

ADVANCED SQL

OBJECTIVES

- ▣ Define terms
- ▣ Write single and multiple table SQL queries
- ▣ Define and use three types of joins
- ▣ Write noncorrelated and correlated subqueries
- ▣ Understand and use SQL in procedural languages (e.g. PHP, PL/SQL)
- ▣ Understand triggers and stored procedures

PROCESSING MULTIPLE TABLES

- ❑ **Join**—a relational operation that causes two or more tables with a common domain to be combined into a single table or view
- ❑ **Equi-join**—a join in which the joining condition is based on equality between values in the common columns; common columns appear redundantly in the result table
- ❑ **Natural join**—an equi-join in which one of the duplicate columns is eliminated in the result table

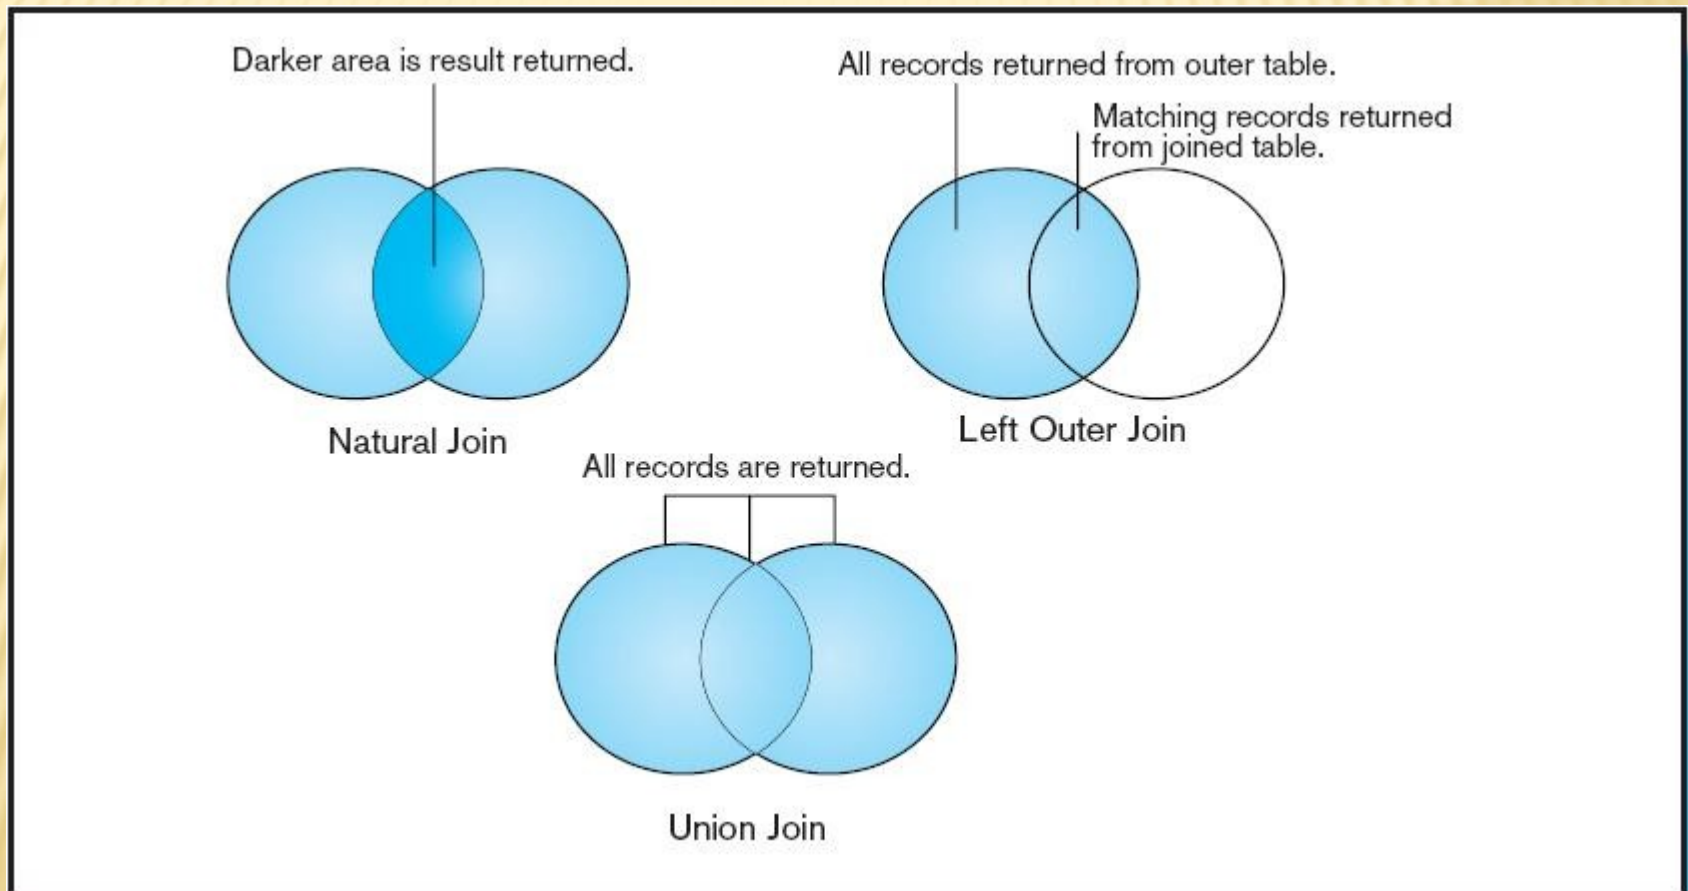
The common columns in joined tables are usually the primary key of the dominant table and the foreign key of the dependent table in 1:M relationships

PROCESSING MULTIPLE TABLES

- Outer join—a join in which rows that do not have matching values in common columns are nonetheless included in the result table (as opposed to *inner* join, in which rows must have matching values in order to appear in the result table)
- Union join—includes all columns from each table in the join, and an instance for each row of each table
- Self join—Matching rows of a table with other rows from the same table

