

国外经典计算机科学教材

An Introduction to Database Systems
(Eighth Edition)

数据库系统导论

(第八版 · 影印版)


[美] C.J. Date 著



中国电力出版社
www.infopower.com.cn

国外经

学教材



An Introduction to Database Systems

(Eighth Edition)

数据库系统导论

(第八版·影印版)

[美] C.J. Date 著



中国电力出版社

www.infopower.com.cn

An Introduction to Database Systems, 8th edition (ISBN 0-321-19784-4)

C.J. Date

Copyright © 2004 Pearson Education, Inc.

Original English Language Edition Published by Addison-Wesley.

All rights reserved.

Reprinting edition published by PEARSON EDUCATION NORTH ASIA LTD and CHINA ELECTRIC POWER PRESS, Copyright © 2006.

本书影印版由 Pearson Education 授权中国电力出版社在中国境内（香港、澳门特别行政区和台湾地区除外）独家出版、发行。

未经出版者书面许可，不得以任何方式复制或抄袭本书的任何部分。

本书封面贴有 Pearson Education 防伪标签，无标签者不得销售。

北京市版权局著作合同登记号 图字：01-2005-5624

For sale and distribution in the People's Republic of China exclusively (except Taiwan, Hong Kong SAR and Macao SAR).

仅限于中华人民共和国境内（不包括中国香港、澳门特别行政区和中国台湾地区）销售发行。

图书在版编目（CIP）数据

数据库系统导论：第 8 版 / （美）德特（Date, C.J.）著.

—影印本. —北京：中国电力出版社，2006

（国外经典计算机科学教材系列）

ISBN 7-5083-3877-4

I. 数... II. 德... III. 数据库系统—高等学校—教材—英文 IV. TP311.13

中国版本图书馆 CIP 数据核字（2005）第 134045 号

丛 书 名：国外经典计算机科学教材系列

书 名：数据库系统导论（第八版·影印版）

编 著：（美）C.J. Date

责任编辑：牛贵华

出版发行：中国电力出版社

地址：北京市三里河路 6 号 邮政编码：100044

电话：（010）88515918 传 真：（010）88518169

印 刷：汇鑫印务有限公司

开本尺寸：185×233

印 张：63.75

书 号：ISBN 7-5083-3877-4

版 次：2006 年 1 月北京第 1 版

2006 年 1 月第 1 次印刷


定 价：89.00 元

版权所有 翻印必究

The material covered in *An Introduction to Database Systems* is organized into six major parts:

- **Part I** (four chapters) provides a broad introduction to the concepts of database systems in general and relational systems in particular. It also introduces the standard database language, SQL.
- **Part II** (six chapters) consists of a detailed and very careful description of the relational model, which is not only the theoretical foundation underlying relational systems but is, in fact, the theoretical foundation for the entire database field.
- **Part III** (four chapters) discusses the general question of database design. Three chapters are devoted to design theory, and the fourth considers semantic modeling and the entity/relationship model.
- **Part IV** (two chapters) is concerned with transaction management (i.e., recovery and concurrency controls).
- **Part V** (eight chapters) shows how relational concepts are relevant to a variety of further aspects of database technology—security, distributed databases, temporal data, decision support, and so on.
- **Part VI** (three chapters) describes the impact of object technology on database systems. Chapter 25 describes object systems specifically; Chapter 26 considers the possibility of a *rapprochement* between object and relational technologies and discusses object/relational systems; and Chapter 27 addresses the relevance to databases of XML.

About the Author



C. J. DATE is an author, lecturer, researcher, and independent consultant specializing in relational database systems. An active member of the database community for nearly 35 years, C. J. Date devotes the major part of his career to exploring, expanding, and expounding the theory and practice of relational technology. He enjoys a reputation second to none for his ability to explain complex technical material in a clear and understandable fashion.

"[C. J. Date's] book is the flag bearer of relational theory and mathematical treatment in general...as well as the runaway leader in discussing the SQL standards. It exercises much more respect for careful language and the importance of concepts and principles in gaining mastery of the field."

