


# SQL

F0E 453  
BAY TREE BOOKSTORE  
USED BOOK

## & ITS APPLICATIONS



Raymond A. Lorie  
Jean-Jacques Daudenarde

# ***SQL and Its Applications***

Raymond A. Lorie  
Jean-Jacques Daudenarde



PRENTICE HALL, Englewood Cliffs, New Jersey 07632

Lorie, Raymond A.

SQL and its applications / Raymond A. Lorie, Jean-Jacques Daudenarde.

p. cm.

Includes bibliographical references and index.

ISBN 0-13-837956-4

1. SQL (Computer program language) 2. Data base management.

I. Daudenarde, Jean-Jacques. II. Title.

QA76.73.S67L67 1991

005.75'6--dc20

90-39784

CIP

Editorial/production supervision: *Mary P. Rottino*

Cover design: *Lundgren Graphics, Inc*

Manufacturing buyers: *Kelly Behr and Susan Brunke*

Cover photo: *Larry Keenan/The Image Bank*



© 1991 by Prentice-Hall, Inc.

A Division of Simon & Schuster

Englewood Cliffs, New Jersey 07632

The publisher offers discounts on this book when ordered in bulk quantities. For more information, write:

Special Sales/College Marketing  
College Technical and Reference Division  
Prentice Hall  
Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

**ISBN 0-13-837956-4**

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada Inc., *Toronto*

Prentice-Hall Hispanoamericana, S.A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Simon & Schuster Asia Pte. Ltd., *Singapore*

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

*To*

*Ileana, Miguel, Ignacio, and Francisco*

*Hélène, Sophie, and Eric*

# Preface

Many programs that run on data processing systems today rely heavily on data that are stored "permanently" on some medium such as magnetic disks. The collection of such preserved data is called a *database*. Any system that helps the user in managing a database is a *database management system* or *DBMS*. Some of the most important functions of such a system deal with the creation of a database, the insertion of information into the database, the retrieval of the stored information, its protection against unauthorized access, and its recovery from system malfunctioning.

The relational model for the management of data was introduced in 1970 in a paper by E. F. Codd, published in the *Communication of the ACM*, 13:6 entitled "A Relational Model for Large Shared Data Banks." This original paper triggered an extensive research in the database field in the 1970's, both theoretical and practical. In the early 1980's, database management systems based on the relational model became available commercially. Starting modestly, it has now become a technology of major importance and several products are available from various companies such as IBM, Oracle, Ingres, and Teradata, to name a few. More and more data processing shops are installing these products, using them for developing applications, and/or querying their databases.

The data language used to interact with the database plays a fundamental role in a DBMS. Although several classes of languages have been proposed in the past, the *Structured Query Language*, or *SQL*, has emerged as the most commonly used relational language. Initially introduced in System R, a research



prototype developed at IBM Research in San Jose, California, SQL is now used in a large number of products, and standardization efforts are under way.

Up to now, most users of relational systems were specialists, accustomed to learn new systems by reading long user's manuals or papers published in the technical literature, or by gaining hands-on experience on systems installed in large data processing accounts. But the situation is changing. First, more users are getting involved with relational systems because more applications are being developed based on the new technology. Second, relational systems are becoming available on much smaller processors, even personal computers, opening the field considerably. The new breed of users consists of application analysts and programmers or end users who are specialists in their fields, but not necessarily in data processing or database technology. These professionals know their information management problems and want to know if and how relational systems - and SQL in particular - can help. This book caters to their needs, as well as the needs of anybody who wants simply to learn about the relational approach to data management at a pragmatic level.

The book is divided into three parts. Part 1 is introductory in nature. Part 2 presents the different constructs of the language. Many short examples are given; they all refer to the same application, making it easy to remember the context and understand what each example is about. While Part 2 studies each SQL statement individually in a convenient order, Part 3 looks at various applications. Each chapter analyzes one of these applications, with emphasis on a certain generic problem. Each time, important questions are considered, such as how to organize the data, or can a certain logical operation be done with a single SQL statement? If not, can a sequence of statements do the job? When is it necessary to write some code in a host language? The same questions will eventually come up during the design of any application, and the case studies presented here should prepare the reader to recognize the multiple trade-offs, and make the right choices.

