

Intro to DB

# CHAPTER 4

## INTERMEDIATE SQL

# Chapter 4: Intermediate SQL

- Join Expressions
- Views
- Transactions
- Integrity Constraints
- SQL Data Types and Schemas
- Authorization

# Joined Relations

- **Join operation**
  - takes two relations and return as a result another relation.
  - a Cartesian product which requires that tuples in the two relations match (under some condition).
  - also specifies the attributes that are present in the result of the join
- The join operations are typically used as subquery expressions in the **from** clause

# Join Operation

- Join matches tuples with the same values specified in the on condition.


```
select *
from instructor join teaches on (instructor.ID=teaches.ID);
```

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

## Join Operation (cont.)

- Let  $r_1(A_1, \dots, A_k), r_2(B_1, \dots, B_n)$   
**select** \*  
**from**  $r_1$  **join**  $r_2$  **on** *condition*;

*is equivalent to*



- You can use expressions (that evaluate to relations) in **from** clauses  
**from**  $E_1, \dots, E_n$

**select distinct** *name, title*  
**from** *instructor* **join** *teaches* **on** (*instructor.ID=teaches.ID*), *course*  
**where** *teaches.course\_id=course.course\_id*

# Joined Relations

- take two relations and return as a result another relation.
- typically used as subquery expressions in the **from** clause
- **Join condition** – defines which tuples in the two relations match, and what attributes are present in the result of the join.
- **Join type** – defines how tuples in each relation that do not match any tuple in the other relation (based on the join condition) are treated.

<i>Join types</i>	<i>Join Conditions</i>
<b>inner join</b> <b>left outer join</b> <b>right outer join</b> <b>full outer join</b>	<b>natural</b> <b>on</b> <predicate> <b>using</b> ( $A_1, A_1, \dots, A_n$ )

# Inner join

- Inner join = join (default)

**select \* from *instructor* join *teaches* on (*instructor.ID=teaches.ID*);**

**= select \* from *instructor* inner join *teaches* on (*instructor.ID=teaches.ID*);**

- ON clause

- an explicit join clause (acts like a **where** clause)

- USING clause

- specifies which columns to test for *equality*

- 

- Above example is equivalent to

**select \* from *instructor* inner join *teaches* using ( *ID* );**

# Natural Join


- (Section 3.3.3)
- *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*;

- Let  $r_1(A_1, \dots, A_k, C_1, \dots, C_m)$ ,  $r_2(B_1, \dots, B_n, C_1, \dots, C_m)$   
**select** \*  
**from**  $r_1$  **natural join**  $r_2$

*is equivalent to*

**select**  $A_1, \dots, A_k, r_1.C_1, \dots, r_1.C_m, B_1, \dots, B_n$






# Natural Join (cont.)

- **Danger in natural join:** beware of unrelated attributes with same name which get equated incorrectly
- List the names of instructors along with the titles of courses that they teach
  - Incorrect version (makes `course.dept_name = instructor.dept_name`)
    - **select distinct** *name, title*  
**from** *instructor* **natural join** *teaches* **natural join** *course*;
  - Correct version
    - **select distinct** *name, title*  
**from** *instructor* **natural join** *teaches, course*  
**where** *teaches.course\_id = course.course\_id*;
  - Another correct version
    - **select distinct** *name, title*  
**from** (*instructor* **natural join** *teaches*)  
**join** *course* **using** (*course\_id*);

# Natural Join (cont.)

- Find the IDs of all students who were taught by an instructor named Einstein in building 301 (remove duplicates in the result).

```
select    distinct takes.ID
from      takes, section, teaches, instructor
where      takes.course_id=section.course_id and takes.sec_id=section.sec_id and
             takes.semester=section.semester and takes.year=section.year and
             section.course_id=teaches.course_id and section.sec_id=teaches.sec_id and
             section.year=teaches.year and section.year=teaches.year and
             teaches.ID=instructor.ID and
             instructor.name='Einstein' and building='301'
```



# Join operation – Example

*course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

*prereq*

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

*course* **natural join** *prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prere_id</i>	<i>course_id</i>
BIO-301	Genetics	Biology	4	BIO-101	BIO-301
CS-190	Game Design	Comp. Sci.	4	CS-101	CS-190

Note: prereq information missing for CS-315 and  
course information missing for CS-437.

# Outer Join

- An extension of the join operation that avoids loss of information.
- Computes the join and then adds tuples from one relation that does not match tuples in the other relation to the result of the join.
- Uses *null* values.

# Left Outer Join

*course* **natural left outer join** *prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prere_id</i>
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-315	Robotics	Comp. Sci.	3	<i>null</i>

= *course* **left outer join** *prereq* **on** *course.course\_id* = *prereq.course\_id*

= *course* **left outer join** *prereq* **using** (*course\_id*)

*course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

*prereq*

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

# Right Outer Join

*course* **natural right outer join** *prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prere_id</i>
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-347	<i>null</i>	<i>null</i>	<i>null</i>	CS-101

*course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

*prereq*

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

# Full Outer Join

*course* **natural full outer join** *prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prere_id</i>
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-315	Robotics	Comp. Sci.	3	<i>null</i>
CS-347	<i>null</i>	<i>null</i>	<i>null</i>	CS-101