—CARL ECKBERG, *San Diego State University*

"[The] 8th Edition is an excellent and comprehensive presentation of the contemporary database field. In particular, Date's chapters on types, relations, object databases, and object-relational databases together provide an exceptionally clear, self-contained exposition of the object-relational approach to databases."

—MARTIN K. SOLOMON, *Florida Atlantic University*

"Chris Date is the computer industry's most respected expert and thinker on database technology, and his book *An Introduction to Database Systems* continues to be the definitive work for those wanting a comprehensive and current guide to database systems."

—COLIN J. WHITE, *President, Intelligent Business Strategies*

"This is the best explanation of concurrency that I have seen in literature, and it covers the ground quite thoroughly."

—BRUCE O. LARSEN, *Stevens Institute of Technology*

"...both an indispensable read and an indispensable reference. No serious information systems or database practitioner should be without this book."

—DECLAN BRADY, *MICS, Systems Architect and Database Specialist, Fujitsu*

"The author's deep insights into the area, informal treatment of profound topics, open-ended discussions of critical issues, comprehensive and up-to-date contents, as well as rich annotations on bibliography have made the book most popular in the database area for more than two decades."

—QIANG ZHU, *The University of Michigan, Dearborn*

"[The book's] appeal is its comprehensiveness and the fact that it is very up-to-date with research developments. The latter factor is due mainly to [Date's] involvement with these developments, which gives him a unique opportunity to write about them."

—DAVID LIVINGSTONE, *University of Northumbria at Newcastle*

出版说明

新世纪的朝阳刚刚露出丝抹微红,如火如荼的全球信息化浪潮便汹涌而至,让人无时无刻不感受到新一轮产业革命的气息。如何在这场变革中占尽先机,既是对民族信息业的挑战,也是机遇。从而,作为民族信息产业发展基石的高等教育事业就被赋予了比以往更重的责任,对培养和造就我国 21 世纪的一代新人提出了更高的要求。但在计算机科学突飞猛进的同时,专业教材的发展却严重滞后,越来越成为人才培养的瓶颈。同时,以美国为代表的西方国家计算机科学教育经历了充分的发展,产生了一批有着巨大影响力的经典教材,因此,以批判、借鉴的态度有选择地引进这些国外经典计算机教材,将促进国内教学体系和国外接轨,大大推动我国计算机教育事业的发展。

中国电力出版社进入计算机图书市场已有近 6 个年头,通过坚持“高端、精品、经典”战略,致力于与国外著名出版机构合作,出版了大批博得计算机业界和教育界赞誉的作品。通过与信息技术教育界人士的广泛沟通,同时依托丰富的出版资源,中国电力出版社适时推出了“国外经典计算机科学教材”的出版计划。本次教材出版计划是美国最大的计算机教育出版机构——Pearson 教育集团(Addison-Wesley、Prentice-Hall 等皆为其下属子公司)合作,依托其数十年积累的大批经典教材资源,确保了教材选题的权威经典。

为保证这套教材的含金量,并做到有的放矢,我们在国内组织了由中国科学院、北京大学等一流院校教师组成的专家指导委员会,对高校课程教学体系做了系统、详细的调查,听取了众多教育专家、行业专家的意见,对教育部的教育规划进行了认真研究,并深入了解国外大学实际教学选用的教材状况,对国外教材做了理性的分析,确立了依托国家教育计划、传播先进教学理念、为培养符合社会需要的高素质创新型人才服务,来作为本次“国外经典计算机科学教材”出版计划的宗旨。