When developing Part 2 we realized that we had to face a certain dilemma. SQL is a language which is still young; and although a large portion of the syntax and semantics of the language is common to all implementations, there are differences. We could have chosen to cover one particular instance of SQL, or we could have chosen to discuss only what is common to all important implementations, or to concentrate on only those constructs that have practically been standardized. Instead, we chose to introduce all the concepts that exist in most SQL. When the syntax - or the semantics - differs, we say so and show one particular instance. This approach allows us to be complete in covering the concepts; we feel this is important. We do not claim to be exhaustive in giving all features of all implementations. Consider, for example, the functions that apply to data of type "character string": once the reader

full exploitation of such relationships is not trivial; Chapter 14 discusses some interesting algorithms.

Finally, Chapter 15 goes through the design of an interactive, menu-driven application; it mainly illustrates the trade-offs between *static* and *dynamic* SQL and presents some details on how to use the dynamic approach. Chances are that you will never have to use dynamic SQL. But, just in case you have to, this chapter introduces the most important concepts.

Chapters 2 to 9 end with short series of exercises. All these refer to a common application context that is described in Chapter 2; the answers are given in Appendix.

### ***Acknowledgments***

We would like to thank our colleagues at IBM Research for keeping up an environment in which relational databases and SQL in particular have been - and still are - the base for interesting research; we have learned a lot from them. We are grateful to Charles Bontempo, Horacio Terrizzano, and a third anonymous reviewer for their many comments and suggestions. Some of these suggestions resulted in significant changes. We also acknowledge the help of Ignacio Terrizzano and David Torney, who read an early manuscript and whose comments have been particularly valuable. Finally, we are indebted to the IBM Corporation, and the management of the Research Division in particular, for providing some of the computer resources needed for the production of this book; we also thank the editors and staff at Prentice Hall for their constant support.

Raymond A. Lorie  
*IBM Corporation*  
*Research Division*  
*San Jose, California*

Jean-Jacques Daudenarde  
*IBM Corporation*  
*Storage Systems Products Division*  
*San Jose, California*

understands that there are such functions - for example to extract a substring - it is less important to give the exhaustive list of functions implemented in the various systems. Such a list is easily found in the reference manual specific to an implementation.

Again, the main goal of the book is to explain the concepts of SQL and show how these concepts can be helpful in solving a wide variety of information problems.

## Chapter organization

The first chapter provides an introduction to the relational way of organizing data. It describes a particular application and shows how its data can be couched in relational terms. Then, SQL is briefly introduced, showing examples of data definition, manipulation, and query capabilities. We believe it is important to know, from the start, how an SQL system handles the requests; for that reason, Chapter 1 covers the embedding of SQL in programs, as well as the issues of *dynamic* versus *static* SQL.

Chapter 2 considers another application and shows how the concept of *normalization* helps the designer to come up with an appropriate data organization.

Part 2 covers the language itself.

Chapter 3 introduces the basics of *data definition*. This is normally the first set of SQL facilities that a user would need when getting familiar with an SQL system in a hands-on session. It also discusses how data can be inserted into the database. At this point, the reader can create a table and populate it with data.

Chapter 4 discusses the basic *query* capabilities of SQL but only those that apply to a single table. Many examples are given; they are all based on the application introduced in Chapter 1.

Chapter 5 covers all other query features; in particular it introduces the notion of *join* and *subquery*.

Chapter 6 revisits the *data manipulation* facilities. They comprise all features that change the contents of the database, such as insertions, deletions, and updates. Since set-oriented data manipulation uses query capabilities to specify the set involved, we had to postpone their coverage until this point.

Chapter 7 studies some *performance* issues. At first, this chapter title may be somewhat surprising. After all, the scope of the book is the language, and



many performance issues are very much linked to the implementation. But there are operations that are intrinsically complex and we want the user to be able to recognize them. Also, some techniques that are used by practically all implementations are somewhat visible at the language level (indexing, locking); they need to be covered here.

Chapter 8 deals with *views*. A view allows the user to see the data in a way that is different from the way the data are actually stored.

Chapter 9 studies the *protection of data*. Data are an important asset and must be protected from any corruption or loss, as well as from unauthorized access. Although systems generally try to protect the data without putting any burden on the user, this cannot be done completely without some semantic knowledge that only the user has. This chapter discusses the SQL constructs used to communicate such knowledge to the system.