Figure 7-2

Visualization of different join types with results returned in shaded area



THE FOLLOWING SLIDES CREATE TABLES FOR THIS ENTERPRISE DATA MODEL

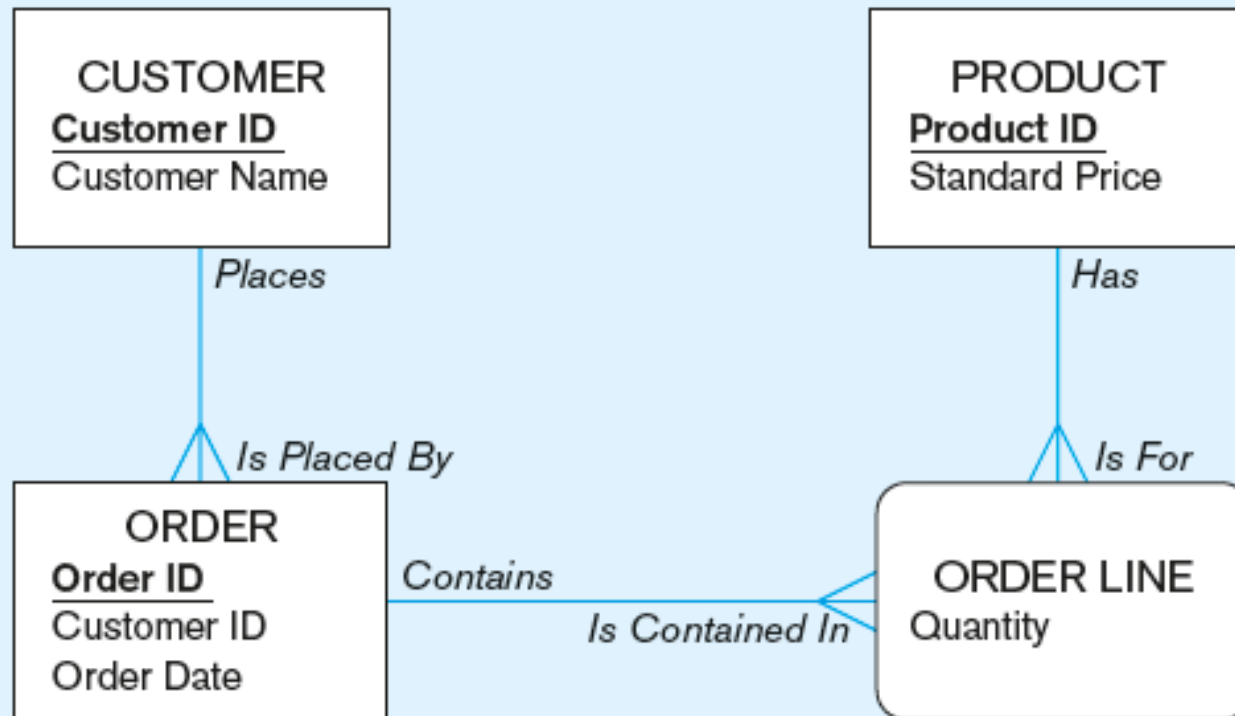


Figure 7-1 Pine Valley Furniture Company Customer_T and Order_T tables with pointers from customers to their orders

The image shows two database tables side-by-side. The left table, 'Order_T', has columns: OrderID, OrderDate, and CustomerID. The right table, 'Customer_T', has columns: CustomerID, CustomerName, CustomerAddress, CustomerCity, CustomerState, and CustomerPostalCode. Arrows indicate a one-to-many relationship where each customer has multiple orders.

OrderID	OrderDate	CustomerID
1001	10/21/2010	1
1002	10/21/2010	8
1003	10/22/2010	15
1004	10/22/2010	5
1005	10/24/2010	3
1006	10/24/2010	2
1007	10/27/2010	11
1008	10/30/2010	12
1009	11/5/2010	4
1010	11/5/2010	1
0		0

CustomerID	CustomerName	CustomerAddress	CustomerCity	CustomerState	CustomerPostalCode
1	Contemporary Casuals	1355 S Hines Blvd	Gainesville	FL	32601-2871
2	Value Furniture	15145 S.W. 17th St.	Plano	TX	75094-7743
3	Home Furnishings	1900 Allard Ave.	Albany	NY	12209-1125
4	Eastern Furniture	1925 Beltline Rd.	Carteret	NJ	07008-3188
5	Impressions	5585 Westcott Ct.	Sacramento	CA	94206-4056
6	Furniture Gallery	325 Flatiron Dr.	Boulder	CO	80514-4432
7	Period Furniture	394 Rainbow Dr.	Seattle	WA	97954-5589
8	California Classics	816 Peach Rd.	Santa Clara	CA	96915-7754
9	M and H Casual Furniture	3709 First Street	Clearwater	FL	34620-2314
10	Seminole Interiors	2400 Rocky Point Dr.	Seminole	FL	34646-4423
11	American Euro Lifestyles	2424 Missouri Ave N.	Prospect Park	NJ	07508-5621
12	Battle Creek Furniture	345 Capitol Ave. SW	Battle Creek	MI	49015-3401
13	Heritage Furnishings	66789 College Ave.	Carlisle	PA	17013-8834
14	Kaneohe Homes	112 Kiowai St.	Kaneohe	HI	96744-2537
15	Mountain Scenes	4132 Main Street	Ogden	UT	84403-4432
(New)					

These tables are used in queries that follow

EQUI-JOIN EXAMPLE

- For each customer who placed an order, what is the customer's name and order number?

```
SELECT Customer_T.CustomerID, Order_T.CustomerID,  
       CustomerName, OrderID  
FROM Customer_T, Order_T  
WHERE Customer_T.CustomerID = Order_T. CustomerID  
ORDER BY OrderID
```

Result:

CUSTOMERID	CUSTOMERID	CUSTOMERNAME	ORDERID
1	1	Contemporary Casuals	1001
8	8	California Classics	1002
15	15	Mountain Scenes	1003
5	5	Impressions	1004
3	3	Home Furnishings	1005
2	2	Value Furniture	1006
11	11	American Euro Lifestyles	1007
12	12	Battle Creek Furniture	1008
4	4	Eastern Furniture	1009
1	1	Contemporary Casuals	1010

10 rows selected.

Customer ID
appears twice in
the result

EQUI-JOIN EXAMPLE – ALTERNATIVE SYNTAX

```
SELECT Customer_T.CustomerID, Order_T.CustomerID,  
       CustomerName, OrderID  
FROM Customer_T INNER JOIN Order_T ON  
       Customer_T.CustomerID = Order_T.CustomerID  
ORDER BY OrderID;
```

INNER JOIN clause is an alternative to WHERE clause, and is used to match primary and foreign keys.

An INNER join will only return rows from each table that have matching rows in the other.

This query produces same results as previous equi-join example.

NATURAL JOIN EXAMPLE

For each customer who placed an order, what is the customer's name and order number?