我们从 2002 年的下半年开始着手这套教材的策划工作,并多次组织了专家研讨会、座谈会等,分析现有教材的优点与不足,采其精华,并力争体现本套教材的质量和特色。

1. 深入理解国内的教学体系结构,并比照国外相同专业的课程设置,既具有现实的适用性,又立足发展眼光,具备一定的前瞻性。
2. 以计算机专业的核心课程为基础,同时配合专业教学计划,争取覆盖专业选修课程和专业任选课程。
3. 选取国外的最新教材版本,同时对照国内同专业课程的学时要求,对不适用的版本进行剔除,充分满足国内教学要求。
4. 根据专业对口和必须具备同课程教学经验的要求,严格挑选译者,并严把质量关,确保教材翻译的高质量。
5. 通过从原出版社网站下载勘误表及与原书作者进行沟通的方式,对原书中的错误一一做了修改。

6. 对教材出版的后期工作,如审校、编辑、排版、印刷进行了严格的质量把关。

经过专家指导委员会的集体讨论,并广泛听取广大高等院校师生的意见,反复比较,从数百种国外教材中遴选出数十种,列入第一阶段的出版计划。这些教材的作者无一不是学富五车的大师,如 Stallings, Date, Ullman, Aho, Bryant, Sedgewick 等,他们的作品均是一版再版,并被众多国外一流大学如 Stanford University, MIT, UC Bekerley, Carnegie Mellon Univeristy, University of Michigan 等采用为教材。拟订的第一阶段出版计划包括 30 种图书,内容覆盖程序设计、数据结构、操作系统、计算机体系结构、数据库、编译原理、软件工程、图形学、通信与网络、离散数学等计算机专业核心基础课程,基本满足国内计算机专业的教学要求。

此外,为了帮助广大任课教师加深对本系列教材的理解,减轻他们的备课难度,我们从国外出版机构引进了大批的课程教学辅助资料,并积极延请国内优秀教师,根据其使用该系列教材中的教学经验,着手编写更加适合国内应用状况的教辅材料。

由于我们对国内高校计算机教育存在认识深度上的不足,在选题、翻译、编辑加工出版等方面的工作中还有许多有待提高之处,恳请广大师生和读者提出批评和建议,并期待有更多的人加入到我们的工作中来。我们的联系方式是:

电子邮件: csbook@cepp.com.cn

联系电话: 010-88515918-300

联系地址: 北京市西城区三里河路 6 号中国电力出版社

邮政编码: 100044

***This book is dedicated to my wife Lindy
and to the memory of my mother Rene—***

***also to the memory of Ted Codd, who, sadly,
passed away as this book was going to press***

Those who cannot remember the past
are condemned to repeat it

Usually quoted in the form:

Those who don't know history are
doomed to repeat it

—George Santayana

I would like to see computer science
teaching set deliberately
in a historical framework. . .

Students need to understand
how the present situation has come
about, what was tried,
what worked and what did not, and
how improvements in hardware
made progress possible. The absence
of this element in

their training causes people to
approach every problem from
first principles. They are apt to
propose solutions that
have been found wanting in the past.

Instead of standing
on the shoulders of their precursors,
they try to go it alone.

—Maurice V. Wilkes

About the Author

C. J. Date is an independent author, lecturer, researcher, and consultant, specializing in relational database technology. He is based in Healdsburg, California.

In 1967, following several years as a mathematical programmer and programming instructor for Leo Computers Ltd. (London, England), Mr. Date moved to the IBM (UK) Development Laboratories, where he worked on the integration of database functionality into PL/I. In 1974 he transferred to the IBM Systems Development Center in California, where he was responsible for the design of a database language known as the Unified Database Language, UDL, and worked on technical planning and externals design for the IBM products SQL/DS and DB2. He left IBM in May, 1983.

Mr. Date has been active in the database field for well over 30 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 North America and also in Europe, Australia, Latin America, and the Far East. In addition to the present book, he is author or coauthor of a number of other database texts, including, from Morgan Kaufmann, *Temporal Data and the Relational Model* (2003) and, from Addison-Wesley, *Foundation for Future Database Systems: The Third Manifesto* (2nd edition, 2000), a detailed proposal for the future direction of the field; *Database: A Primer* (1983), which treats database systems from the nonspecialist's point of view; a series of *Relational Database Writings* books (1986, 1990, 1992, 1995, 1998), which deal with various aspects of relational technology in depth; and another series of books on specific systems and languages—*A Guide to DB2* (4th edition, 1993), *A Guide to SYBASE and SQL Server* (1992), *A Guide to SQL/DS* (1988), *A Guide to INGRES* (1987), and *A Guide to the SQL Standard* (4th edition, 1997). His books have been translated into several languages, including Braille, Chinese, Dutch, French, German, Greek, Italian, Japanese, Korean, Polish, Portuguese, Russian, and Spanish.