*course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

*prereq*

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

# Joined Relations – Examples

*course* **inner join** *prereq* **on**

*course.course\_id = prereq.course\_id*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prere_id</i>	<i>course_id</i>
BIO-301	Genetics	Biology	4	BIO-101	BIO-301
CS-190	Game Design	Comp. Sci.	4	CS-101	CS-190

*course* **left outer join** *prereq* **on**

*course.course\_id = prereq.course\_id*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prere_id</i>	<i>course_id</i>
BIO-301	Genetics	Biology	4	BIO-101	BIO-301
CS-190	Game Design	Comp. Sci.	4	CS-101	CS-190
CS-315	Robotics	Comp. Sci.	3	<i>null</i>	<i>null</i>



# Integrity Constraints

- Integrity constraints guard against accidental damage to the database

- 

- Examples

- A checking account must have a balance greater than \$10,000.00
  - A salary of a bank employee must be at least \$4.00 an hour
  - A customer must have a (non-null) phone number

- Constraints on a single relation

- **not null**
  - **unique**
  - **primary key**
  - **check** ( $P$ ), where  $P$  is a predicate

- Referential integrity


# Not Null Constraint

- Declare *name* for *budget* to be **not null**

*name* **varchar(20) not null**

*budget* **numeric(12,2) not null**

- Example



# The Unique & Primary Key Constraints

- **unique** (  $A_1, A_2, \dots, A_m$  )

- **primary key** (  $A_1, A_2, \dots, A_m$  )

- $A_1, A_2, \dots, A_m$  form the primary key for the table

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

# The check clause

- **check** ( $P$ ), where  $P$  is a predicate
- Example:
  - ensure that semester is one of fall, winter, spring or summer

```
create table section (  
    course_id varchar (8),  
    sec_id varchar (8),  
    semester varchar (6),  
    year numeric (4,0),  
    building varchar (15),  
    room_number varchar (7),  
    time slot id varchar (4),  
    primary key (course_id, sec_id, semester, year),  
    check (semester in ('Fall', 'Winter', 'Spring', 'Summer'))  
)
```

# Referential Integrity

- “A value that appears in one relation for a given set of attributes also appears for a certain set of attributes in another relation.”
  - Example:
    - If “Biology” is a department name appearing in one of the tuples in the *course* relation,
    - then there exists a tuple in the *department* relation for “Biology”.
- specified as part of **create table** statement
- **Foreign Key**
  - the attributes that comprise the foreign key, and
  - the name of the relation referenced by the foreign key
  -


course_id	title	dept_name	credits
BIO-101	Intro. to Biology	Biology	4
BIO-301	Genetics	Biology	4
BIO-399	Computational Biology	Biology	3
CS-101	Intro. to Computer Science	Comp. Sci.	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3
CS-319	Image Processing	Comp. Sci.	3
CS-347	Database System Concepts	Comp. Sci.	3
EE-181	Intro to Digital Systems	Elec. Eng.	3
FINL 201	Investment Banking	Finance	3

dept_name	building	budget
Biology	Watson	90000
Comp. Sci.	Taylor	100000
Elec. Eng.	Taylor	85000
Finance	Painter	120000
History	Painter	50000
Music	Packard	80000
Physics	Watson	70000

# Cascading Actions in Referential Integrity

- **create table** *course* (  
    *course\_id* **char**(5) **primary key**,  
    *title* **varchar**(20),  
    *dept\_name* **varchar**(20) **references** *department*  
)
- **create table** *course* (  
    ...  
    *dept\_name* **varchar**(20),  
    **foreign key** (*dept\_name*) **references** *department*  
        **on delete cascade**  
        **on update cascade**,  
    ...  
)
-


# Primary & Foreign Keys in Schema Diagram



# Authorization Specification in SQL

- The **grant** statement is used to confer authorization  
    **grant** <privilege list>  
    **on** <relation name or view name>  
    **to** <user list>
- <user list> is:
  - a user-id
  - **public**, which allows all valid users the privilege granted
  - a role (explained later)


▪





# Privileges in SQL

- **select**: allows read access to relation
  - (or the ability to query using the view)
  - Example: grant users  $U_1$ ,  $U_2$ , and  $U_3$  **select** authorization on the *branch* relation:



- **insert**: the ability to insert tuples
- **update**: the ability to update using the SQL update statement
- **delete**: the ability to delete tuples.
- **all privileges**: used as a short form for all the allowable privileges
- **index**: allows creation and deletion of indices.
- **resources**: allows creation of new relations.
- **alteration**: allows addition or deletion of attributes in a relation.
- **drop**: allows deletion of relations.

# Revoking Authorization in SQL

- The **revoke** statement is used to revoke authorization.

**revoke** <privilege list>

**on** <relation name or view name>


**from** <user list>

- Example:

**revoke select on *branch* from  $U_1, U_2, U_3$**


- <privilege-list> may be **all** to revoke all privileges the revokee may hold.
- All privileges that depend on the privilege being revoked are also revoked.
- If the same privilege was granted twice to the same user by different grantees, the user may retain the privilege after the revocation.

# Authorization-Grant Graph



# Roles

- Roles are used to represent a group of users and their privileges
  - **create role** *grader*;
- Privileges can be granted to roles:
  - **grant select on** *student* **to** *grader*
  - **grant select, update on** *takes* **to** *grader*
- Roles can be granted to users, as well as to other roles
  - **grant** *grader* **to** Amit;
  - **create role** *dean*;
  - **grant** *grader* **to** *dean*;
  - **grant** *dean* **to** Kim;



**END OF CHAPTER 4**