Part 3 of the book is dedicated to case studies.

Chapter 10 discusses a classical data processing example, with emphasis on simple *transactions* and *batch* processing. Chapter 11, on the contrary, focuses on *complex queries*. These two case studies, together with the examples of Part 2 should enable most users to develop straightforward applications.

The next chapters look at more advanced applications which all exhibit some kind of specific characteristics. Together, they enable the reader to build an arsenal of techniques that will be extremely useful in broad areas of application.

Chapters 12 and 13 look at problems which deal with interesting types of data. Chapter 12 is concerned with the management of *textual information*; Chapter 13 addresses the management of *object-oriented data*. Object data play a very important role in engineering applications, graphics, and office automation; the chapter looks at how to group the information into highly structured clusters. Although, in these chapters, the techniques are presented in specific application contexts, they are quite general in nature. The most likely situation is that some of these problems will come up when you design real-life applications, even if the context is apparently quite different.

Chapter 14 is concerned with a class of *graph* problems where a relationship is defined between instances of the same entity. Such a relationship can be represented by arcs between nodes in a graph. For example, a city A is connected to another city B, and B is connected to C, which is itself connected to D, or A, or B; or a part uses other parts in a *bill of material* application. The

# Contents

<b>Preface</b> .....	<b>xi</b>
<b>PART 1. INTRODUCTION</b> .....	<b>1</b>
<b>Chapter 1. Relational Data Management</b> .....	<b>1</b>
1.1 Introducing data management concepts .....	2
1.2 The relational model .....	6
1.3 Introduction to SQL .....	8
1.3.1 Creating a table .....	8
1.3.2 Inserting a row .....	8
1.3.3 Exploiting relationships .....	9
1.3.4 The <b>select</b> block .....	9
1.3.5 Aggregates .....	11
1.3.6 Sets .....	11
1.3.7 Subqueries .....	12
1.3.8 Joins .....	13
1.3.9 Order .....	15
1.4 Overview of an SQL system .....	16
1.4.1 Writing specific SQL applications .....	17
1.4.2 Program variables .....	18
1.4.3 Using a program .....	18
1.4.4 Static SQL .....	21
1.4.5 Dynamic SQL .....	21

1.4.6 Error handling .....	22
1.5 Summary .....	23
<b>Chapter 2. Database Design</b> .....	<b>25</b>
2.1 First Normal Form .....	27
2.2 Looking at the data in multiple ways .....	30
2.3 Functional dependencies .....	31
2.4 Second Normal Form .....	32
2.5 Third Normal Form .....	33
2.6 More than one key .....	34
2.7 Multivalued dependencies .....	35
2.8 Summary .....	37
2.9 Schema for the exercises .....	38
2.10 Exercises .....	40
<b>PART 2. THE LANGUAGE</b> .....	<b>43</b>
<b>Chapter 3. Constructing a Table</b> .....	<b>43</b>
3.1 Table creation .....	43
3.1.1 Table names .....	43
3.1.2 Column names .....	44
3.1.3 Data types .....	45
3.1.4 Missing information .....	46
3.1.5 Examples .....	47
3.1.6 Table deletion .....	48
3.2 Data insertion .....	49
3.2.1 Inserting from a program .....	50
3.3 Summary .....	51
3.4 Exercises .....	51
<b>Chapter 4. Elementary Queries</b> .....	<b>52</b>
4.1 Projection list .....	54
4.1.1 Selecting columns .....	54
4.1.2 Selecting computed values .....	54
4.1.3 Operators .....	55
4.1.4 Duplicate or distinct rows .....	57
4.2 Filters .....	58
4.2.1 Simple comparison expression .....	58
4.2.2 Operators on character strings .....	59
4.2.3 Operators on dates .....	60
4.2.4 Expressions .....	60
4.2.5 Using <b>between</b> .....	61
4.2.6 The <b>in</b> operator .....	62

