

# Database Management Systems

Madhavan Mukund

<https://www.cmi.ac.in/~madhavan>

Sai University

Lecture 9, 13 September 2023

- Join — cartesian product combined with selection

$$\sigma_p (r_1 \times r_2)$$

$$r_1 \bowtie_p r_2$$

# Joins in SQL

- Join — cartesian product combined with selection
- Three specific types of join

- Natural join — equality on same name attributes

- Outer join

- Inner join



keep only one copy



# Joined Relations

- **Join operations** take two relations and return as a result another relation.
- A join operation is a Cartesian product which requires that tuples in the two relations match (under some condition). It 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
- Three types of joins:
  - Natural join
  - Inner join
  - Outer join



# Natural Join in SQL


- Natural join matches tuples with the same values for all common attributes, and retains only one copy of each common column.
- List the names of instructors along with the course ID of the courses that they taught
  - `select name, course_id`  
`from students, takes` ←  $\times$  *product*  
`where student.ID = takes.ID;`  $\sigma$
- Same query in SQL with “natural join” construct
  - `select name, course_id`  
`from student natural join takes;`



## Natural Join in SQL (Cont.)

- The **from** clause in can have multiple relations combined using natural join:

```
select  $A_1, A_2, \dots, A_n$   
from  $r_1$  natural join  $r_2$  natural join .. natural join  $r_n$   
where  $P$ ;
```



$$2+3+5$$
$$(2+3)+5$$
$$2+(3+5)$$

$$(r_1 \bowtie r_2) \bowtie r_3$$
$$r_1 \bowtie (r_2 \bowtie r_3)$$



# Student Relation

<i>ID</i>	<i>name</i>	<i>dept_name</i>	<i>tot_cred</i>
00128	Zhang	Comp. Sci.	102
12345	Shankar	Comp. Sci.	32
19991	Brandt	History	80
23121	Chavez	Finance	110
44553	Peltier	Physics	56
45678	Levy	Physics	46
54321	Williams	Comp. Sci.	54
55739	Sanchez	Music	38
70557	Snow	Physics	0
76543	Brown	Comp. Sci.	58
76653	Aoi	Elec. Eng.	60
98765	Bourikas	Elec. Eng.	98
98988	Tanaka	Biology	120



# Takes Relation

<i>ID</i>	<i>course_id</i>	<i>sec_id</i>	<i>semester</i>	<i>year</i>	<i>grade</i>
00128	CS-101	1	Fall	2017	A
00128	CS-347	1	Fall	2017	A-
12345	CS-101	1	Fall	2017	C
12345	CS-190	2	Spring	2017	A
12345	CS-315	1	Spring	2018	A
12345	CS-347	1	Fall	2017	A
19991	HIS-351	1	Spring	2018	B
23121	FIN-201	1	Spring	2018	C+
44553	PHY-101	1	Fall	2017	B-
45678	CS-101	1	Fall	2017	F
45678	CS-101	1	Spring	2018	B+
45678	CS-319	1	Spring	2018	B
54321	CS-101	1	Fall	2017	A-
54321	CS-190	2	Spring	2017	B+
55739	MU-199	1	Spring	2018	A-
76543	CS-101	1	Fall	2017	A
76543	CS-319	2	Spring	2018	A
76653	EE-181	1	Spring	2017	C
98765	CS-101	1	Fall	2017	C-
98765	CS-315	1	Spring	2018	B
98988	BIO-101	1	Summer	2017	A
98988	BIO-301	1	Summer	2018	<i>null</i>





