




CD-ROM includes
trial version of
Oracle®
Standard Edition

Introduction to Relational Databases and SQL Programming



Christopher Allen
Simon Chatwin
Catherine A. Creary

Mc
Graw
Hill

Technology
Education

Introduction	xvii
Chapter 1	
■ Introducing Relational Databases	
What Exactly Is a Database?	1
Tables	1
Rows/Records	2
Columns/Fields	2
How Is a Database Different from a Spreadsheet?	3
Many Rows	3
Many Users Simultaneously	4
Security	4
Relational Abilities	5
Constraints to Ensure Data Quality	6
Case Study—Clapham Specialty Store	6
Designing Your First Database	7
How Will Knowing This Help You?	8
When Developing Software	8
When Doing Database Administration	8
When Doing Business Analysis	8
If You Just Want to Know How to Use Databases Better	9
History of SQL	9
Chapter 1 Review	11
Chapter 2	
■ Storing and Retrieving Data: The Basics	
16	
Prepare to Work with a Database	17
Creating a Table	17
Inserting Records	18
Selecting Records	18
Dropping a Table	18
Creating Tables	20
Guidelines for Naming Tables and Columns	21
Creating a More Involved Table	24
Determining a Table's Structure	30
NULL and NOT NULL Columns	31
Inserting Data—Additional Techniques	32
How to Insert Records Containing NULL Values	32

Creating and Populating a Table	35
How to Insert Data That Contains Apostrophes	37
Viewing Data from a Table—Additional Techniques	37
Selecting Specific Columns	38
Changing Column Order	38
Performing Math Using Data in a Table	39
Connecting Two or More Pieces of Text Together	41
Assigning Aliases to Columns	42
Changing the Data Values You View	43
Chapter 2 Review	45
Chapter 3	
■ Performing Advanced Data Manipulation	
52	
SQL Command Categories	53
Data Definition	53
Data Manipulation	53
Data Control	53
Data Retrieval	54
Transaction Control	54
Limiting Which Records You Select	54
Filtering Records Based on Numbers	55
Filtering Records Based on Text	59
Filtering Records Based on Dates	60
Viewing Records in a Different Order	62
Sorting on Individual Columns	63
Sorting on Multiple Columns	63
Showing Only Unique Values	64
Selecting from the DUAL Table	66
Modifying Data in a Table	67
Removing Records from a Table	68
Deleting Rows Matching Specific Criteria	68
Deleting All Rows	68
Transaction Control	69
Undoing DML Transactions	69
Making Data Available to Others	72
Implicit and Explicit COMMITs	73
Performing Advanced Data Manipulation	74
Chapter 3 Review	75

Chapter 4		
■ Controlling SQL*Plus	80	
Editing Prior Commands	81	
Using a Text Editor	81	
Using the EDIT Command	82	
Line-Level Editing	83	
Copying and Pasting	85	
Using Your Mouse to Edit Text	86	
Clearing the SQL*Plus Screen	87	
Customizing the SQL*Plus Environment	87	
Customizing Using the SQL*Plus Menu	87	
Customizing Using Commands	89	
Saving Environment Customizations	89	
Producing More Readable Output	90	
Formatting Numbers in SQL*Plus	90	
Formatting Text in SQL*Plus	91	
Formatting Column Headings in SQL*Plus	92	
Spooling Output to Disk	94	
SQL Script Files	95	
Creating a Script File	95	
Running a Script File	96	
Using Variables in Script Files	96	
Chapter 4 Review	99	
Chapter 5		
■ Oracle SQL Built-in Functions	104	
Implementing Commonly Used Single-Row Functions	105	
System Variables	106	
Number Functions	108	
Text Functions	110	
Using Single-Row Functions	120	
Date	121	
Data Conversion	126	
Other Functions	132	
Implementing Commonly Used Group Functions	139	
Grouping Data via the GROUP BY Clause	142	
Including and Excluding Grouped Rows via the HAVING Clause	143	
Using Group Functions	145	
Chapter 5 Review	147	
Chapter 6		
■ Indexes, Joins, and Subqueries	152	
Creating the Test Tables	153	
Indexes	155	
Indexes in Databases	155	
How to Create Indexes	156	
Types of Indexes	157	
B-Tree Indexes	157	
Bitmap Indexes	159	
Bitmap Versus B-Tree Indexes	160	
Function-Based Indexes	160	
When to Use Indexes	161	
Relationships Between Tables	162	
Creating an Index	163	
Writing SELECT Statements to Display Data from More Than One Table	164	
Types of Joins	167	
Set Operators	173	
Writing Subqueries	177	
What Is a Subquery?	177	
Types of Problems Subqueries Can Solve	177	
Single-Row Subqueries	177	
Multirrow Subqueries	179	
Multicolumn Subqueries	180	
Correlated Subqueries	181	
Chapter 6 Review	183	
Chapter 7		
■ Creating a Program with PL/SQL	188	
What Is PL/SQL?	189	
Describing PL/SQL	192	
Who's Who in SQL, PL/SQL, and SQL*Plus	193	
Stored Procedures, Functions, and Triggers	194	
Stored Procedures and SQL Scripts	196	
Structure of a PL/SQL Block	196	
Header Section	197	
Declaration Section	197	
Execution Section	198	
Exception Section	198	
Creating a Simple PL/SQL Procedure	199	
Calling Procedures and Functions	200	
PL/SQL Variables and Constants	201	
Declaring PL/SQL Variables	202	
Declaring PL/SQL Constants	202	
Assigning Values to Variables	203	
Using Variables	204	
Control Structures in PL/SQL	205	
IF Statement	206	
Loops	208	
Cursors	210	
Nested Loops and Cursor Example	215	
Error Handling	217	
Exceptions	217	
System-Defined Exceptions	218	
Programmer-Defined Exceptions	220	
Creating a Programmer-Defined Exception	221	
Chapter 7 Review	223	

