



Chapter 3: Introduction to SQL

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan

See www.db-book.com for conditions on re-use



Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
 - Data Definition
 - Basic Query Structure
 - Additional Basic Operations
 - Set Operations
 - Null Values
 - Aggregate Functions
 - Modification of the Database
 - Nested Subqueries



History

- IBM Sequel language developed as part of System R project at the IBM San Jose Research Laboratory
- Renamed Structured Query Language (SQL)
- ANSI and ISO standard SQL:
 - SQL-86
 - SQL-89
 - **SQL-92**
 - SQL:1999 (language name became Y2K compliant!)
 - SQL:2003
- Commercial systems offer most, if not all, SQL-92 features, plus varying feature sets from later standards and special proprietary features.
 - Not all examples here may work on your particular system.



Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
 - domain types
 - create table with integrity constraints
 - drop table / alter table
- Basic Query Structure
- Additional Basic Operations
- Set Operations
- Null Values
- Aggregate Functions
- Modification of the Database
- Nested Subqueries



Domain Types in SQL

- **char(*n*)**. *Fixed length* character string, with user-specified length *n*.
- **varchar(*n*)**. *Variable length* character strings, with user-specified *maximum length n*.
- **int**. Integer (a finite subset of the integers that is machine-dependent).
- **smallint**. Small integer (a machine-dependent subset of the integer domain type).
- **numeric(*p,d*)**. Fixed point number, with user-specified precision of *p* digits, with *d* digits to the right of decimal point.
- **real, double precision**. Floating point and double-precision floating point numbers, with machine-dependent precision.
- **float(*n*)**. Floating point number, with user-specified precision of at least *n* digits.
- More are covered in Chapter 4.



Create Table Construct

- An SQL relation is defined using the **create table** command:

```
create table  $r$  ( $A_1 D_1, A_2 D_2, \dots, A_n D_n,$   
                (integrity-constraint1),  
                ...,  
                (integrity-constraintk))
```

- r is the name of the relation
- each A_i is an attribute name in the schema of relation r
- D_i is the data type of values in the domain of attribute A_i

- Example:

```
create table instructor (  
    ID          char(5),  
    name        varchar(20) not null,  
    dept_name varchar(20),  
    salary     numeric(8,2))
```

- **insert into instructor values** ('10211', 'Smith', 'Biology', 66000);
- **insert into instructor values** ('10211', null, 'Biology', 66000);



Integrity Constraints in Create Table

- **not null**
- **primary key** (A_1, \dots, A_n)
- **foreign key** (A_m, \dots, A_n) **references** r

Example: Declare *ID* as the primary key for *instructor*

```
create table instructor (  
    ID          char(5),  
    name       varchar(20) not null,  
    dept_name varchar(20),  
    salary     numeric(8,2),  
    primary key (ID),  
    foreign key (dept_name) references department)
```

primary key declaration on an attribute automatically ensures **not null**



Drop and Alter Table Constructs

- **drop table**

- **alter table**

- **alter table r add $A D$**

- ▶ where A is the name of the attribute to be added to relation r and D is the domain of A .
 - ▶ All tuples in the relation are assigned *null* as the value for the new attribute.

- **alter table r drop A**

- ▶ where A is the name of an attribute of relation r
 - ▶ Dropping of attributes not supported by many databases.



Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
- Basic Query Structure
 - select / from / where / natural join
- Additional Basic Operations
- Set Operations
- Null Values
- Aggregate Functions
- Nested Subqueries
- Modification of the Database



Basic Query Structure

- A typical SQL query has the form:

select A_1, A_2, \dots, A_n
from r_1, r_2, \dots, r_m
where P

- A_i represents an attribute
 - R_i represents a relation
 - P is a predicate.
- The result of an SQL query is a relation.
 - NOTE: SQL names are case insensitive



The select Clause

- The **select** clause list the attributes desired in the result of a query
 - corresponds to the **projection** operation of the relational algebra

- ```
select dept_name
from instructor
```



# The select Clause

- The **select** clause list the attributes desired in the result of a query
  - corresponds to the **projection** operation of the relational algebra

- **select distinct** *dept\_name*  
**from** *instructor*



# The select Clause

- The **select** clause list the attributes desired in the result of a query
  - corresponds to the **projection** operation of the relational algebra
- **select distinct dept\_name**  
**from instructor**
- SQL allows duplicates in relations as well as in query results.
- To force the elimination of duplicates, insert the keyword **distinct** after select.
- Find the names of all departments with instructor, and remove duplicates  
The keyword **all** specifies that duplicates not be removed.