## student natural join takes

Only one  
copy  
of ID

ID	name	dept_name	tot_cred	course_id	sec_id	semester	year	grade
00128	Zhang	Comp. Sci.	102	CS-101	1	Fall	2017	A
00128	Zhang	Comp. Sci.	102	CS-347	1	Fall	2017	A-
12345	Shankar	Comp. Sci.	32	CS-101	1	Fall	2017	C
12345	Shankar	Comp. Sci.	32	CS-190	2	Spring	2017	A
12345	Shankar	Comp. Sci.	32	CS-315	1	Spring	2018	A
12345	Shankar	Comp. Sci.	32	CS-347	1	Fall	2017	A
19991	Brandt	History	80	HIS-351	1	Spring	2018	B
23121	Chavez	Finance	110	FIN-201	1	Spring	2018	C+
44553	Peltier	Physics	56	PHY-101	1	Fall	2017	B-
45678	Levy	Physics	46	CS-101	1	Fall	2017	F
45678	Levy	Physics	46	CS-101	1	Spring	2018	B+
45678	Levy	Physics	46	CS-319	1	Spring	2018	B
54321	Williams	Comp. Sci.	54	CS-101	1	Fall	2017	A-
54321	Williams	Comp. Sci.	54	CS-190	2	Spring	2017	B+
55739	Sanchez	Music	38	MU-199	1	Spring	2018	A-
76543	Brown	Comp. Sci.	58	CS-101	1	Fall	2017	A
76543	Brown	Comp. Sci.	58	CS-319	2	Spring	2018	A
76653	Aoi	Elec. Eng.	60	EE-181	1	Spring	2017	C
98765	Bourikas	Elec. Eng.	98	CS-101	1	Fall	2017	C-
98765	Bourikas	Elec. Eng.	98	CS-315	1	Spring	2018	B
98988	Tanaka	Biology	120	BIO-101	1	Summer	2017	A
98988	Tanaka	Biology	120	BIO-301	1	Summer	2018	null



# Dangerous in Natural Join

- Beware of unrelated attributes with same name which get equated incorrectly
- Example -- List the names of students instructors along with the titles of courses that they have taken

- Correct version

*(student IN takes) IN P COURSE* ✓  
`select name, title  
from (student natural join takes), course  
where takes.course_id = course.course_id;`

- Incorrect version

*student IN (takes X IN P COURSE)*  
`select name, title  
from (student natural join takes) natural join course;`

- This query omits all (student name, course title) pairs where the student takes a course in a department other than the student's own department.
- The correct version (above), correctly outputs such pairs.



# 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.
- Three forms of outer join:
  - left outer join
  - right outer join
  - full outer join

$\pi_{ID}(\text{student} \bowtie \text{takes})$   
|| ?

$\pi_{ID}(\text{students})$

Blue \ Green = Students not taking a course



# Outer Join Examples

- Relation *course*

<u>course_id</u>	title	dept_name	credits
BIO-501	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

- Relation *prereq*

<u>course_id</u>	prereq_id
BIO-501	BIO-101
CS-190	CS-101
CS-347	CS-101

- Observe that

*course* information is missing for CS-437

*prereq* information is missing for CS-315



# Left Outer Join

- left* *right*
- *course* natural **left outer join** *prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_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>

- In relational algebra: *course* ⋈ *prereq*





# Right Outer Join

- *course natural right outer join prereq*

*course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_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

- In relational algebra: *course* ⋈ *prereq*

*prereq*

⋈



# Full Outer Join

- *course natural full outer join prereq*

left outer  
right outer

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_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

- In relational algebra: *course*  $\bowtie$  *prereq*



# Joined Types and Conditions

- **Join operations** take two relations and return as a result another relation.
- These additional operations are 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>
<b>inner join</b>
<b>left outer join</b>
<b>right outer join</b>
<b>full outer join</b>

<i>Join conditions</i>
<b>natural</b>
<b>on</b> <predicate>
<b>using</b> $(A_1, A_2, \dots, A_n)$



$t_1$  natural join  $t_2$  - match all columns

$t_1$  natural join  $t_2$  using  $(c_1, c_2)$

- restricts match to  
 $c_1, c_2$

$$\sigma_p(r_1 \times r_2) \rightarrow r_1 \bowtie_p r_2$$

$t_1$  join  $t_2$  on condition

from  $t_1, t_2$   
Where condition



## Joined Relations – Examples

- *course natural right outer join prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_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 full outer join prereq using (course\_id)*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_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