Mr. Date has also produced over 300 technical articles and research papers and has made a variety of original contributions to database theory. For several years, he was a regular columnist for the magazine *Database Programming & Design*. He also contributes regularly to the website <http://dbdebunk.com>. His professional seminars on database technology, offered both in North America and overseas, are widely considered to be second to none for the quality of the subject matter and the clarity of the exposition.

Mr. Date holds an Honours Degree in Mathematics from Cambridge University, England (BA 1962, MA 1966) and the honorary degree of Doctor of Technology from De Montfort University, England (1994).

Preface to the Eighth Edition

This book is a comprehensive introduction to the now very large field of database systems. It provides a solid grounding in the foundations of database technology and gives some idea as to how the field is likely to develop in the future. The book is meant primarily as a textbook, not a work of reference (though I hope and believe it can be useful as a reference also, to some extent). The emphasis throughout is on **insight** and **understanding**, not just on formalisms.

PREREQUISITES

The book as a whole is meant for anyone professionally interested in computing in some way who wants to gain an understanding of what database systems are all about. I assume you have at least a basic knowledge of both:

- The storage and file management capabilities (indexing, etc.) of a modern computer system
- The features of at least one high-level programming language (Java, Pascal, PL/I, etc.)

Regarding the first of these prerequisites, however, please note that a detailed tutorial on such matters can be found in the online Appendix D, “Storage Structures and Access Methods.”

STRUCTURE

I have to say that I am a little embarrassed at the size of this book. The fact is, however, that database technology has become a very large field, and it is not possible to do it justice in fewer than 1,000 pages or so (indeed, most of the book’s competitors are also about 1,000 pages). Be that as it may, the book overall is divided into six major parts:

- I. Basic Concepts
- II. The Relational Model
- III. Database Design
- IV. Transaction Management

V. Further Topics

VI. Objects, Relations, and XML

Each part in turn is divided into several chapters:

- Part I (four chapters) provides a broad introduction to the concepts of database systems in general and relational systems in particular. It also introduces the standard database language **SQL**.
- Part II (six chapters) consists of a detailed and very careful description of the **relational model**, which is not only the theoretical foundation underlying relational systems, but is in fact the theoretical foundation for the entire database field.
- Part III (four chapters) discusses the general question of **database design**; three chapters are devoted to design theory, the fourth considers semantic modeling and the entity/relationship model.
- Part IV (two chapters) is concerned with **transaction management** (i.e., recovery and concurrency controls).
- Part V (eight chapters) is a little bit of a *potpourri*. In general, however, it shows how relational concepts are relevant to a variety of further aspects of database technology—**security, distributed databases, temporal data, decision support**, and so on.
- Finally, Part VI (three chapters) describes the impact of **object technology** on database systems. Chapter 25 describes **object systems** specifically; Chapter 26 considers the possibility of a *rapprochement* between object and relational technologies and discusses **object/relational** systems; and Chapter 27 addresses the relevance to databases of **XML**.

There are also four appendixes. The first is an overview of a dramatic new and radically different implementation technology called **The TransRelational™ Model**; the next gives a **BNF grammar** for SQL expressions; the third contains a **glossary** of the more important abbreviations, acronyms, and symbols introduced in the body of the text; and the last is, as already explained, an online tutorial on **storage structures and access methods**.

ONLINE MATERIALS

The following instructor supplements are available only to qualified instructors. For information on accessing them, please contact your local Addison-Wesley Sales Representative, or send e-mail to aw.cse@aw.com.