Join involves multiple tables in FROM clause

```
SELECT Customer_T.CustomerID, CustomerName, OrderID
FROM Customer_T NATURAL JOIN Order_T ON
Customer_T.CustomerID = Order_T.CustomerID;
```

ON clause performs the equality check for common columns of the two tables

Note: from Fig. 7-1, you see that only 10 Customers have links with orders

→ Only 10 rows will be returned from this INNER join

OUTER JOIN EXAMPLE

List the customer name, ID number, and order number for all customers. Include customer information even for customers that do have an order.

```
SELECT Customer_T.CustomerID, CustomerName, OrderID  
FROM Customer_T LEFT OUTER JOIN Order_T  
WHERE Customer_T.CustomerID = Order_T.CustomerID;
```

LEFT OUTER JOIN
clause causes
customer data to
appear even if there is
no corresponding
order data

Unlike INNER join,
this will include
customer rows with
no matching order
rows

Outer Join Results

Unlike INNER join, this will include customer rows with no matching order rows

CUSTOMERID	CUSTOMERNAME	ORDERID
1	Contemporary Casuals	1001
1	Contemporary Casuals	1010
2	Value Furniture	1006
3	Home Furnishings	1005
4	Eastern Furniture	1009
5	Impressions	1004
6	Furniture Gallery	
7	Period Furniture	
8	California Classics	1002
9	M & H Casual Furniture	
10	Seminole Interiors	
11	American Euro Lifestyles	1007
12	Battle Creek Furniture	1008
13	Heritage Furnishings	
14	Kaneohe Homes	
15	Mountain Scenes	1003
16 rows selected.		

Syntax-LOJ

*SELECT columns FROM left_table LEFT OUTER
JOIN right_table ON
left_table.common_column =
right_table.common_column;*

Consider two tables: Employees & Departments.

```
SELECT Employees.EmployeeName,  
Departments.DepartmentName FROM Employees LEFT  
OUTER JOIN Departments ON  
Employees.DepartmentID = Departments.DepartmentID;
```

A Left Outer Join returns all records from the left table and the matched records from the right table. If there is no match, the result is NULL on the side of the right table.

Syntax-ROJ

```
SELECT columns FROM left_table RIGHT  
OUTER JOIN right_table ON  
left_table.common_column =  
right_table.common_column;  
Consider two tables: Employees & Departments.
```

```
SELECT Employees.EmployeeName,  
Departments.DepartmentName FROM Employees RIGHT  
OUTER JOIN Departments ON Employees.DepartmentID  
= Departments.DepartmentID;
```

A Right Outer Join returns all records from the right table and the matched records from the left table. If there is no match, the result is NULL on the side of the left table..

MULTIPLE TABLE JOIN

EXAMPLE
Assemble all information necessary to create an invoice for order number 1006

```
SELECT Customer_T.CustomerID, CustomerName, CustomerAddress,  
       CustomerCity, CustomerState, CustomerPostalCode, Order_T.OrderID,  
       OrderDate, OrderedQuantity, ProductDescription, StandardPrice,  
       (OrderedQuantity * ProductStandardPrice)  
FROM Customer_T, Order_T, OrderLine_T, Product_T  
WHERE Order_T.CustomerID = Customer_T.CustomerID  
      AND Order_T.OrderID = OrderLine_T.OrderID  
      AND OrderLine_T.ProductID = Product_T.ProductID  
      AND Order_T.OrderID = 1006;
```

Four tables
involved in
this join

Each pair of tables requires an equality-check condition in the WHERE clause, matching primary keys against foreign keys

Figure 7-4 Results from a four-table join (edited for readability)

From CUSTOMER_T table

CUSTOMERID	CUSTOMERNAME	CUSTOMERADDRESS	CUSTOMER CITY	CUSTOMER STATE	CUSTOMER POSTALCODE
2	Value Furniture	15145 S. W. 17th St.	Plano	TX	75094 7743
2	Value Furniture	15145 S. W. 17th St.	Plano	TX	75094 7743
2	Value Furniture	15145 S. W. 17th St.	Plano	TX	75094 7743

ORDERID	ORDERDATE	ORDERED QUANTITY	PRODUCTNAME	PRODUCT STANDARDPRICE	(QUANTITY* STANDARDPRICE)
1006	24-OCT -10	1	Entertainment Center	650	650
1006	24-OCT -10	2	Writer's Desk	325	650
1006	24-OCT -10	2	Dining Table	800	1600

From ORDER_T table

From PRODUCT_T table

SELF-JOIN EXAMPLE

Query: What are the employee ID and name of each employee and the name of his or her supervisor (label the supervisor's name Manager)?

```
SELECT E.EmployeeID, E.EmployeeName, M.EmployeeName AS Manager
FROM Employee_T E, Employee_T M
WHERE E.EmployeeSupervisor = M.EmployeeID;
```