## Joined Relations – Examples

- *course* **inner join** *prereq* on  
*course.course\_id = prereq.course\_id*

*optional*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_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

- What is the difference between the above, and a natural join?
- *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>prereq_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>



## Joined Relations – Examples

- *course natural right outer join prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_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 full outer join prereq using (course\_id)*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_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

# Views in SQL

- Views are virtual tables

# Views in SQL

- **Views** are virtual tables
- Hide sensitive information from some users — hide salary

```
select ID, name, dept_name  
from instructor
```

# Views in SQL

- **Views** are virtual tables
- Hide sensitive information from some users — hide salary

```
select ID, name, dept_name  
from instructor
```

- Create convenient “intermediate tables”

```
select instructor.name, course.title  
from instructor, course natural join teaches
```



# View Definition and Use

- A view of instructors without their salary

```
create view faculty as  
  select ID, name, dept_name  
  from instructor
```

- Find all instructors in the Biology department

```
select name  
from faculty  
where dept_name = 'Biology'
```

- Create a view of department salary totals

```
create view departments_total_salary(dept_name, total_salary) as  
  select dept_name, sum (salary)  
  from instructor  
  group by dept_name;
```





## Views Defined Using Other Views

- One view may be used in the expression defining another view
- A view relation  $v_1$  is said to *depend directly* on a view relation  $v_2$  if  $v_2$  is used in the expression defining  $v_1$
- A view relation  $v_1$  is said to *depend on* view relation  $v_2$  if either  $v_1$  depends directly to  $v_2$  or there is a path of dependencies from  $v_1$  to  $v_2$
- A view relation  $v$  is said to be *recursive* if it depends on itself.



## Views Defined Using Other Views

- create view *physics\_fall\_2017* as  
select *course.course\_id, sec\_id, building, room\_number*  
from *course, section*  
where *course.course\_id = section.course\_id*  
and *course.dept\_name = 'Physics'*  
and *section.semester = 'Fall'*  
and *section.year = '2017'*;
- create view *physics\_fall\_2017\_watson* as  
select *course\_id, room\_number*  
from *physics\_fall\_2017*  
where *building= 'Watson'*;



# View Expansion

- Expand the view :

```
create view physics_fall_2017_watson as  
select course_id, room_number  
from physics_fall_2017  
where building = 'Watson'
```

- To:

```
create view physics_fall_2017_watson as  
select course_id, room_number  
from (select course.course_id, building, room_number  
from course, section  
where course.course_id = section.course_id  
and course.dept_name = 'Physics'  
and section.semester = 'Fall'  
and section.year = '2017')
```

where *building* = 'Watson';



## View Expansion (Cont.)

- A way to define the meaning of views defined in terms of other views.
- Let view  $v_1$  be defined by an expression  $e_1$  that may itself contain uses of view relations.
- View expansion of an expression repeats the following replacement step:
  - repeat**
    - Find any view relation  $v_i$  in  $e_1$
    - Replace the view relation  $v_i$  by the expression defining  $v_i$
  - until** no more view relations are present in  $e_1$
- As long as the view definitions are not recursive, this loop will terminate



# Materialized Views

- Certain database systems allow view relations to be physically stored.
  - Physical copy created when the view is defined.
  - Such views are called **Materialized view**:
- If relations used in the query are updated, the materialized view result becomes out of date
  - Need to **maintain** the view, by updating the view whenever the underlying relations are updated.