- An Instructor's Manual provides guidance on how to use the book as a basis for teaching a database course. It consists of a series of notes, hints, and suggestions on each part, chapter, and appendix, as well as other supporting material.
- Answers to exercises (included in Instructor's Manual)
- Lecture slides in PowerPoint format

The following supplements are available to all readers of this book at <http://www.aw.com/cssupport>.

- Appendix D on storage structures and access methods (as already mentioned)
- Answers to a selected subset of the exercises

HOW TO READ THIS BOOK

The book overall is meant to be read in sequence more or less as written, but you can skip later chapters, and later sections within chapters, if you choose. A suggested plan for a first reading would be:

- Read Chapters 1 and 2 “once over lightly.”
- Read Chapters 3 and 4 very carefully (except perhaps for Sections 4.6 and 4.7).
- Read Chapter 5 “once over lightly.”
- Read Chapters 6, 7, 9, and 10 carefully, but skip Chapter 8 (except perhaps for Section 8.6 on SQL).
- Read Chapter 11 “once over lightly.”
- Read Chapters 12 and 14 carefully,¹ but skip Chapter 13.
- Read Chapters 15 and 16 carefully (except perhaps for Section 15.6 on two-phase commit).
- Read subsequent chapters selectively (but in sequence), according to taste and interest.

Each chapter opens with an introduction and closes with a summary; in addition, most chapters include exercises, and the online answers often give additional information about the topic of the exercise. Most chapters also include a set of references, many of which are annotated. This structure allows the subject matter to be treated in a layered fashion, with the most important concepts and results being presented “in line” in the main body of the text and various subsidiary issues and more complex aspects being deferred to the exercises or answers or references (as appropriate). *Note:* References are identified by two-part numbers in square brackets. For example, the reference “[3.1]” refers to the first item in the list of references at the end of Chapter 3: namely, a paper by E. F. Codd published in *CACM* 25, No. 2, in February, 1982. (For an explanation of abbreviations used in references—e.g., “CACM”—see Appendix C.)

COMPARISON WITH EARLIER EDITIONS

We can summarize the major differences between this edition and its immediate predecessor as follows:

- *Part I:* Chapters 1–4 cover roughly the same ground as Chapters 1–4 in the seventh edition, but they have been significantly revised at the detail level. In particular,

¹ You could also read Chapter 14 earlier if you like, possibly right after Chapter 4.

Chapter 4, the introduction to SQL, has been upgraded to the level of the current standard SQL:1999, as indeed has SQL coverage throughout the entire book. (This fact all by itself caused major revisions to more than half the chapters from the seventh edition.) *Note:* Facilities likely to be included in the next version of the standard—which will probably be ratified in late 2003—are also mentioned where appropriate.

- *Part II:* Chapters 5–10, on the relational model, are a totally rewritten, considerably expanded, and very much improved version of Chapters 5–9 from the seventh edition. In particular, the material on types—also known as domains—has been expanded into a chapter of its own (Chapter 5), and the material on integrity (Chapter 9) has been completely restructured and rethought. I hasten to add that the changes in these chapters do not represent changes in the underlying concepts but, rather, changes in how I have chosen to present them, based on my practical experience in teaching this material in live presentations.

Note: Some further words of explanation are in order here. Earlier editions of the book used SQL as a basis for teaching relational concepts, in the belief that it was easier on the student to show the concrete before the abstract. Unfortunately, however, the gulf between SQL and the relational model grew and continued to grow, ultimately reaching a point where I felt it would be actively misleading to use SQL for such a purpose any longer. The sad truth is that SQL is now so far from being a true embodiment of relational principles—it suffers from so many sins of both omission and commission—that I would frankly prefer not to discuss it at all! However, SQL is obviously important from a commercial point of view; thus, every database professional needs to have some familiarity with it, and it would just not be appropriate to ignore it in a book of this nature. I therefore settled on the strategy of including (a) a chapter on SQL basics in Part I of the book and (b) individual sections in other chapters, as applicable, describing those aspects of SQL that are specific to the subject of the chapter in question. In this way the book still provides comprehensive—indeed, extensive—coverage of SQL material, but puts that coverage into what I feel is the proper context.