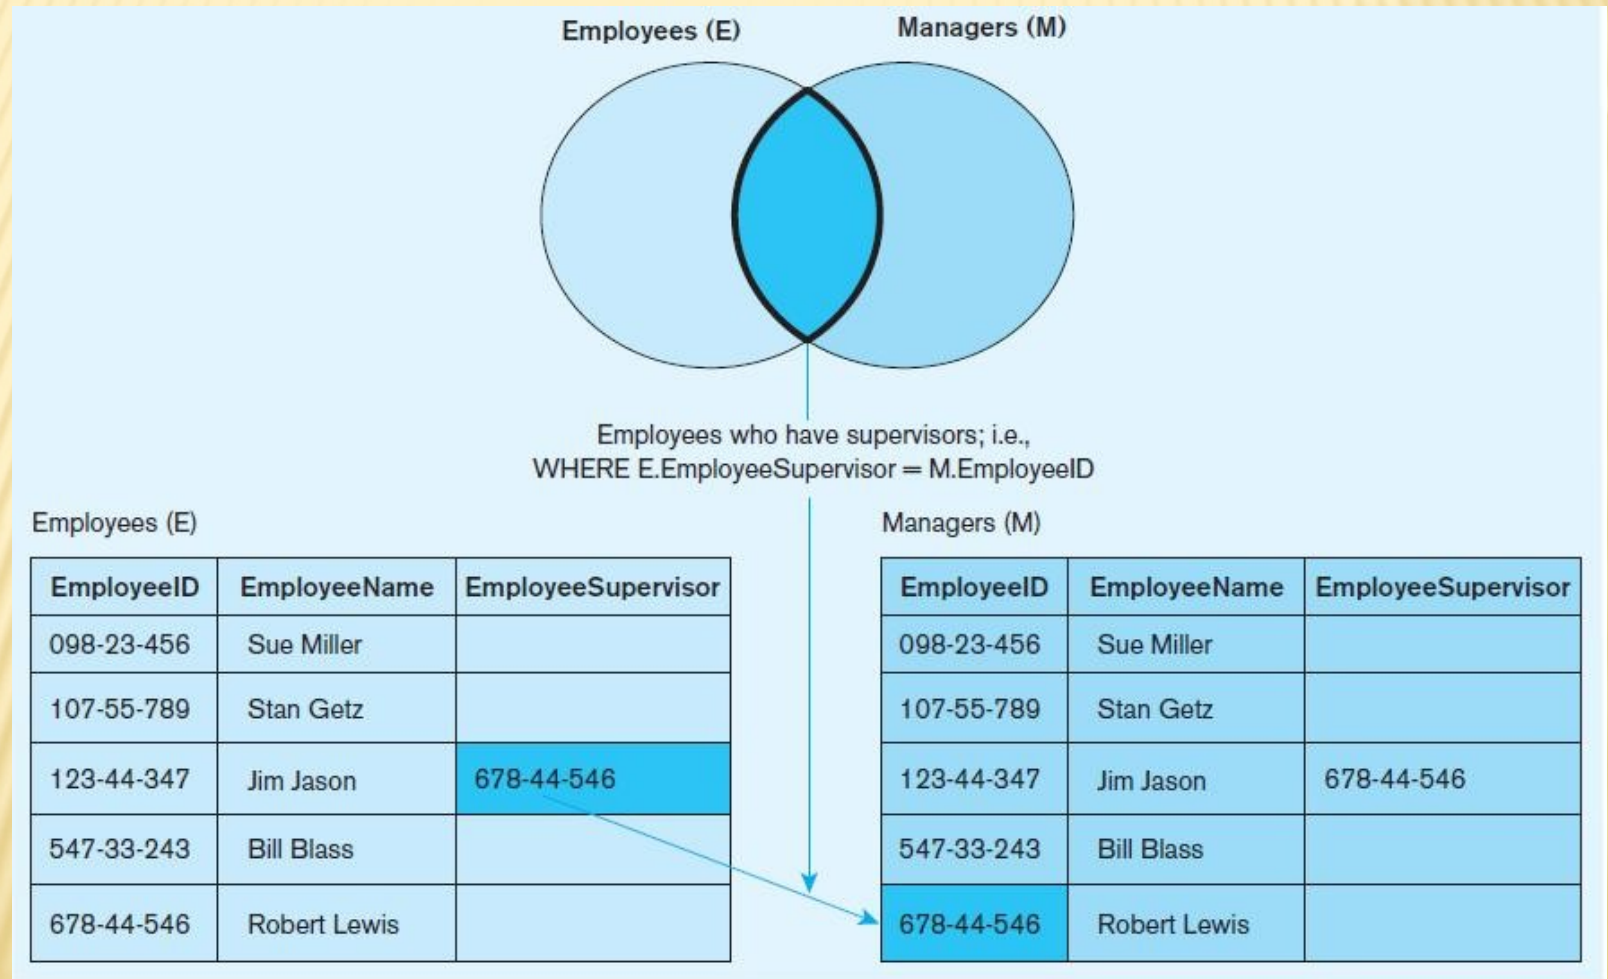
The same table is used on both sides of the join; distinguished using table aliases

Result:

EMPLOYEEID	EMPLOYEE NAME	MANAGER
123-44-347	Jim Jason	Robert Lewis

Self-joins are usually used on tables with unary relationships

Figure 7-5 Example of a self-join



PROCESSING MULTIPLE TABLES

USING SUBQUERIES

- ▮ Subquery—placing an inner query (SELECT statement) inside an outer query
- ▮ Options:
 - ▮ In a condition of the WHERE clause
 - ▮ As a “table” of the FROM clause
 - ▮ Within the HAVING clause
- ▮ Subqueries can be:
 - ▮ Noncorrelated—executed once for the entire outer query
 - ▮ Correlated—executed once for each row returned by the outer query

SUBQUERY EXAMPLE

□ Show all customers who have placed an order

```
SELECT CustomerName  
FROM Customer_T  
WHERE CustomerID IN
```

The IN operator will test to see if the CUSTOMER_ID value of a row is included in the list returned from the subquery

```
(SELECT DISTINCT CustomerID  
FROM Order_T);
```

Subquery is embedded in parentheses. In this case it returns a list that will be used in the WHERE clause of the outer query

Result:

<u>CUSTOMER_NAME</u>
Contemporary Casuals
Value Furniture
Home Furnishings
Eastern Furniture
Impressions
California Classics
American Euro Lifestyles
Battle Creek Furniture
Mountain Scenes
9 rows selected.

JOIN VS. SUBQUERY

Some queries could be accomplished by either a join or a subquery

Query: What are the name and address of the customer who placed order number 1008?

```
SELECT CustomerName, CustomerAddress, CustomerCity,  
       CustomerState, CustomerPostalCode  
FROM Customer_T, Order_T  
WHERE Customer_T.CustomerID = Order_T. CustomerID  
      AND OrderID = 1008;
```

Join version

Subquery version

```
SELECT CustomerName, CustomerAddress, CustomerCity,  
       CustomerState, CustomerPostalCode  
FROM Customer_T  
WHERE Customer_T.CustomerID =  
      (SELECT Order_T.CustomerID  
       FROM Order_T  
        WHERE OrderID = 1008);
```

Figure 7-6 Graphical depiction of two ways to answer a query with different types of joins