## Update of a View

- Add a new tuple to *faculty* view which we defined earlier

**insert into** *faculty*

**values** ('30765', 'Green', 'Music');

- This insertion must be represented by the insertion into the *instructor* relation
  - Must have a value for salary.
- Two approaches
  - Reject the insert
  - Inset the tuple

('30765', 'Green', 'Music', null)

into the *instructor* relation



## Some Updates Cannot be Translated Uniquely

- **create view** *instructor\_info* **as**  
    **select** *ID, name, building*  
    **from** *instructor, department*  
    **where** *instructor.dept\_name= department.dept\_name;*
- **insert into** *instructor\_info*  
    **values** ('69987', 'White', 'Taylor');
- Issues
  - Which department, if multiple departments in Taylor?
  - What if no department is in Taylor?



## And Some Not at All

- **create view** *history\_instructors* **as**  
**select** \*  
**from** *instructor*  
**where** *dept\_name*= 'History';
- What happens if we insert  
('25566', 'Brown', 'Biology', 100000)  
into *history\_instructors*?





# View Updates in SQL

- Most SQL implementations allow updates only on simple views
  - The **from** clause has only one database relation.
  - The **select** clause contains only attribute names of the relation, and does not have any expressions, aggregates, or **distinct** specification.
  - Any attribute not listed in the **select** clause can be set to null
  - The query does not have a **group** by or **having** clause.



## Constraints on a Single Relation

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



# Not Null Constraints

- **not null**
  - Declare *name* and *budget* to be **not null**  
*name* **varchar(20) not null**  
*budget* **numeric(12,2) not null**



# Unique Constraints

- **unique** (  $A_1, A_2, \dots, A_m$  )
  - The unique specification states that the attributes  $A_1, A_2, \dots, A_m$  form a candidate key.
  - Candidate keys are permitted to be null (in contrast to primary keys).



# The check clause

- The **check** (P) clause specifies a predicate P that must be satisfied by every tuple in a relation.
- 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

- Ensures that 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 *instructor* relation, then there exists a tuple in the *department* relation for “Biology”.
- Let A be a set of attributes. Let R and S be two relations that contain attributes A and where A is the primary key of S. A is said to be a **foreign key** of R if for any values of A appearing in R these values also appear in S.



## Referential Integrity (Cont.)


- Foreign keys can be specified as part of the SQL **create table** statement
  - foreign key** (*dept\_name*) **references** *department*
- By default, a foreign key references the primary-key attributes of the referenced table.
- SQL allows a list of attributes of the referenced relation to be specified explicitly.
  - foreign key** (*dept\_name*) **references** *department*  
(*dept\_name*)



## Cascading Actions in Referential Integrity

- When a referential-integrity constraint is violated, the normal procedure is to reject the action that caused the violation.
- An alternative, in case of delete or update is to cascade

```
create table course (  
  (...  
  dept_name varchar(20),  
  foreign key (dept_name) references department  
  on delete cascade  
  on update cascade,  
  ...)
```



- Instead of cascade we can use :
  - set null,
  - set default





## Built-in Data Types in SQL

- **date:** Dates, containing a (4 digit) year, month and date
  - Example: **date** '2005-7-27'
- **time:** Time of day, in hours, minutes and seconds.
  - Example: **time** '09:00:30'      **time** '09:00:30.75'
- **timestamp:** date plus time of day
  - Example: **timestamp** '2005-7-27 09:00:30.75'
- **interval:** period of time
  - Example: interval '1' day
  - Subtracting a date/time/timestamp value from another gives an interval value
  - Interval values can be added to date/time/timestamp values

- Many other features
  - Transactions
  - Assertions and triggers
  - ...

- Many other features
  - Transactions
  - Assertions and triggers
  - ...
- Can call SQL from other programming languages
  - Almost every language has library functions to invoke SQL
  - Transfer data between online forms and databases
  - ...

# Security — SQL injection attacks

- User input can be malicious commands to corrupt database
- Always validate data entered in a form before passing on to SQL

# Security — SQL injection attacks

- User input can be malicious commands to corrupt database
- Always validate data entered in a form before passing on to SQL