**select all dept\_name**  
**from instructor**



# The select Clause (Cont.)

- An asterisk in the select clause denotes “all attributes”

```
select *
from instructor
```

- The **select** clause can contain arithmetic expressions involving the operation, +, −, \*, and /, and operating on constants or attributes of tuples.

- **select** *ID, name, salary/12*  
**from** *instructor*



# The where Clause

- The **where** clause specifies conditions that the result must satisfy
  - Corresponds to the ***selection predicate*** of the relational algebra.
- To find all instructors in Comp. Sci. dept with salary > 80000

```
select name
from instructor
where dept_name = 'Comp. Sci.' and salary > 80000
```
- Comparison results can be combined using the logical connectives **and**, **or**, and **not**.
- Comparisons can be applied to results of arithmetic expressions.



# The from Clause

- The **from** clause lists the relations involved in the query
  - Corresponds to the ***Cartesian product*** operation of the relational algebra.
- Find the Cartesian product *instructor X teaches*  
**select \***  
**from *instructor, teaches***
  - generates every possible instructor – teaches pair, with all attributes from both relations.
- Cartesian product not very useful directly, but useful combined with where-clause condition (selection operation in relational algebra).





# Joins

- For all instructors who have taught courses, find their names and the course ID of the courses they taught.

```
select name, course_id
from instructor, teaches
where instructor.ID = teaches.ID
```

*instructor*

| ID    | name       | dept_name  | salary |
|-------|------------|------------|--------|
| 10101 | Srinivasan | Comp. Sci. | 65000  |
| 12121 | Wu         | Finance    | 90000  |
| 15151 | Mozart     | Music      | 40000  |
| 22222 | Einstein   | Physics    | 95000  |
| 32343 | El Said    | History    | 60000  |
| 33456 | Guth       | Art        | 35000  |

*teaches*

| ID    | course_id | sec_id | semester | year |
|-------|-----------|--------|----------|------|
| 10101 | CS-101    | 1      | Fall     | 2009 |
| 10101 | CS-315    | 1      | Spring   | 2010 |
| 10101 | CS-347    | 1      | Fall     | 2009 |
| 12121 | FIN-201   | 1      | Spring   | 2010 |
| 15151 | MU-199    | 1      | Spring   | 2010 |
| 22222 | PHY-101   | 1      | Fall     | 2009 |



# Try Writing Some Queries in SQL

- Can you write the following queries in SQL?
  - Find the IDs of the courses taught by Srinivasan.
  - Find all the departments that offered courses in the Fall of 2009.

*instructor*

| <i>ID</i> | <i>name</i> | <i>dept_name</i> | <i>salary</i> |
|-----------|-------------|------------------|---------------|
| 10101     | Srinivasan  | Comp. Sci.       | 65000         |
| 12121     | Wu          | Finance          | 90000         |
| 15151     | Mozart      | Music            | 40000         |
| 22222     | Einstein    | Physics          | 95000         |
| 32343     | El Said     | History          | 60000         |

*teaches*

| <i>ID</i> | <i>course_id</i> | <i>sec_id</i> | <i>semester</i> | <i>year</i> |
|-----------|------------------|---------------|-----------------|-------------|
| 10101     | CS-101           | 1             | Fall            | 2009        |
| 10101     | CS-315           | 1             | Spring          | 2010        |
| 10101     | CS-347           | 1             | Fall            | 2009        |
| 12121     | FIN-201          | 1             | Spring          | 2010        |
| 15151     | MU-199           | 1             | Spring          | 2010        |
| 22222     | PHY-101          | 1             | Fall            | 2009        |



# Try Writing Some Queries in SQL

- Can you write the following queries in SQL?
  - Find the IDs of the courses taught by Srinivasan.  
**select distinct** course\_id  
**from** teachers, instructor  
**where** instructor.name = `Srinivasan' **and** instructor.ID = teachers.ID
  - Find all the departments that offered courses in the Fall of 2009.