(a) Join query approach

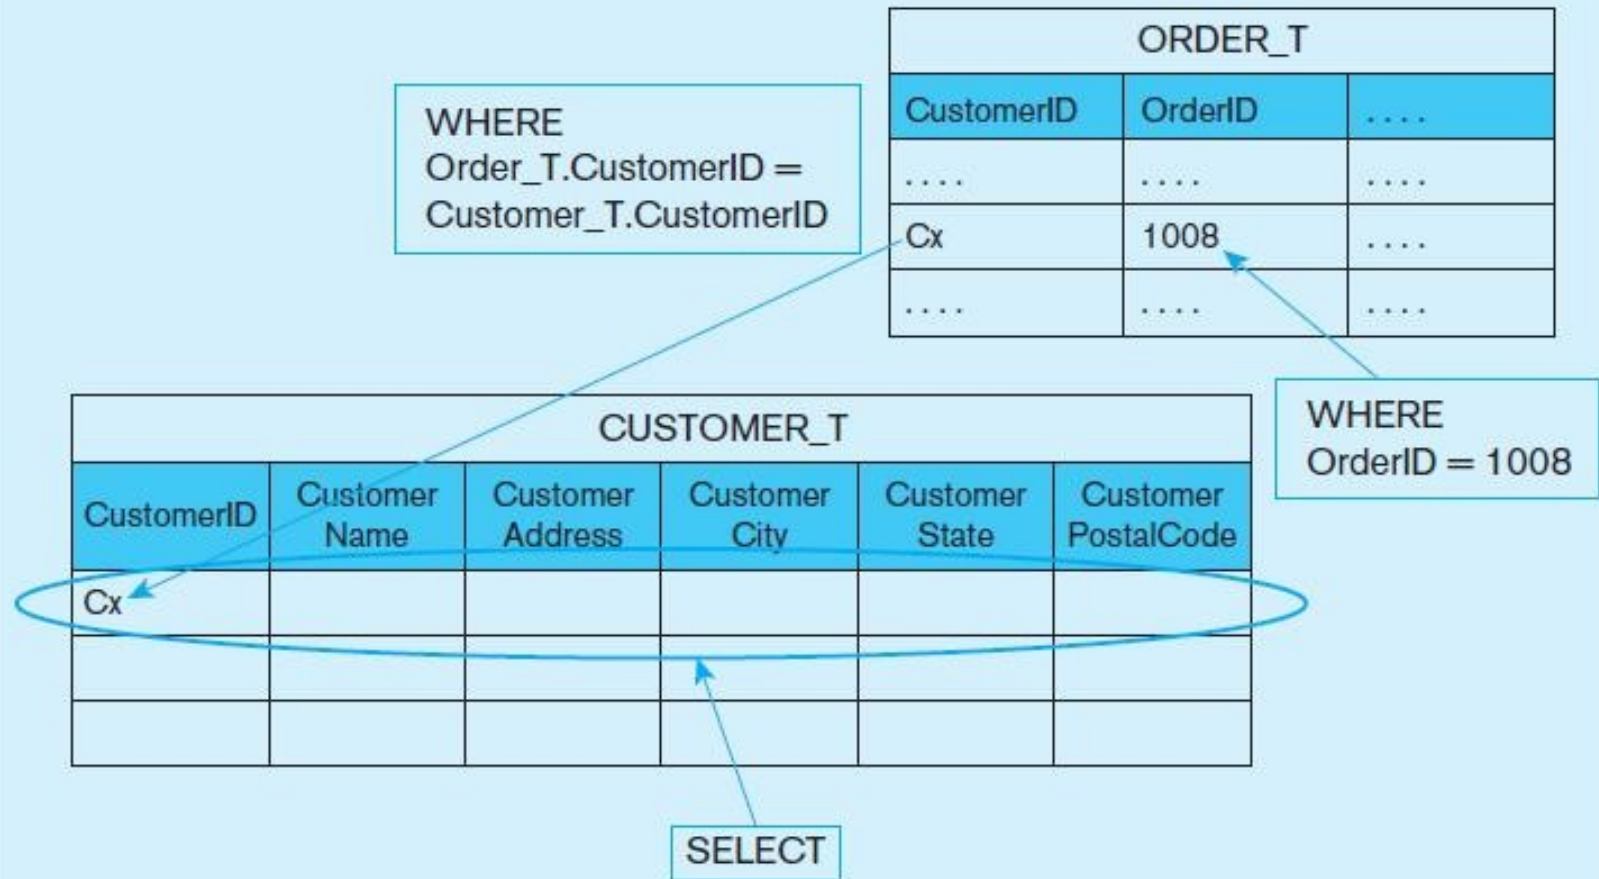
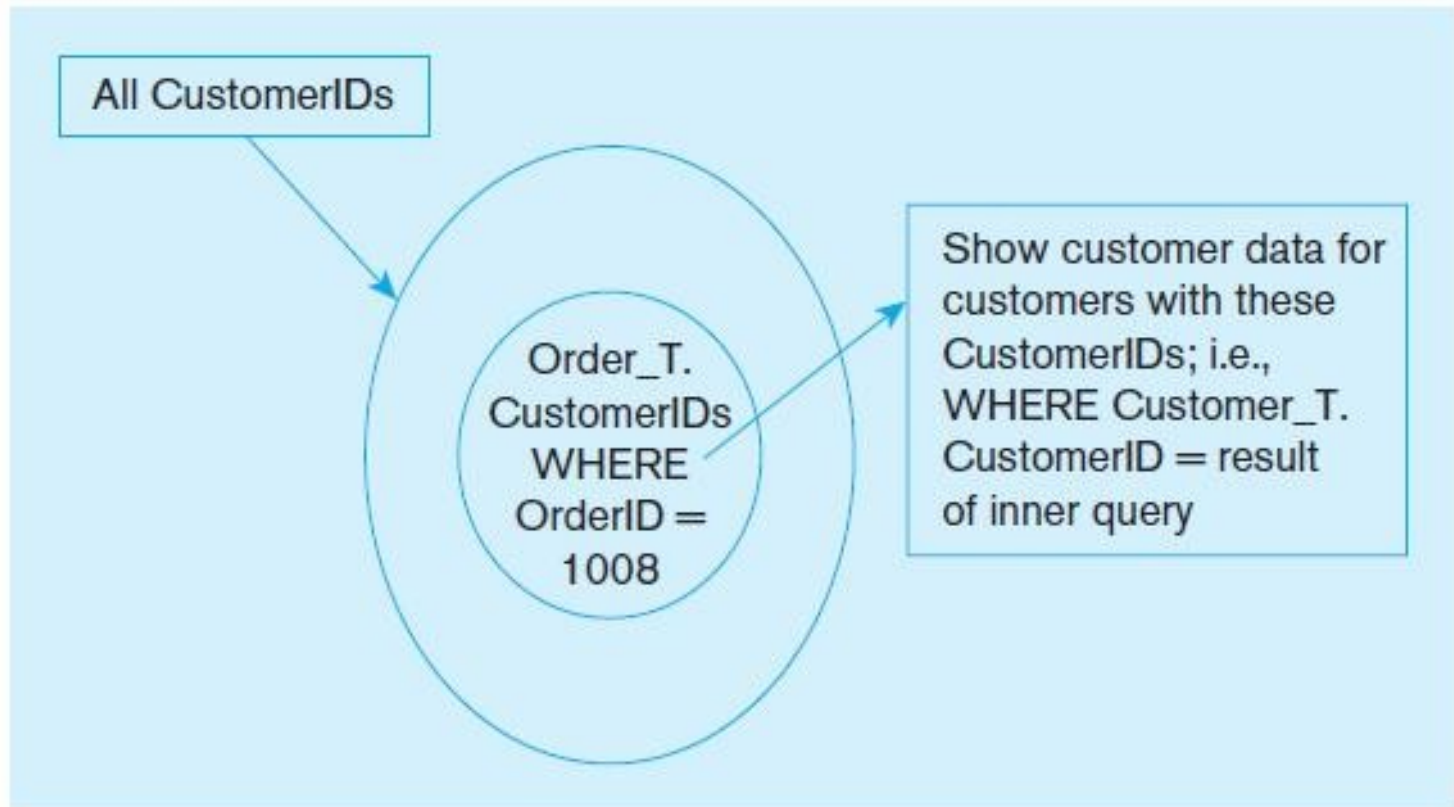


Figure 7-6 Graphical depiction of two ways to answer a query with different types of joins

(b) Subquery approach



CORRELATED VS. NONCORRELATED SUBQUERIES

- ❑ Noncorrelated subqueries:
 - ❑ Do not depend on data from the outer query
 - ❑ Execute once for the entire outer query
- ❑ Correlated subqueries:
 - ❑ Make use of data from the outer query
 - ❑ Execute once for each row of the outer query
 - ❑ Can use the EXISTS operator

Figure 7-8a Processing a noncorrelated subquery

What are the names of customers who have placed orders?

```
SELECT CustomerName
FROM Customer_T
WHERE CustomerID IN
```

```
(SELECT DISTINCT CustomerID
FROM Order_T);
```

1. The subquery (shown in the box) is processed first and an intermediate results table created:

CUSTOMERID

1
8
15
5
3
2
11
12
4

9 rows selected.