Chapter 8

■ Reading a Data Model 230

Overview of Data Model Design	231
Purpose and Benefits of Models	231
Relational Integrity: Quality Data	232
Types of Data Models	233
Conceptual Model	233
Logical Data Model	234
Physical Data Model	234
Reading an Entity Relationship Diagram	236
Entities	236
Attributes	237
Relationships	238
Cardinality and Optionality Notations	240
Reading an Entity Relationship Diagram	241
Chapter 8 Review	242

Chapter 9

■ Basics of Designing a Database's Structure 246

The Business Specification: Let the Data Tell You	
Where It Goes	247
Selecting the Database's Grain	247
Entities and Attributes	248
Identifying Records Reliably: Primary Keys	251
Why Do You Need a Primary Key?	251
Composite Primary Keys	251
Natural Primary Keys vs. Surrogate	
Primary Keys	252
Relationships: Referring to Data	
in Other Tables	254
Common Data Model Standards	255
Crow's Foot (IE)	256
IDEFIX	256
Relationships: Cardinality and Optionality	258
One-to-Many	258
Many-to-Many	259
One-to-One	260
Optionality	261
Dependency: Identifying Relationships	263
Recursive vs. Binary	264
Modeling Multiple Categories:	
Supertype and Subtypes	264
Creating Basic Data Models	265
Categories, Supertypes, and Subtypes	265
Implementing Super/Subtypes	
in a Physical Model	267
Chapter 9 Review	268

Chapter 10

■ Normalization 272

The Process of Normalization	273
Dependency	273
Dependents and Determinants	273
The First Three Normal Forms	275
First Normal Form: Eliminate	
Repeating Groups	275
Second Normal Form: Eliminate	
Redundant Data	277
Third Normal Form: Eliminate Attributes Not	
Dependent on the Primary Key	278
Apply the Normal Forms	
to a Database Model	279
The Fourth and Fifth Normal Forms	279
Fourth Normal Form: Isolate Independent	
Multiple Relationships	280
Fifth Normal Form: Isolate Semantically	
Related Multiple Relationships	282
The Rules You Really Need	283
Anomalies in the Data	283
Normalizing the Data	284
Insert	285
Delete	285
Update	285
A Tax on Being Law-Abiding	286
Moving from Logical to Physical Models	286
Choosing Your Engine	286
Changing Terminology	287
Translating Super- and Subtypes	287
Chapter 10 Review	289

Chapter 11

■ Analyzing Data Quality Issues 294

Datatypes and Missing Data: Quality Basics	295
Handling Missing Values	295
Apples and Oranges: Defining Datatypes	296
Choosing a Datatype	298
Creating a Table and Inserting Data	299
Converting Datatypes: Weak and	
Strong Typing	301
Data Domains: Sanity Checks	301
Domains as Sets of Values	301
Column and Table Constraints	302
Column Constraints	303
Table Constraints	305
Primary Key Constraints and Indexes	307
Uniqueness and How to Enforce It	307
Alternate Keys	310
Other Indexes	311

Foreign Key Constraints:	
Values from Other Tables	312
Adding the Constraint	313
Implementing Cardinality and Optionality	314
Cascading Effects	316
Cascade Delete on a Recursive Relationship	319
The Cascades That Don't	321
Creating the Movie Database	322
Declarative Relational Integrity: Pros and Cons	326
Declaring Foreign Key Constraints	327
Triggers and Procedural Code	327
Naming Constraints: Make It Easy for the	
Programmers	329
Naming Tables and Columns	330
Naming Check Constraints	330
Naming Foreign Key Constraints	330
Naming Indexes	331
Chapter 11 Review	332
 Chapter 12	
■ Other Useful Oracle Techniques 338	
Transferring Data Between Tables	339
Transferring Data Using INSERT	342
Creating a New Table Based on an Existing One	343
Renaming Tables	344
Altering a Table's Structure	345
Adding Columns	345
Changing Column Datatypes	345
Changing NULL Options	346
Views	348
Creating a View	349
Updateable Views	350
Dropping Views	351
Top N Analysis	352
Creating a View on a Table	353
Other Database Objects	358
Sequences	358
Synonyms	362
Chapter 12 Review	365
 Appendix	
■ On the CD-ROMs 370	
About Oracle 9i Standard Edition	
for Windows	370
System Requirements	370
Registering with the Oracle Technology Network	371
Installing Oracle 9i Standard Edition	372
 ■ Glossary 373	
 ■ Index 385	