*instructor*

| ID    | name       | dept_name  | salary |
|-------|------------|------------|--------|
| 10101 | Srinivasan | Comp. Sci. | 65000  |
| 12121 | Wu         | Finance    | 90000  |
| 15151 | Mozart     | Music      | 40000  |
| 22222 | Einstein   | Physics    | 95000  |
| 32343 | El Said    | History    | 60000  |

*teaches*

| ID    | course_id | sec_id | semester | year |
|-------|-----------|--------|----------|------|
| 10101 | CS-101    | 1      | Fall     | 2009 |
| 10101 | CS-315    | 1      | Spring   | 2010 |
| 10101 | CS-347    | 1      | Fall     | 2009 |
| 12121 | FIN-201   | 1      | Spring   | 2010 |
| 15151 | MU-199    | 1      | Spring   | 2010 |
| 22222 | PHY-101   | 1      | Fall     | 2009 |



# Try Writing Some Queries in SQL

## ■ Can you write the following queries in SQL?

- Find the IDs of the courses taught by Srinivasan.

```
select distinct course_id
from teachers, instructor
where instructor.name = `Srinivasan' and instructor.ID = teachers.ID
```

- Find all the departments that offered courses in the Fall of 2009.

```
select distinct dept_name
from instructor, teachers
where semester = `Fall' and year = 2009 and teaches.ID = instructor.ID
```

*instructor*

| ID    | name       | dept_name  | salary |
|-------|------------|------------|--------|
| 10101 | Srinivasan | Comp. Sci. | 65000  |
| 12121 | Wu         | Finance    | 90000  |
| 15151 | Mozart     | Music      | 40000  |
| 22222 | Einstein   | Physics    | 95000  |
| 32343 | El Said    | History    | 60000  |

*teaches*

| ID    | course_id | sec_id | semester | year |
|-------|-----------|--------|----------|------|
| 10101 | CS-101    | 1      | Fall     | 2009 |
| 10101 | CS-315    | 1      | Spring   | 2010 |
| 10101 | CS-347    | 1      | Fall     | 2009 |
| 12121 | FIN-201   | 1      | Spring   | 2010 |
| 15151 | MU-199    | 1      | Spring   | 2010 |
| 22222 | PHY-101   | 1      | Fall     | 2009 |



# Natural Join

- Natural join matches tuples with the same values for all common attributes, and retains only one copy of each common column
- **select \***  
**from instructor natural join teaches;**

| ID    | name       | dept_name  | salary | course_id | sec_id | semester | year |
|-------|------------|------------|--------|-----------|--------|----------|------|
| 10101 | Srinivasan | Comp. Sci. | 65000  | CS-101    | 1      | Fall     | 2009 |
| 10101 | Srinivasan | Comp. Sci. | 65000  | CS-315    | 1      | Spring   | 2010 |
| 10101 | Srinivasan | Comp. Sci. | 65000  | CS-347    | 1      | Fall     | 2009 |
| 12121 | Wu         | Finance    | 90000  | FIN-201   | 1      | Spring   | 2010 |
| 15151 | Mozart     | Music      | 40000  | MU-199    | 1      | Spring   | 2010 |
| 22222 | Einstein   | Physics    | 95000  | PHY-101   | 1      | Fall     | 2009 |
| 32343 | El Said    | History    | 60000  | HIS-351   | 1      | Spring   | 2010 |
| 45565 | Katz       | Comp. Sci. | 75000  | CS-101    | 1      | Spring   | 2010 |
| 45565 | Katz       | Comp. Sci. | 75000  | CS-319    | 1      | Spring   | 2010 |
| 76766 | Crick      | Biology    | 72000  | BIO-101   | 1      | Summer   | 2009 |
| 76766 | Crick      | Biology    | 72000  | BIO-301   | 1      | Summer   | 2010 |



# Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
- Basic Query Structure
- Additional Basic Operations
  - rename / string operations / ordering / predicates
- Set Operations
- Null Values
- Aggregate Functions
- Nested Subqueries
- Modification of the Database



# The Rename Operation (1/2)

- The SQL allows renaming relations and attributes using the **as** clause:

*old-name as new-name*

- E.g.,

- **select** *ID, name, salary/12 as monthly\_salary*  
**from** *instructor*



# The Rename Operation (2/2)

- Find the names of all instructors who have a higher salary than some instructor in 'Comp. Sci'.

*instructor*

| <i>ID</i> | <i>name</i> | <i>dept_name</i> | <i>salary</i> |
|-----------|-------------|------------------|---------------|
| 10101     | Srinivasan  | Comp. Sci.       | 65000         |
| 12121     | Wu          | Finance          | 90000         |
| 15151     | Mozart      | Music            | 40000         |
| 22222     | Einstein    | Physics          | 95000         |
| 32343     | El Said     | History          | 60000         |
| 33456     | Geller      | Physics          | 85000         |





# The Rename Operation (2/2)

- Find the names of all instructors who have a higher salary than some instructor in 'Comp. Sci'.

*instructor*

| <i>ID</i> | <i>name</i> | <i>dept_name</i> | <i>salary</i> |
|-----------|-------------|------------------|---------------|
| 10101     | Srinivasan  | Comp. Sci.       | 65000         |
| 12121     | Wu          | Finance          | 90000         |
| 15151     | Mozart      | Music            | 40000         |
| 22222     | Einstein    | Physics          | 95000         |
| 32343     | El Said     | History          | 60000         |
| 32456     | Gilbert     | Music            | 87000         |

- **select distinct** *T.name*  
**from** *instructor as T, instructor as S*  
**where** *T.salary > S.salary and S.dept\_name = 'Comp. Sci.'*
- Keyword **as** is optional and may be omitted  
*instructor as T*  $\equiv$  *instructor T*



# String Operations

- SQL includes a string-matching operator for comparisons on character strings. The operator “like” uses patterns that are described using two special characters:
  - percent (%). The % character matches any substring.
  - underscore (\_). The \_ character matches any character.
- Find the names of all instructors whose name includes the substring “dar”.

```
select name
from instructor
where name like '%dar%'
```



# Ordering the Display of Tuples

- List in alphabetic order the names of all instructors  
**select distinct** *name*  
**from** *instructor*  
**order by** *name*
- We may specify **desc** for descending order or **asc** for ascending order, for each attribute; ascending order is the default.
  - Example: **order by** *name desc*
- Can sort on multiple attributes
  - Example: **order by** *dept\_name, name*



# Where Clause Predicates

- SQL includes a **between** comparison operator
- Example: Find the names of all instructors with salary between \$90,000 and \$100,000 (that is,  $\geq \$90,000$  and  $\leq \$100,000$ )
  - **select** *name*  
**from** *instructor*  
**where** *salary* **between** 90000 **and** 100000
- Tuple comparison
  - **select** *name, course\_id*  
**from** *instructor, teaches*  
**where** (*instructor.ID, dept\_name*) = (*teaches.ID, 'Biology'*);



# Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
- Basic Query Structure
- Additional Basic Operations
- Set Operations
  - union / intersect / except
- Null Values
- Aggregate Functions
- Nested Subqueries
- Modification of the Database



# Set Operations

- Find courses that ran in Fall 2009 or in Spring 2010
- Find courses that ran in Fall 2009 and in Spring 2010
- Find courses that ran in Fall 2009 but not in Spring 2010



# Set Operations

- Find courses that ran in Fall 2009 or in Spring 2010

(**select** *course\_id* **from** *section* **where** *sem* = 'Fall' **and** *year* = 2009)

**union**

(**select** *course\_id* **from** *section* **where** *sem* = 'Spring' **and** *year* = 2010)

