 (new) is an introduction to relational query optimization. Finally, Chapter 17 is a revised and extended version of the old Chapter 14 on further normalization.

*Part 4:* This part is also new. It consists of three new chapters. Chapter 18 discusses transaction processing (in particular, recovery and concurrency considerations); Chapter 19 describes the problems of security and integrity in a database system; and Chapter 20 discusses a number of auxiliary (frontend) systems—report writers, business graphics systems, natural language interfaces, application generators, and so on.

*Part 5:* This part replaces Parts 3–5 from the third edition. It consists of three long chapters, one for each of the major nonrelational approaches (inverted list, hierarchic, and network). Each chapter first discusses the relevant data model and then goes on to describe a typical system supporting that model. Chapter 21 (new) covers DATACOM/DB; Chapter 22 (revised and condensed version of the old Chapters 15–22) covers IMS; and Chapter 23 (revised and condensed version of the old Chapters 23–26) covers IDMS. Note that Chapter 23 in particular now describes an (important) concrete system instead of just the “official” CODASYL specifications, which were never fully implemented. Chapters 21 and 23 also include material on the new “relational” extensions to DATACOM/DB and IDMS.

*Part 6:* This part is also new. Chapter 24 describes distributed database systems and also briefly discusses database machines. Chapter 25 presents an introduction to semantic modeling; in particular, it describes the entity/relationship model and the extended relational model RM/T.

In addition, the old chapters 27 and 28 have been dropped, though much of the material from those chapters is preserved at various points throughout the book.

Why all these changes? There are several major reasons:

1. First, the promise of relational systems—foretold as far back as the first edition of this book (1975)—is at last beginning to come to fruition. As already mentioned, almost every modern database product is relational, and almost all database research is relational-based. Furthermore, most of the older (nonrelational) products are currently being extended to provide some kind of relational support. Given the intended audience for this book, therefore, it no longer seems necessary or appropriate to give “equal time” to the nonrelational approaches.

2. Following on from the previous point: Given the widespread acceptance of the relational approach, a solid understanding of the fundamentals of that approach is now more important than ever, for both researchers and practitioners—in fact, for anyone involved in any aspect of database management, in any capacity whatsoever. The relational portions of the book have therefore been completely rewritten, and at the same time consolidated and greatly expanded, with the aim of providing a basis for that solid understanding.

3. In addition, the “coming of age” of relational systems has made it possible to base most examples and discussions on commercially available products, instead of just on research prototypes as in earlier editions. (It goes without saying, of course, that commercial products are subject to continual change; details of specific systems are thus not guaranteed to be accurate as of the most recent release. Furthermore, inclusion or exclusion of a particular product in this book should not be construed as either endorsement or otherwise of the product in question—the purpose of including any specific system is merely to use it to make some specific point or illustrate some specific idea.)

4. The reduced coverage of nonrelational systems has made it possible to include material on a number of topics for which there was no room in the third edition. Specifically, the book now includes at least an introduction to all of the following: query optimization, recovery, concurrency, security, integrity, forms-based systems, natural language interfaces, report writing, graphics interfaces, application generators, semantic modeling (including entity/relationship modeling), distributed databases, and database machines. Note: Many of these topics are treated in more detail in Volume II (see below).

One result of all these changes, incidentally, is that the prefaces to the previous editions are no longer very relevant. They have therefore been dropped from this edition.

As indicated above, this book has a sequel, namely *An Introduction to Database Systems: Volume II* (first edition, Addison-Wesley, 1982)—referred to throughout the present text as simply *Volume II*. Volume II provides a more thorough treatment of some of the topics introduced in Parts 3, 4, and 6 of the present book (hereinafter referred to as *Volume I*). However, I hope the discussions above make

it clear that this fourth edition of Volume I is far more than just a rehash of material from the previous (third) edition and Volume II; on the contrary, it contains a great deal of new material, and a significantly different emphasis on much of the existing material. Equally, of course, it is not intended to replace Volume II, which goes into considerably more detail (on its own topics) than is necessary or appropriate for the present volume.