CustomerIDs
from orders



Show
names

2. The outer query returns the requested customer information for each customer included in the intermediate results table:

CUSTOMERNAME

Contemporary Casuals
Value Furniture
Home Furnishings
Eastern Furniture
Impressions
California Classics
American Euro Lifestyles
Battle Creek Furniture
Mountain Scenes
9 rows selected.

A noncorrelated subquery processes completely before the outer query begins

CORRELATED SUBQUERY

EXAMPLE

Show all orders that include furniture finished in natural ash

The EXISTS operator will return a TRUE value if the subquery resulted in a non-empty set, otherwise it returns a FALSE

```
SELECT DISTINCT OrderID FROM OrderLine_T
WHERE EXISTS
  (SELECT *
   FROM Product_T
   WHERE ProductID = OrderLine_T.ProductID
   AND Productfinish = 'Natural Ash');
```

➔ A correlated subquery always refers to an attribute from a table referenced in the outer query

The subquery is testing for a value that comes from the outer query

Figure 7-8b Processing a correlated subquery

What are the order IDs for all orders that have included furniture finished in natural ash?

```
SELECT DISTINCT OrderID FROM OrderLine_T
WHERE EXISTS
  (SELECT *
   FROM Product_T
    WHERE ProductID = OrderLine_T.ProductID
      AND ProductFinish = 'Natural Ash');
```

Subquery refers to outer-
query data, so executes once
for each row of outer query

OrderID	ProductID	OrderedQuantity
1001	1	1
1001	2	2
1001	4	1
1002	3	5
1003	3	3
1004	6	2
1004	8	2
1005	4	4
1005	4	1
1005	5	2
1007	1	3
1007	2	2
1008	3	3
1008	8	3
1009	4	2
1009	7	3
1010	8	10
*	0	0

	ProductID	ProductDescription	ProductFinish	ProductStandardPrice	ProductLineID
▶	1	End Table	Cherry	\$175.00	10001
⊞	2 → 2	Coffee Table	Natural Ash	\$200.00	20001
⊞	4 → 3	Computer Desk	Natural Ash	\$375.00	20001
⊞	4	Entertainment Center	Natural Maple	\$650.00	30001
⊞	5	Writer's Desk	Cherry	\$325.00	10001
⊞	6	8-Drawer Dresser	White Ash	\$750.00	20001
⊞	7	Dining Table	Natural Ash	\$800.00	20001
⊞	8	Computer Desk	Walnut	\$250.00	30001
*	(AutoNumber)			\$0.00	

1. The first order ID is selected from OrderLine_T: OrderID = 1001.
2. The subquery is evaluated to see if any product in that order has a natural ash finish. Product 2 does, and is part of the order. EXISTS is valued as *true* and the order ID is added to the result table.
3. The next order ID is selected from OrderLine_T: OrderID = 1002.
4. The subquery is evaluated to see if the product ordered has a natural ash finish. It does. EXISTS is valued as *true* and the order ID is added to the result table.
5. Processing continues through each order ID. Orders 1004, 1005, and 1010 are not included in the result table because they do not include any furniture with a natural ash finish. The final result table is shown in the text on page 302.

Note: only the
orders that
involve
products with
Natural Ash
will be
included in
the final
results

ANOTHER SUBQUERY EXAMPLE

□ Show all products whose standard price is higher than the average price

Subquery forms the derived table used in the FROM clause of the outer query

One column of the subquery is an aggregate function that has an alias name. That alias can then be referred to in the outer query

```
SELECT ProductDescription, ProductStandardPrice, AvgPrice
FROM
  (SELECT AVG(ProductStandardPrice) AvgPrice FROM Product_T),
  Product_T
WHERE ProductStandardPrice > AvgPrice;
```

The WHERE clause normally cannot include aggregate functions, but because the aggregate is performed in the subquery its result can be used in the outer query's WHERE clause

UNION QUERIES

- Combine the output (union of multiple queries) together into a single result table

```
SELECT C1.CustomerID, CustomerName, OrderedQuantity,  
'Largest Quantity' AS Quantity  
FROM Customer_T C1, Order_T O1, OrderLine_T Q1  
WHERE C1.CustomerID = O1.CustomerID  
AND O1.OrderID = Q1.OrderID  
AND OrderedQuantity =  
(SELECT MAX(OrderedQuantity)  
FROM OrderLine_T)
```

First query

Combine → UNION

```
SELECT C1.CustomerID, CustomerName, OrderedQuantity,  
'Smallest Quantity'  
FROM Customer_T C1, Order_T O1, OrderLine_T Q1  
WHERE C1.CustomerID = O1.CustomerID  
AND O1.OrderID = Q1.OrderID  
AND OrderedQuantity =  
(SELECT MIN(OrderedQuantity)  
FROM OrderLine_T)  
ORDER BY 3;
```

Second query

Figure 7-9 Combining queries using UNION

```
SELECT C1.CustomerID, CustomerName, OrderedQuantity, 'Largest Quantity' AS Quantity
FROM Customer_T C1, Order_T O1, OrderLine_T Q1
WHERE C1.CustomerID = O1.CustomerID
      AND O1.OrderID = Q1.OrderID
      AND OrderedQuantity =
          (SELECT MAX(OrderedQuantity)
           FROM OrderLine_T)
```

1. In the above query, the subquery is processed first and an intermediate results table created. It contains the maximum quantity ordered from OrderLine_T and has a value of 10.
2. Next the main query selects customer information for the customer or customers who ordered 10 of any item. Contemporary Casuals has ordered 10 of some unspecified item.

```
SELECT C1.CustomerID, CustomerName, OrderedQuantity, 'Smallest Quantity'
FROM Customer_T C1, Order_T O1, OrderLine_T Q1
WHERE C1.CustomerID = O1.CustomerID
      AND O1.OrderID = Q1.OrderID
      AND OrderedQuantity =
          (SELECT MIN(OrderedQuantity)
           FROM OrderLine_T)

ORDER BY 3;
```

1. In the second main query, the same process is followed but the result returned is for the minimum order quantity.
2. The results of the two queries are joined together using the UNION command.
3. The results are then ordered according to the value in OrderedQuantity. The default is ascending value, so the orders with the smallest quantity, 1, are listed first.