- Find courses that ran in Fall 2009 and in Spring 2010

- Find courses that ran in Fall 2009 but not in Spring 2010



# Set Operations

- Find courses that ran in Fall 2009 or in Spring 2010

(**select** *course\_id* **from** *section* **where** *sem* = 'Fall' **and** *year* = 2009)

**union**

(**select** *course\_id* **from** *section* **where** *sem* = 'Spring' **and** *year* = 2010)

- Find courses that ran in Fall 2009 and in Spring 2010

(**select** *course\_id* **from** *section* **where** *sem* = 'Fall' **and** *year* = 2009)

**intersect**

(**select** *course\_id* **from** *section* **where** *sem* = 'Spring' **and** *year* = 2010)

- Find courses that ran in Fall 2009 but not in Spring 2010





# Set Operations

- Find courses that ran in Fall 2009 or in Spring 2010

**(select *course\_id* from *section* where *sem* = 'Fall' and *year* = 2009)**

**union**

**(select *course\_id* from *section* where *sem* = 'Spring' and *year* = 2010)**

- Find courses that ran in Fall 2009 and in Spring 2010

**(select *course\_id* from *section* where *sem* = 'Fall' and *year* = 2009)**

**intersect**

**(select *course\_id* from *section* where *sem* = 'Spring' and *year* = 2010)**

- Find courses that ran in Fall 2009 but not in Spring 2010

**(select *course\_id* from *section* where *sem* = 'Fall' and *year* = 2009)**

**except**

**(select *course\_id* from *section* where *sem* = 'Spring' and *year* = 2010)**



# Set Operations

- Set operations **union**, **intersect**, and **except**
  - Each of the above operations automatically ***eliminates duplicates***
- To retain all duplicates use the corresponding multiset versions **union all**, **intersect all** and **except all**.



# Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
- Basic Query Structure
- Additional Basic Operations
- Set Operations
- Null Values
- Aggregate Functions
- Nested Subqueries
- Modification of the Database



# Null Values

- It is possible for tuples to have a null value, denoted by *null*, for some of their attributes
- *null* signifies an **unknown value** or that **a value does not exist**.
- The result of any arithmetic expression involving *null* is *null*
  - Example:  $5 + \text{null}$  returns null
- The predicate **is null** can be used to check for null values.
  - Example: Find all instructors whose salary is null.

```
select name
from instructor
where salary is null
```



# Null Values and Three Valued Logic

- Any comparison with ***null*** returns ***unknown***
  - Example:  $5 < null$  or  $null \diamond null$  or  $null = null$
- Three-valued logic using the truth value *unknown*:
  - OR:  $(unknown \text{ or } true) =$   
 $(unknown \text{ or } false) =$   
 $(unknown \text{ or } unknown) =$
  - AND:  $(true \text{ and } unknown) =$   
 $(false \text{ and } unknown) =$   
 $(unknown \text{ and } unknown) =$
  - NOT:  $(\text{not } unknown) =$



# Null Values and Three Valued Logic

- Any comparison with ***null*** returns ***unknown***
  - Example:  $5 < \text{null}$  or  $\text{null} \diamond \text{null}$  or  $\text{null} = \text{null}$
- Three-valued logic using the truth value *unknown*:
  - OR:  $(\text{unknown} \text{ or } \text{true}) = \text{true}$ ,  
 $(\text{unknown} \text{ or } \text{false}) = \text{unknown}$   
 $(\text{unknown} \text{ or } \text{unknown}) = \text{unknown}$
  - AND:  $(\text{true} \text{ and } \text{unknown}) = \text{unknown}$ ,  
 $(\text{false} \text{ and } \text{unknown}) = \text{false}$ ,  
 $(\text{unknown} \text{ and } \text{unknown}) = \text{unknown}$
  - NOT:  $(\text{not } \text{unknown}) = \text{unknown}$
  - “*P* is unknown” evaluates to true if predicate *P* evaluates to *unknown*
- Result of **where** clause predicate is treated as *false* if it evaluates to *unknown*



# Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
- Basic Query Structure
- Additional Basic Operations
- Set Operations
- Null Values
- Aggregate Functions
  - group by / having
- Modification of the Database
- Nested Subqueries



# Aggregate Functions

- These functions operate on the multiset of values of a column of a relation, and return a value

**avg:** average value

**min:** minimum value

**max:** maximum value

**sum:** sum of values

**count:** number of values





# Aggregate Functions (Cont.)

- Find the average salary of instructors in the Computer Science department
  - **select avg** (*salary*)  
**from** *instructor*  
**where** *dept\_name*= 'Comp. Sci.';
- Find the total number of instructors who teach a course in the Spring 2010 semester: assume table **teaches**(**ID**, **semester**, **year**)
- Find the number of tuples in the *course* relation
  - **select count** (\*)  
**from** *course*;



# Aggregate Functions (Cont.)

- Find the average salary of instructors in the Computer Science department
  - **select avg** (*salary*)  
**from** *instructor*  
**where** *dept\_name*= 'Comp. Sci.';
- Find the total number of instructors who teach a course in the Spring 2010 semester: assume table **teaches**(**ID**, **semester**, **year**)
  - **select count** (**distinct** *ID*)  
**from** *teaches*  
**where** *semester* = 'Spring' **and** *year* = 2010
- Find the number of tuples in the *course* relation
  - **select count** (\*)  
**from** *course*;



# Aggregate Functions – Group By

- Find the average salary of instructors in each department
  - **select** *dept\_name*, **avg** (*salary*)  
**from** *instructor*  
**group by** *dept\_name*;

| <i>ID</i> | <i>name</i> | <i>dept_name</i> | <i>salary</i> |
|-----------|-------------|------------------|---------------|
| 76766     | Crick       | Biology          | 72000         |
| 45565     | Katz        | Comp. Sci.       | 75000         |
| 10101     | Srinivasan  | Comp. Sci.       | 65000         |
| 83821     | Brandt      | Comp. Sci.       | 92000         |
| 98345     | Kim         | Elec. Eng.       | 80000         |
| 12121     | Wu          | Finance          | 90000         |
| 76543     | Singh       | Finance          | 80000         |
| 32343     | El Said     | History          | 60000         |
| 58583     | Califieri   | History          | 62000         |
| 15151     | Mozart      | Music            | 40000         |
| 33456     | Gold        | Physics          | 87000         |
| 22222     | Einstein    | Physics          | 95000         |



# Aggregate Functions – Group By

- Find the average salary of instructors in each department

- **select** *dept\_name*, **avg** (*salary*)  
**from** *instructor*  
**group by** *dept\_name*;

| <i>ID</i> | <i>name</i> | <i>dept_name</i> | <i>salary</i> |
|-----------|-------------|------------------|---------------|
| 76766     | Crick       | Biology          | 72000         |
| 45565     | Katz        | Comp. Sci.       | 75000         |
| 10101     | Srinivasan  | Comp. Sci.       | 65000         |
| 83821     | Brandt      | Comp. Sci.       | 92000         |
| 98345     | Kim         | Elec. Eng.       | 80000         |
| 12121     | Wu          | Finance          | 90000         |
| 76543     | Singh       | Finance          | 80000         |
| 32343     | El Said     | History          | 60000         |
| 58583     | Califieri   | History          | 62000         |
| 15151     | Mozart      | Music            | 40000         |
| 33456     | Gold        | Physics          | 87000         |
| 22222     | Einstein    | Physics          | 95000         |

| <i>dept_name</i> | <i>avg_salary</i> |
|------------------|-------------------|
| Biology          | 72000             |
| Comp. Sci.       | 77333             |
| Elec. Eng.       | 80000             |
| Finance          | 85000             |
| History          | 61000             |
| Music            | 40000             |
| Physics          | 91000             |



# Aggregation (Cont.)

- */\* erroneous query \*/*  
**select** *dept\_name, ID, avg (salary)*  
**from** *instructor*  
**group by** *dept\_name;*