- *Part III:* Chapters 10–13 are a mostly cosmetic revision of Chapters 9–12 from the seventh edition. There are detail-level improvements throughout, however.

Note: Again some further explanation is in order. Some reviewers of earlier editions complained that database design issues were treated too late. But it is my feeling that students are not ready to design databases properly or to appreciate design issues fully until they have some understanding of what databases are and how they are used; in other words, I believe it is important to spend some time on the relational model and related matters before exposing the student to design questions. Thus, I still believe Part III is in the right place. (That said, I do recognize that many instructors prefer to treat the entity/relationship material much earlier. To that end, I have tried to make Chapter 14 more or less self-contained, so that they can bring it in immediately after, say, Chapter 4.)

- *Part IV:* The two chapters of this part, Chapters 15 and 16, are completely rewritten, extended, and improved versions of Chapters 14 and 15 from the seventh edition. In

particular, Chapter 16 now includes a careful analysis of—and some unorthodox conclusions regarding—the so-called ACID properties of transactions.

- *Part V*: Chapter 20 on type inheritance and Chapter 23 on temporal databases have been totally rewritten to reflect recent research and developments in those areas. Revisions to other chapters are mostly cosmetic, though there are improvements in explanations and examples throughout and new material here and there.
- *Part VI*: Chapters 25 and 26 are improved and expanded versions of Chapters 24 and 25 from the seventh edition. Chapter 27 is new.

Finally, Appendix A is also new, while Appendixes B and C are revised versions of Appendixes A and C, respectively, from the seventh edition (the material from Appendix B in that edition has been incorporated into the body of the book). Appendix D is a significantly revised version of Appendix A from the *sixth* edition.

WHAT MAKES THIS BOOK DIFFERENT?

Every database book on the market has its own individual strengths and weaknesses, and every writer has his or her own particular ax to grind. One concentrates on transaction management issues; another stresses entity/relationship modeling; another looks at everything through an SQL lens; yet another takes a pure “object” point of view; still another views the field exclusively in terms of some commercial product; and so on. And, of course, I am no exception to this rule—I too have an ax to grind: what might be called the **foundation** ax. I believe very firmly that we must get the foundation right, and understand it properly, before we try to build on that foundation. This belief on my part explains the heavy emphasis in this book on the relational model; in particular, it explains the length of Part II—the most important part of the book—where I present my own understanding of the relational model as carefully as I can. I am interested in foundations, not fads and fashions. Products change all the time, but principles endure.

In this regard, I would like to draw your attention to the fact that there are several important (“foundation”) topics for which this book, virtually alone among the competition, includes an entire in-depth chapter (or an appendix, in one case). The topics in question include:

- Types
- Integrity
- Views
- Missing information
- Inheritance
- Temporal databases
- The TransRelational™ Model

In connection with that same point (the importance of foundations), I have to admit that the overall tone of the book has changed over the years. The first few editions were mostly descriptive in nature; they described the field as it actually was in practice, “warts and all.” Later editions, by contrast, were much more *prescriptive*; they talked about the way the field *ought* to be and the way it ought to develop in the future, if we did things right. And the present edition is certainly prescriptive in this sense (so it is a text with an attitude!). Since the first part of that “doing things right” is surely educating oneself as to what those right things are, I hope this new edition can help in that endeavor.

Yet another related point: As you might know, I recently published, along with my colleague Hugh Darwen, another “prescriptive” book, *Foundation for Future Database Systems: The Third Manifesto* (reference [3.3] in the present book).² That book, which we call *The Third Manifesto* or just the *Manifesto* for short, builds on the relational model to offer a detailed technical proposal for future database systems; it is the result of many years of teaching and thinking about such matters on the part of both Hugh and myself. And, not surprisingly, the ideas of the *Manifesto* permeate the present book. Which is not to say the *Manifesto* is a prerequisite to the present book—it is not; but it *is* directly relevant to much that is in the present book, and further related information is often to be found therein.