*Acknowledgments:* As always, it is a great pleasure to acknowledge my debt to the many people involved, directly or indirectly, in the production of this book. First, the text has benefited enormously from the comments of numerous students on seminars I have taught based on this material. It has also benefited from the comments of several reviewers of the third edition: Randell Flint, Alan Hevner, Larry Rowe, Sharon Salveter, Stewart Shen, Paul Spirakis, Philip Stone, Richard Walters, and most especially Dave DeWitt. My thanks also to the following, each of whom reviewed at least some portion of the manuscript of this edition or made technical material available or otherwise helped me find answers to my many technical questions: Paul Butterworth, Roger Buchanan, Don Casey, Ted Codd, Dave DeWitt, Guy Egerton, Ron Fagin, Randell Flint, Robin Haines, Doug Hembry, Peter Kreps, Jack Nevison, Larry Rowe, Phil Shaw, Mike Stonebraker, and Sharon Weinberg. I would also like to thank Relational Technology Inc. for permission to base certain portions of Chapter 20 on an article that I wrote for them some time ago on the Visual Programming tools of INGRES. Finally, this book would never have been possible without the support and assistance of my long-suffering family, nor (obviously) without the professional expertise of my editor Elydia Siegel and everyone involved at Addison-Wesley. To them all, my heartfelt appreciation and thanks.

*Saratoga, California*

C. J. Date

# Contents

---

## PART 1

### Basic Concepts

1

## 1

### An Overview of Database Management

3

1.1	An introductory example	3
1.2	What is a database system?	5
1.3	Operational data	9
1.4	Why database?	12
1.5	Data independence	15
1.6	Relational systems and others	19
	Exercises	21
	References and bibliography	23
	Answers to selected exercises	25

## 2

### An Architecture for a Database System

29

2.1	Purpose	29
2.2	The three levels of the architecture	29
2.3	The external level	32
2.4	The conceptual level	35
2.5	The internal level	36
2.6	Mappings	37
2.7	The database administrator	37
2.8	The database management system	40
2.9	Data communications	40
	Exercises	41
	References and bibliography	41

<b>3</b>	<b>The Internal Level</b>	<b>45</b>
3.1	Introduction	45
3.2	Database access: an overview	47
3.3	Page sets and files	51
3.4	Indexing	58
3.5	Hashing	67
3.6	Pointer chains	72
3.7	Compression techniques	75
3.8	Concluding remarks	80
	Exercises	80
	References and bibliography	82
	Answers to selected exercises	88

---

PART 2

<b>A</b>	<b>Sample Relational System</b>	<b>93</b>
----------	---------------------------------	-----------

<b>4</b>	<b>An Overview of DB2</b>	<b>95</b>
4.1	Background	95
4.2	Relational databases	96
4.3	The SQL language	100
4.4	Major system components	107
	References and bibliography	113

<b>5</b>	<b>Data Definition</b>	<b>117</b>
5.1	Introduction	117
5.2	Base tables	117
5.3	Indexes	121
5.4	Concluding remarks	123
	Exercises	123
	Answers to selected exercises	124

<b>6</b>	<b>Data Manipulation</b>	<b>127</b>
6.1	Introduction	127
6.2	Simple queries	128
6.3	Join queries	132
6.4	Built-in functions	136
6.5	Advanced features	139
6.6	Update operations	147
6.7	Concluding remarks	151
	Exercises	154
	Answers to selected exercises	157

<b>7</b>	<b>The System Catalog</b>	<b>165</b>
7.1	Introduction	165
7.2	Querying the catalog	167
7.3	Updating the catalog	168
	Exercises	169
	Answers to selected exercises	170

<b>8</b>	<b>Views</b>	<b>173</b>
8.1	Introduction	173
8.2	View definition	176
8.3	DML operations on views	178
8.4	Logical data independence	182
8.5	Advantages of views	184
	Exercises	185
	Answers to selected exercises	185