4.2.7 Complex filters	62
4.2.8 Effect of <b>null</b>	62
4.3 Program interface	63
4.3.1 One-row query	64
4.3.2 Multiple-row query	65
4.3.3 Retrieving null values	65
4.4 Aggregates	67
4.4.1 General aggregation	67
4.4.2 Aggregation with null values	69
4.5 <b>Group by</b>	70
4.6 <b>Order by</b>	72
4.7 Summary	73
4.8 Exercises	75
<b>Chapter 5. Complex Queries</b>	<b>76</b>
5.1 The concept of join	77
5.1.1 Remarks	82
5.1.2 Null values in join	83
5.2 Uses of joins	83
5.2.1 Cartesian product	83
5.2.2 Decoding join	84
5.2.3 Multiplicative join	86
5.2.4 Selective join	88
5.2.5 Joins: a summary	89
5.3 Subqueries	90
5.3.1 Subquery returning a single row	90
5.3.2 Correlated subquery returning a single row	91
5.3.3 Subquery returning a set	92
5.3.4 Where the join does not help	95
5.3.5 More complex examples	96
5.4 Union	98
5.5 Summary	101
5.6 Exercises	102
<b>Chapter 6. Changing a Table</b>	<b>104</b>
6.1 Set insert	104
6.2 Delete and set delete	106
6.3 Update and set update	107
6.4 Delete based on a cursor	108
6.5 Update based on cursor	111
6.6 Adding a column to a table	113
6.7 Summary	113
6.8 Exercises	114

<b>Chapter 7. Controlling the Performance</b>	<b>115</b>
7.1 Introduction	115
7.2 Storage of rows	117
7.2.1 Files	117
7.2.2 Storing records in a file	118
7.3 Indexes	121
7.3.1 Index performance	122
7.3.2 Index maintenance	124
7.4 Hashing	124
7.5 Support of <b>order by</b>	126
7.5.1 Using a sort	126
7.5.2 Using an index	127
7.6 Query with a filter	130
7.7 Support of <b>group by</b>	132
7.8 Summary: indexes	133
7.9 Joins	134
7.9.1 Nested loops	134
7.9.2 The merge-join algorithm	136
7.10 Subqueries	137
7.11 Index creation	139
7.12 Summary	141
7.13 Exercises	141
 <b>Chapter 8. Views</b>	 <b>144</b>
8.1 Introduction	144
8.2 The general concept of view	147
8.3 Several uses of views	147
8.3.1 Selecting rows and columns	147
8.3.2 Computed columns	150
8.3.3 Joining tables	150
8.3.4 Aggregate functions, grouping	151
8.4 Summary	152
8.5 Exercises	153
 <b>Chapter 9. Data Protection</b>	 <b>154</b>
9.1 Introduction	154
9.2 Protection against user errors	155
9.3 Referential integrity	156
9.4 The notion of transaction	160
9.4.1 Unit of work, transaction	161
9.5 Protection against effects of concurrency	162
9.6 Locking	163
9.6.1 A simple lock protocol	164

9.6.2 Lock modes	164
9.6.3 Lock granularity	164
9.6.4 Practical advice	166
9.6.5 Deadlock	167
9.7 Protection against failures	168
9.7.1 Recovering from soft failures	168
9.7.2 Logging	169
9.7.3 Recovery from hard failures	170
9.8 Protection against unauthorized access	170
9.8.1 Granting a privilege	171
9.8.2 Revoking a privilege	174
9.9 Summary	174
9.10 Exercises	175

### **PART 3. APPLICATIONS** 177

#### **Chapter 10. Account Receivable** 177

10.1 Overview	177
10.2 Operations	179
10.3 The database design	179
10.4 Sales transaction	182
10.5 Batch processing	184
10.6 Recovery	189
10.7 Summary	190

#### **Chapter 11. Resource Allocation** 193

11.1 Find an object in a set	193
11.2 Find a sequence of objects	195
11.3 Storing allocated resources only	198
11.4 From discrete to continuous resources	200
11.5 Summary	203

#### **Chapter 12. Document Management** 204

12.1 Storing a document	206
12.1.1 Formatted data	206
12.2 Unformatted data	208
12.2.1 Long strings	208
12.2.2 Storing a very long string	209
12.3 Structuring the information	210
12.4 Text indexing	212
12.5 Progressive refinement of a search	217
12.6 Document control	219
12.7 Summary	222