# Aggregation (Cont.)

- Attributes in **select** clause outside of aggregate functions must appear in **group by** list
  - */\* erroneous query \*/*  
**select** *dept\_name, ID, avg (salary)*  
**from** *instructor*  
**group by** *dept\_name;*



# Aggregate Functions – Having Clause

- Find the names and average salaries of all departments whose average salary is greater than 42000



# Aggregate Functions – Having Clause

- Find the names and average salaries of all departments whose average salary is greater than 42000

```
select dept_name, avg (salary)
from instructor
group by dept_name
having avg (salary) > 42000;
```





# Aggregate Functions – Having Clause

- Find the names and average salaries of all departments whose average salary is greater than 42000

```
select dept_name, avg (salary)
from instructor
group by dept_name
having avg (salary) > 42000;
```

Note: predicates in the **having** clause are applied after the formation of groups whereas predicates in the **where** clause are applied before forming groups



# Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
- Basic Query Structure
- Additional Basic Operations
- Set Operations
- Null Values
- Aggregate Functions
- Modification of the Database
  - deletion / insertion / update
- Nested Subqueries



# Modification of the Database – Deletion

- Delete all instructors

**delete from** *instructor*

- Delete all instructors from the Finance department

**delete from** *instructor*

**where** *dept\_name*= 'Finance';



# Modification of the Database – Insertion

- Add a new tuple to *course*

**insert into** *course*

**values** ('CS-437', 'Database Systems', 'Comp. Sci.', 4);

- or equivalently

**insert into** *course* (*course\_id*, *title*, *dept\_name*, *credits*)

**values** ('CS-437', 'Database Systems', 'Comp. Sci.', 4);

- Add a new tuple to *student* with *tot\_creds* set to null

**insert into** *student*

**values** ('3003', 'Green', 'Finance', *null*);



# Modification of the Database – Updates

- Increase salaries of instructors whose salary is over \$100,000 by 3%, and all others receive a 5% raise

- Write two **update** statements:

```
update instructor
 set salary = salary * 1.03
 where salary > 100000;
update instructor
 set salary = salary * 1.05
 where salary <= 100000;
```

- The order is important
- Can be done better using the **case** statement (next slide)



# Case Statement for Conditional Updates

- Same query as before but with case statement

```
update instructor
 set salary = case
 when salary <= 100000 then salary * 1.05
 else salary * 1.03
 end
```



# Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
- Basic Query Structure
- Additional Basic Operations
- Set Operations
- Null Values
- Aggregate Functions
- Modification of the Database
- Nested Subqueries
  - all / some / exists / not exists / derived relations / with clause



# Nested Subqueries

- SQL provides a mechanism for the nesting of subqueries.
- A **subquery** is a **select-from-where** expression that is nested within another query.





# Example Query

- Find courses offered in Fall 2009 and in Spring 2010

**section(course\_id, semester, year)**

- Find courses offered in Fall 2009 but not in Spring 2010



# Example Query

- Find courses offered in Fall 2009 and in Spring 2010

```
select distinct course_id section(course_id, semester, year)
from section
where semester = 'Fall' and year = 2009 and semester = 'Spring'
and year = 2010;
```

- Find courses offered in Fall 2009 but not in Spring 2010



# Example Query

- Find courses offered in Fall 2009 and in Spring 2010

```
select distinct course_id section(course_id, semester, year)
from section
where semester = 'Fall' and year= 2009 and semester = 'Spring'
and year= 2010;
```

- Find courses offered in Fall 2009 but not in Spring 2010



# Example Query

- Find courses offered in Fall 2009 and in Spring 2010

```
select distinct course_id section(course_id, semester, year)
from section
where semester = 'Fall' and year= 2009 and semester = 'Spring'
and year= 2010;
```

```
select distinct course_id
from section
where semester = 'Fall' and year= 2009 and
 course_id in (select course_id
 from section
 where semester = 'Spring' and year= 2010);
```