Note: with UNION queries, the quantity and data types of the attributes in the SELECT clauses of both queries must be identical

TIPS FOR DEVELOPING QUERIES

- ▮ Be familiar with the data model (entities and relationships)
- ▮ Understand the desired results
- ▮ Know the attributes desired in result
- ▮ Identify the entities that contain desired attributes
- ▮ Review ERD
- ▮ Construct a WHERE equality for each link
- ▮ Fine tune with GROUP BY and HAVING clauses if needed
- ▮ Consider the effect on unusual data

QUERY EFFICIENCY CONSIDERATIONS

- ▮ Instead of `SELECT *`, identify the specific attributes in the `SELECT` clause; this helps reduce network traffic of result set
- ▮ Limit the number of subqueries; try to make everything done in a single query if possible
- ▮ If data is to be used many times, make a separate query and store it as a view

GUIDELINES FOR BETTER QUERY DESIGN

- ▮ Understand how indexes are used in query processing
- ▮ Keep optimizer statistics up-to-date
- ▮ Use compatible data types for fields and literals
- ▮ Write simple queries
- ▮ Break complex queries into multiple simple parts
- ▮ Don't nest one query inside another query
- ▮ Don't combine a query with itself (if possible avoid self-joins)

GUIDELINES FOR BETTER QUERY DESIGN (CONT.)

- ▮ Create temporary tables for groups of queries
- ▮ Combine update operations
- ▮ Retrieve only the data you need
- ▮ Don't have the DBMS sort without an index
- ▮ Learn!
- ▮ Consider the total query processing time for ad hoc queries

ENSURING TRANSACTION

INTEGRITY

- ▮ Transaction = A discrete unit of work that must be completely processed or not processed at all
 - ▮ May involve multiple updates
 - ▮ If any update fails, then all other updates must be cancelled
- ▮ SQL commands for transactions
 - ▮ BEGIN TRANSACTION/END TRANSACTION
 - ▮ Marks boundaries of a transaction
 - ▮ COMMIT
 - ▮ Makes all updates permanent
 - ▮ ROLLBACK
 - ▮ Cancels updates since the last COMMIT

Figure 7-10 An SQL Transaction sequence (in pseudocode)

BEGIN transaction

INSERT OrderID, Orderdate, CustomerID into Order_T;

INSERT OrderID, ProductID, OrderedQuantity into OrderLine_T;

INSERT OrderID, ProductID, OrderedQuantity into OrderLine_T;

INSERT OrderID, ProductID, OrderedQuantity into OrderLine_T;

END transaction

Valid information inserted.
COMMIT work.



All changes to data
are made permanent.

Invalid ProductID entered.



Transaction will be ABORTED.
ROLLBACK all changes made to Order_T.



All changes made to Order_T
and OrderLine_T are removed.
Database state is just as it was
before the transaction began.

DATA DICTIONARY FACILITIES

- System tables that store metadata
- Users usually can view some of these tables
- Users are restricted from updating them
- Some examples in Oracle 11g

Table	Description
DBA_TABLES	Describes all tables in the database
DBA_TAB_COMMENTS	Comments on all tables in the database
DBA_CLUSTERS	Describes all clusters in the database
DBA_TAB_COLUMNS	Describes columns of all tables, views, and clusters
DBA_COL_PRIVS	Includes all grants on columns in the database
DBA_COL_COMMENTS	Comments on all columns in tables and views
DBA_CONSTRAINTS	Constraint definitions on all tables in the database
DBA_USERS	Information about all users of the database

TRIGGERS AND ROUTINES

- ▮ **Triggers**—routines that execute in response to a database event (INSERT, UPDATE, or DELETE)
- ▮ **Routines**
 - ▮ Program modules that execute on demand
- ▮ **Functions**—routines that return values and take input parameters
- ▮ **Procedures**—routines that do not return values and can take input or output parameters

Figure 7-11 Triggers contrasted with stored procedures (based on Mullins 1995)