<b>9</b>	<b>Embedded SQL</b>	<b>189</b>
9.1	Introduction	189
9.2	Operations not involving cursors	192
9.3	Operations involving cursors	193
9.4	A comprehensive example	196
9.5	Dynamic SQL	199
9.6	Concluding remarks	201
	Exercises	202
	Answers to selected exercises	203

---

## PART 3

<b>Relational Database Management</b>	<b>207</b>
---------------------------------------	------------

<b>10</b>	<b>An Overview of INGRES</b>	<b>209</b>
10.1	Background	209
10.2	Data Definition	210
10.3	Data manipulation: retrieval operations	214
10.4	Data manipulation: update operations	221
10.5	Views	223
10.6	Embedded QUEL	224
	Exercises	227
	References and bibliography	228
	Answers to selected exercises	228

<b>11</b>	<b>Relational Data Structure</b>	<b>233</b>
11.1	An introductory example	233
11.2	Domains	235
11.3	Relations	239
11.4	Relational databases	245
	References and bibliography	245
<b>12</b>	<b>Relational Integrity Rules</b>	<b>249</b>
12.1	Introduction	249
12.2	Primary keys	249
12.3	Foreign keys	250
12.4	The two integrity rules	252
12.5	Implications of the two integrity rules	253
	References and bibliography	255
<b>13</b>	<b>Relational Algebra</b>	<b>257</b>
13.1	Introduction	257
13.2	A syntax for the relational algebra	260
13.3	Traditional set operations	262
13.4	Special relational operations	264
13.5	Examples	270
13.6	Concluding remarks	270
	Exercises	272
	References and bibliography	273
	Answers to selected exercises	276
<b>14</b>	<b>Relational Calculus</b>	<b>281</b>
14.1	Introduction	281
14.2	Tuple-oriented relational calculus	283
14.3	Examples	287
14.4	Relational calculus vs. relational algebra	290
14.5	Domain-oriented relational calculus	294
14.6	Query-By-Example	297
	Exercises	301
	References and bibliography	302
	Answers to selected exercises	303
<b>15</b>	<b>Relational Systems</b>	<b>313</b>
15.1	The relational model: a summary	313
15.2	Essentiality	315

15.3	Definition of a relational system	320
15.4	Enforcing the relational integrity rules	322
15.5	Some current relational products	325
	Exercises	327
	References and bibliography	327
	Answers to selected exercises	329

## 16

<b>Query Optimization</b>	333
16.1 Introduction	333
16.2 A simple example	334
16.3 The optimization process: an overview	335
16.4 Optimization in System R	340
16.5 Optimization in INGRES	343
16.6 Implementing the relational operators	345
References and bibliography	350

## 17

<b>Further Normalization</b>		361
17.1	<b>Introduction</b>	361
17.2	<b>Functional dependence</b>	364
17.3	<b>First, second, and third normal forms</b>	366
17.4	<b>Boyce/Codd normal form</b>	374
17.5	<b>Good and bad decompositions</b>	378
17.6	<b>Fourth normal form</b>	381
17.7	<b>Fifth normal form</b>	385
17.8	<b>Concluding remarks</b>	390
	<b>Exercises</b>	392
	<b>References and bibliography</b>	394
	<b>Answers to selected exercises</b>	402

---

### PART 4

<b>The Database Environment</b>	411
---------------------------------	-----

## 18

<b>Recovery and Concurrency</b>	413
18.1 <b>Introduction</b>	413
18.2 <b>Transaction recovery</b>	413
18.3 <b>System and media recovery</b>	417
18.4 <b>Three concurrency problems</b>	419
18.5 <b>Locking</b>	422
18.6 <b>Concluding remarks</b>	427
<b>Exercises</b>	428
<b>References and bibliography</b>	430
<b>Answers to selected exercises</b>	434