Note: Reference [3.3] uses a language called **Tutorial D** for illustrative purposes, and the present book does the same. **Tutorial D** syntax and semantics are intended to be more or less self-explanatory (the language might be characterized, loosely, as “Pascal-like”), but individual features are explained when they are first used if such explanation seems necessary.

A CLOSING REMARK

I would like to close these prefatory notes with the following lightly edited extract from another preface—Bertrand Russell’s own preface to *The Bertrand Russell Dictionary of Mind, Matter and Morals* (ed., Lester E. Denonn), Citadel Press, 1993, reprinted here by permission:

I have been accused of a habit of changing my opinions . . . I am not myself in any degree ashamed of [that habit]. What physicist who was already active in 1900 would dream of boasting that his opinions had not changed during the last half century? . . . [The] kind of philosophy that I value and have endeavoured to pursue is scientific, in the sense that there is some definite knowledge to be obtained and that new discoveries can make the admission of former error inevitable to any candid mind. For what I have said, whether early or late, I do not claim the kind of truth which theologians claim for their creeds. I claim only, at best, that the opinion expressed was a sensible one to hold at the time . . . I should be much surprised if subsequent research did not show that it needed to be modified. [Such opinions were not] intended as pontifical pronouncements, but only as the best I could do at the time towards the promotion of clear and accurate thinking. Clarity, above all, has been my aim.

² There is a website, too: <http://www.thethirdmanifesto.com>. See also <http://www.dbdebunk.com> for much related material.

If you compare earlier editions of this book with this eighth edition, you will find that I too have changed my opinions on many matters (and no doubt will continue to do so). I hope you will accept the remarks just quoted as adequate justification for this state of affairs. I share Bertrand Russell's perception of what the field of scientific inquiry is all about, but he expresses that perception far more eloquently than I could.

ACKNOWLEDGMENTS

Once again it is a pleasure to acknowledge my debt to the many people involved, directly or indirectly, in the production of this book:

- First of all, I must thank my friends David McGoveran and Nick Tindall for their major involvement in this edition; David contributed the first draft of Chapter 22 on decision support, and Nick contributed the first draft of Chapter 27 on XML. I must also thank my friend and colleague Hugh Darwen for major help (in a variety of forms) with all SQL portions of the manuscript. Nagraj Alur and Fabian Pascal also provided me with a variety of technical background material. A special vote of thanks goes to Steve Tarin for inventing the technology described in Appendix A and for his help in getting me to understand it fully.
- Second, the text has benefited from the comments of students on the seminars I have been teaching over the past several years. It has also benefited enormously from the comments of, and discussions with, numerous friends and reviewers, including Hugh Darwen, IBM; Guy de Tré, Ghent University; Carl Eckberg, San Diego State University; Cheng Hsu, Rensselaer Polytechnic Institute; Abdul-Rahman Itani, The University of Michigan-Dearborn; Vijay Kanabar, Boston University; Bruce O. Larsen, Stevens Institute of Technology; David Livingstone, University of Northumbria at Newcastle; David McGoveran, Alternative Technologies; Steve Miller, IBM; Fabian Pascal, independent consultant; Martin K. Solomon, Florida Atlantic University; Steve Tarin, Required Technologies; and Nick Tindall, IBM. Each of these people reviewed at least some part of the manuscript or made technical material available or otherwise helped me find answers to my many technical questions, and I am very grateful to all of them.
- I would also like to thank my wife Lindy for contributing the cover art once again and for her support throughout this and all my other database-related projects over the years.
- Finally, I am grateful (as always) to everyone at Addison-Wesley—especially Maite Suarez-Rivas and Katherine Harutunian—for all of their encouragement and support throughout this project, and to my packager Elisabeth Beller for another sterling job.

Healdsburg, California
2003

C. J. Date