## **x     SQL and Its Applications**

<b>Chapter 13. An Object-oriented Application</b>	<b>223</b>
13.1 A graphics application	223
13.1.1 Describing icons	225
13.1.2 Describing a drawing	228
13.2 The notion of object	229
13.3 Techniques for managing object data	230
13.3.1 Use of identifiers	230
13.3.2 Fetching an object	232
13.4 Summary	236
<b>Chapter 14. Graph Problems</b>	<b>238</b>
14.1 Path to the top of the tree	240
14.2 Finding all nodes under a node	241
14.3 Leaves of the tree	248
14.4 Acyclic graphs (bill of material)	248
14.5 Cyclic graphs (network problems)	249
14.6 Towards more support in SQL	251
14.7 Summary	252
<b>Chapter 15. Developing a User Interface</b>	<b>253</b>
15.1 Static SQL for parameterized queries	254
15.2 The use of wild characters	255
15.3 Supporting a set of parameters	257
15.3.1 Static versus dynamic SQL	260
15.3.2 Creating temporary tables	261
15.4 Ad hoc statement	263
15.4.1 Selection of multiple rows	263
15.4.2 Multiple execution of the same statement	266
15.4.3 Other cases	267
15.5 Ad hoc query for a known environment	267
15.6 Summary	269
<b>APPENDIX. Solution of Exercises</b>	<b>271</b>
<b>INDEX</b>	<b>293</b>

# ***Part 1. Introduction***

## **Chapter 1**

### **Relational Data Management**

Any human organization, from the simplest to the most complex, relies on information. In fact, any activity generates information about what happened, where, when, how, and so on. Any piece of information may be expressed as a sentence in a natural language. This representation may be satisfactory for some applications; for others it is not. In fact, for most applications, a better organization consists of representing the information as a collection of well-structured data. Additionally, since data generated today may be important for future activities and decisions, they must be recorded permanently, or at least until explicitly destroyed. The science - or art - of organizing permanent data is called *data management*.

It is fairly easy to save data on a permanent medium such as magnetic tape or disk. In general, the programming languages such as Cobol, C, and Fortran provide commands to save the contents of program variables on permanent storage and to read them back into memory. However, these commands generally allow the user to save the data in a sequential manner, one data element after another or one record (a collection of data elements) after another. Sometimes, more sophisticated commands allow a program to store data in a way that supports direct access to a record, given the record number. Nevertheless, although such data organizations are extremely useful, they provide only low-level functions, leaving most of the data management problems to the user.

## 2 SQL and Its Applications

Fortunately, the state of the art in data management has improved tremendously in the last decade. Years of research and development in theoretical as well as practical aspects of data management have yielded impressive results: today, sophisticated languages and systems exist, which simplify the user's job by making it possible to think about data at a much more conceptual level without worrying about a myriad of details.

The most successful of these systems today are based on what is called the **relational model of data** and rely on SQL (Structured Query Language). Their success is in part due to the simplicity of the model and the power of the operations available to store, manipulate, retrieve, and control the data. Another reason for their success is the internal sophistication of the systems that support the relational model and SQL in particular; it is this sophistication that enables high-level functions to be executed very efficiently at an affordable cost.

This chapter will guide the reader unfamiliar with data management through four sections. The first section uses an example to illustrate the most important concepts of data management without referring to any computer implementation. The second section introduces the relational model and describes the data of an application in relational terms. The third section goes through some SQL scenarios for data manipulation and queries. Finally, the last section briefly discusses the facilities generally provided by a data management system and by most SQL systems in particular.

### 1.1 Introducing data management concepts

Let us consider an organization and analyze the data needed to ensure its functioning. The organization is a company which rents tools. The company purchases tools and rents them to customers.

Actions generate the data that need to be kept in order to do business. When a tool is purchased, its existence must be recorded. This is done by maintaining a file of cards, each card recording the information associated with a tool. Since it is important to be able to refer to a particular tool in an unambiguous way, each tool must have a unique name. A simple way of naming a new tool is to give it a sequential number; the first tool receives the number 1, the next tool, 2, and so on. Each tool card with its tool number is stored in the file - let us call it the TOOL file - in order of increasing tool number. (By the way, this makes it very easy to find the next available number since one can just add 1 to the number on the last card in the file.) In addition to the tool number, the tool card will also contain data associated with the tool, such as