- Find courses offered in Fall 2009 but not in Spring 2010



# Example Query

- Find courses offered in Fall 2009 and in Spring 2010

```
select distinct course_id section(course_id, semester, year)
from section
where semester = 'Fall' and year= 2009 and semester = 'Spring'
and year= 2010;
```

```
select distinct course_id
from section
where semester = 'Fall' and year= 2009 and
 course_id in (select course_id
 from section
 where semester = 'Spring' and year= 2010);
```

- Find courses offered in Fall 2009 but not in Spring 2010

```
select distinct course_id
from section
where semester = 'Fall' and year= 2009 and
 course_id not in (select course_id
 from section
 where semester = 'Spring' and year= 2010);
```



# Set Comparison

- Find names of instructors with salary greater than that of some (at least one) instructor in the Biology department.
- Same query using  $>$  **some** clause



# Set Comparison

- Find names of instructors with salary greater than that of some (at least one) instructor in the Biology department.

```
select distinct T.name
from instructor as T, instructor as S
where T.salary > S.salary and S.dept name = 'Biology';
```

- Same query using > **some** clause



# Set Comparison

- Find names of instructors with salary greater than that of some (at least one) instructor in the Biology department.

```
select distinct T.name
from instructor as T, instructor as S
where T.salary > S.salary and S.dept name = 'Biology';
```

- Same query using > **some** clause

```
select name
from instructor
where salary > some (select salary
from instructor
where dept name = 'Biology');
```





# Example Query

- Find the names of all instructors whose salary is greater than the salary of all instructors in the Biology department.



# Example Query

- Find the names of all instructors whose salary is greater than the salary of all instructors in the Biology department.

```
select name
from instructor
where salary > all (select salary
 from instructor
 where dept name = 'Biology');
```



# Test for Empty Relations

- The **exists** construct returns the value **true** if the argument subquery is nonempty.
- **exists**  $r \Leftrightarrow r \neq \emptyset$
- **not exists**  $r \Leftrightarrow r = \emptyset$



# Correlation Variables

- Yet another way of specifying the query “Find all courses taught in both the Fall 2009 semester and in the Spring 2010 semester”

```
select course_id
from section as S
where semester = 'Fall' and year = 2009 and
 exists (select *
 from section as T
 where semester = 'Spring' and year = 2010
 and S.course_id = T.course_id);
```

- **Correlated subquery**
- **Correlation name** or **correlation variable**



# Not Exists

- Find all students who have taken all courses offered in the Biology department.

- Note that  $X - Y = \emptyset \Leftrightarrow X \subseteq Y$

- *Note:* Cannot write this query using = **all** and its variants



# Not Exists

- Find all students who have taken all courses offered in the Biology department.

```
select distinct S.ID, S.name
from student as S
where not exists ((select course_id
 from course
 where dept_name = 'Biology')
except
 (select T.course_id
 from takes as T
 where S.ID = T.ID));
```

- Note that  $X - Y = \emptyset \Leftrightarrow X \subseteq Y$
- *Note:* Cannot write this query using = **all** and its variants



# Derived Relations

- SQL allows a subquery expression to be used in the **from** clause
- Find the average instructors' salaries of those departments where the average salary is greater than \$42,000."

```
select dept_name, avg_salary
from (select dept_name, avg (salary) as avg_salary
 from instructor
 group by dept_name)
where avg_salary > 42000;
```

- Note that we do not need to use the **having** clause



# With Clause

- The **with** clause provides a way of defining a temporary view whose definition is available only to the query in which the **with** clause occurs.
- Find all departments with the maximum budget

```
with max_budget (value) as
 (select max(budget)
 from department)
select budget
from department, max_budget
where department.budget = max_budget.value;
```





# Chapter 3: Introduction to SQL

- Overview of the SQL Query Language
- Data Definition
  - domain types / create table with integrity constraints
  - drop table / alter table
- Basic Query Structure
  - select / from / where / natural join
- Additional Basic Operations
  - rename / string operations / ordering / predicates
- Set Operations
  - union / intersect / except
- Null Values
- Aggregate Functions
  - group by / having
- Modification of the Database
  - deletion / insertion / update
- Nested Subqueries
  - all / some / exists / not exists / derived relations / with clause