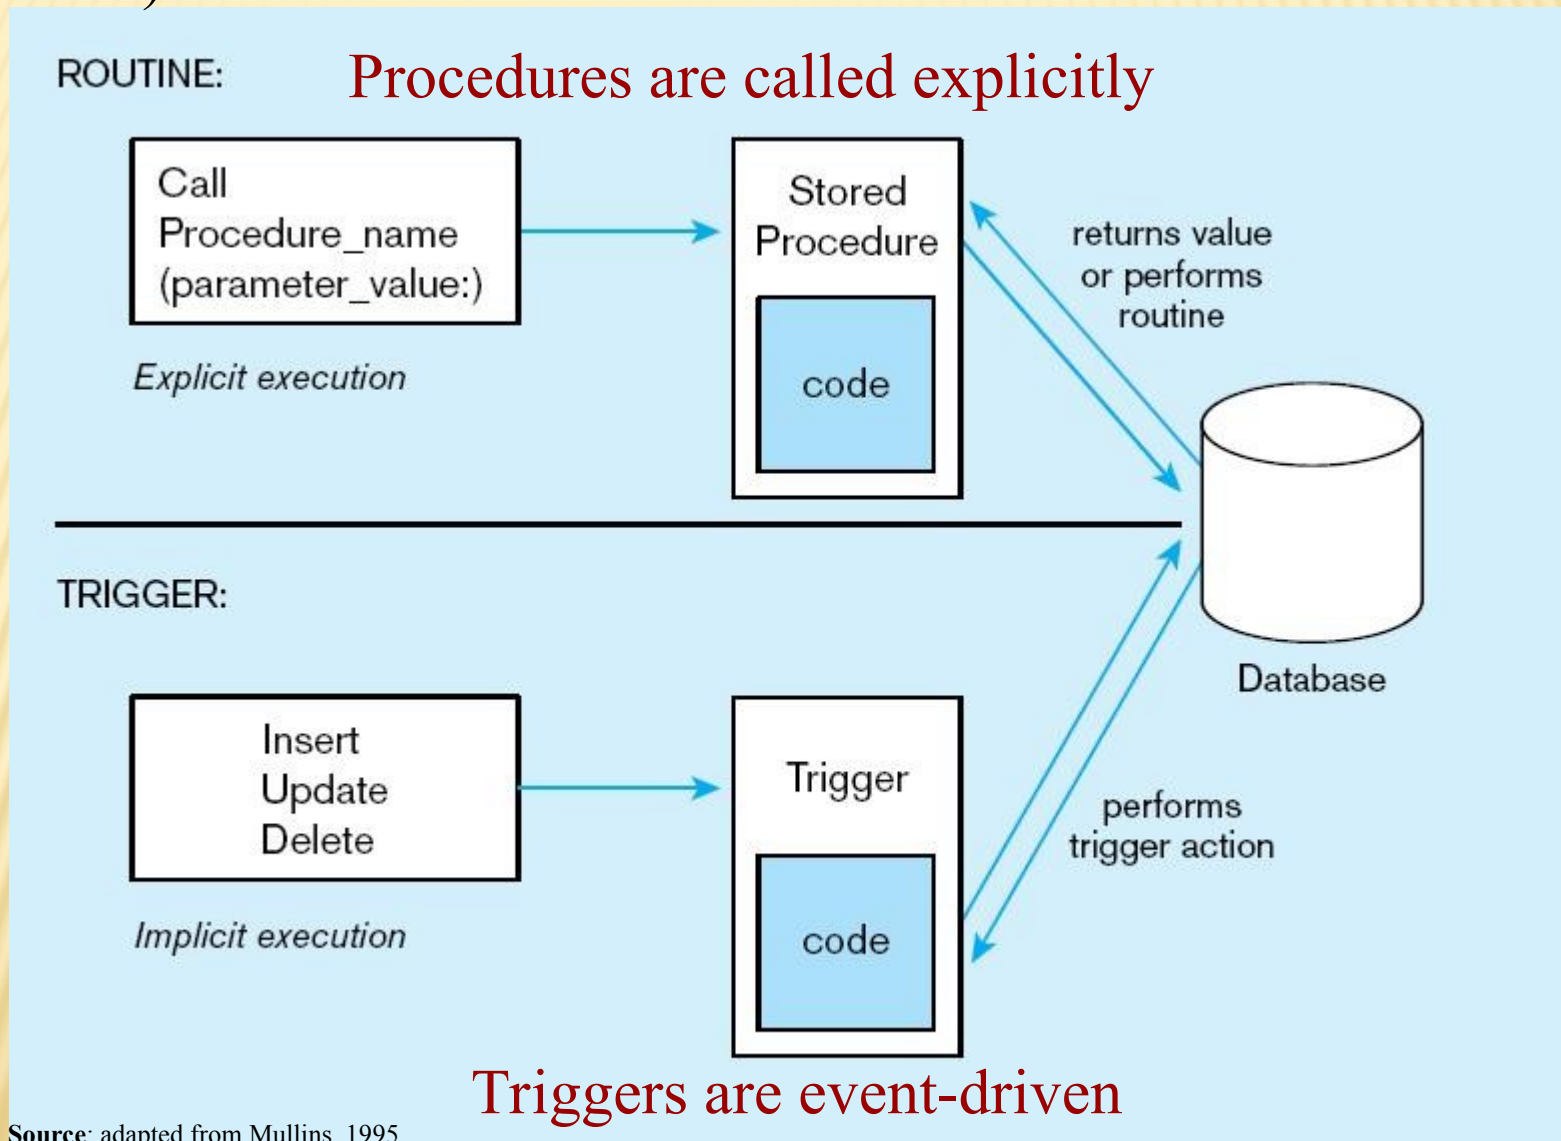


Figure 7-12 Trigger syntax in SQL:2008

```
CREATE TRIGGER trigger_name
    {BEFORE | AFTER | INSTEAD OF} {INSERT | DELETE | UPDATE} ON
    table_name
    [FOR EACH {ROW | STATEMENT}] [WHEN (search condition)]
    <triggered SQL statement here>;
```

Figure 7-13 Syntax for creating a routine, SQL:2008

```
{CREATE PROCEDURE | CREATE FUNCTION} routine_name
([parameter [{,parameter} . . .]])
[RETURNS data_type result_cast] /* for functions only */
[LANGUAGE {ADA | C | COBOL | FORTRAN | MUMPS | PASCAL | PLI | SQL}]
[PARAMETER STYLE {SQL | GENERAL}]
[SPECIFIC specific_name]
[DETERMINISTIC | NOT DETERMINISTIC]
[NO SQL | CONTAINS SQL | READS SQL DATA | MODIFIES SQL DATA]
[RETURNS NULL ON NULL INPUT | CALLED ON NULL INPUT]
[DYNAMIC RESULT SETS unsigned_integer] /* for procedures only */
[STATIC DISPATCH] /* for functions only */
[NEW SAVEPOINT LEVEL | OLD SAVEPOINT LEVEL]
routine_body
```

TABLE 7-2 Comparison of Vendor Syntax Differences in Stored Procedures

The vendors' syntaxes differ in stored procedures more than in ordinary SQL. For an illustration, here is a chart that shows what CREATE PROCEDURE looks like in three dialects. We use one line for each significant part, so you can compare dialects by reading across the line.

SQL:1999/IBM	MICROSOFT/SYBASE	ORACLE
CREATE PROCEDURE	CREATE PROCEDURE	CREATE PROCEDURE
Sp_proc1	Sp_proc1	Sp_proc1
(param1 INT)	@param1 INT	(param1 IN OUT INT)
MODIFIES SQL DATA BEGIN DECLARE num1 INT;	AS DECLARE @num1 INT	AS num1 INT; BEGIN
IF param1 <> 0	IF @param1 <> 0	IF param1 <> 0
THEN SET param1 = 1;	SELECT @param1 = 1;	THEN param1 :=1;
END IF		END IF;
UPDATE Table1 SET column1 = param1;	UPDATE Table1 SET column1 = @param1	UPDATE Table1 SET column1 = param1;
END		END

Source: Data from *SQL Performance Tuning* (Gulutzan and Pelzer, Addison-Wesley, 2002). Viewed at www.tdan.com/i023fe03.htm, June 6, 2007 (no longer available from this site).

EMBEDDED AND DYNAMIC SQL

▮ Embedded SQL

- ▮ Including hard-coded SQL statements in a program written in another language such as C or Java

▮ Dynamic SQL

- ▮ Ability for an application program to generate SQL code on the fly, as the application is running

REASONS TO EMBED SQL IN 3GL

- ❑ Can create a more flexible, accessible interface for the user
- ❑ Possible performance improvement
- ❑ Database security improvement; grant access only to the application instead of users



This work is protected by United States copyright laws and is provided solely for the use of instructors in teaching their courses and assessing student learning. Dissemination or sale of any part of this work (including on the World Wide Web) will destroy the integrity of the work and is not permitted. The work and materials from it should never be made available to students except by instructors using the accompanying text in their classes. All recipients of this work are expected to abide by these restrictions and to honor the intended pedagogical purposes and the needs of other instructors who rely on these materials.

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.

Copyright © 2014 Pearson Education, Inc.