

# RDBMS and SQL

Madhavan Mukund

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

Lecture 8, 17 October 2024

# Joins in SQL

- Join — cartesian product combined with selection

# Joins in SQL

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

# Natural join in SQL

- Match tuples with the same values for *all* common attributes
- Retain only *one copy* of each common column.

# Natural join in SQL

- Match tuples with the same values for *all* common attributes
- Retain 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
  where student.ID = takes.ID;
```

# Natural join in SQL

- Match tuples with the same values for *all* common attributes
- Retain 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
  where student.ID = takes.ID;
```

- Same query in SQL with **natural join**  

```
select name, course_id
  from student natural join takes;
```

# Natural join in SQL

- Match tuples with the same values for *all* common attributes
- Retain 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
  where student.ID = takes.ID;
```

- Same query in SQL with **natural join**  

```
select name, course_id
  from student natural join takes;
```

- Can join multiple relations at a time

```
select A1, A2, ..., Am
  from r1 natural join r2
     natural join ...
     natural join rn
  where P ;
```



# 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

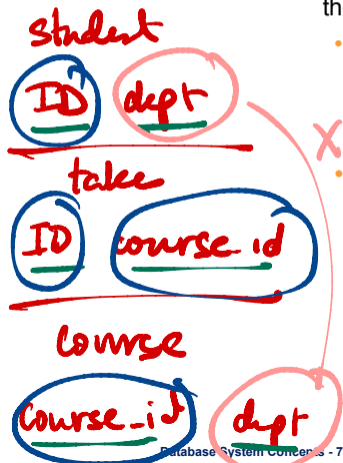
- Correct version

```
select name, title
from student natural join takes, course ✓
where takes.course_id = course.course_id;
```

- Incorrect version

```
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.



# Women's T20 WC

List all players & runs scored

## Teams

PlayerID	Name	Country

## Runs Scored

PlayerID	Runs

# Join Teams & RunsScored

select \* from

Teams natural join RunsScored

PlayerID	Name	County	Runs

Misses out  
players who  
have not played

## Outer join

If an entry is missing in the other table, add null

Select \* from

RuasSwed natural join Teams



# 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 — extra values on left
  - right outer join — extra values on right
  - full outer join — extra values on both sides



# Outer Join Examples

- Relation *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

Handwritten red checkmarks are next to the first two rows, and a red question mark is next to the third row.

- Relation *prereq*

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

The row for CS-347 is circled in red.

- Observe that

*course* information is missing for CS-437

*prereq* information is missing for CS-315





# Left Outer Join

- *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*  $\bowtie$  *prereq*



# Right Outer Join

- *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

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



# Full Outer Join

- *course* natural full 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>
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> $(A_1, A_2, \dots, A_n)$



## 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

*default*

- **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>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>

# Views in SQL

- Views are virtual tables

$r \leftarrow \dots$

# 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
```

*faculty (ID, name, dept\_name)*

*] Query to populate view*

- 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;
```

*no gap*



## 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.

Avoid cyclic dependences





## 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'*;



# 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.



## 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
  - Cascading updates to maintain referential integrity
  - Assertions and triggers
  - Transactions
  - ...

- Many other features
  - Cascading updates to maintain referential integrity
  - Assertions and triggers
  - Transactions
  - ...
- 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



- Set of attributes that one needs to keep track of

# Relational database design

- Set of attributes that one needs to keep track of
- Why not combine into a single table?



# Relational database design

<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	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

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

# Relational database design

<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	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

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

- Combine these into a single table?

# Relational database design

<i>ID</i>	<i>name</i>	<i>salary</i>	<i>dept_name</i>	<i>building</i>	<i>budget</i>
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	<del>80000</del>
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

# Relational database design

- Redundant storage

<i>ID</i>	<i>name</i>	<i>salary</i>	<i>dept_name</i>	<i>building</i>	<i>budget</i>
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

# Relational database design

- Redundant storage
- Maintaining consistency
  - Updates
  - Inserts and deletes

<i>ID</i>	<i>name</i>	<i>salary</i>	<i>dept_name</i>	<i>building</i>	<i>budget</i>
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

Instructions,  
Department


← decompose Starting Point

# Decomposition and information

- `(customer_name,regd_phone,regd_email)`

# Decomposition and information

- (customer\_name, regd\_phone, regd\_email)
- Decompose as (customer\_name, regd\_phone) and (customer\_name, regd\_email)



S	P1
S	P2

S	E1
S	E2

Srinivasan, Phone 1, Email 1  
Srinivasan, Phone 2, Email 2

# Decomposition and information

- `(customer_name,regd_phone,regd_email)`
- Decompose as `(customer_name,regd_phone)` and `(customer_name,regd_email)`
- Name is not unique — loss of **information**



# Decomposition and information

- $(customer\_name, regd\_phone, regd\_email)$
- Decompose as  $(customer\_name, regd\_phone)$  and  $(customer\_name, regd\_email)$
- Name is not unique — loss of **information**
- Recombining decomposed relation should not add tuples

S	P1	E1	
S	P1	E2	X
S	P2	E1	X
S	P2	E2	

# Decomposition and information

- $(\text{customer\_name}, \text{regd\_phone}, \text{regd\_email})$
- Decompose as  $(\text{customer\_name}, \text{regd\_phone})$  and  $(\text{customer\_name}, \text{regd\_email})$
- Name is not unique — loss of **information**
- Recombining decomposed relation should not add tuples
- **Lossless decomposition**
  - Decompose  $R$  as  $R_1$  and  $R_2$
  - Want  $R = R_1 \bowtie R_2$

If  $R_1 \cap R_2$  is empty

$R_1 \bowtie R_2$  is  $R_1 \times R_2$

# Functional dependencies

- $A_1, A_2, \dots, A_k \rightarrow B_1, B_2, \dots, B_m$ 
  - LHS attributes uniquely fix RHS attributes
  - Must hold for **every instance**  
— semantic property of attributes

<i>ID</i>	<i>name</i>	<i>salary</i>	<i>dept_name</i>	<i>building</i>	<i>budget</i>
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

# Functional dependencies

- $A_1, A_2, \dots, A_k \rightarrow B_1, B_2, \dots, B_m$ 
  - LHS attributes uniquely fix RHS attributes
  - Must hold for **every instance**  
— semantic property of attributes
- Need not correspond to superkeys
  - `dept_name`  $\rightarrow$  `building`
  - `dept_name`  $\rightarrow$  `budget`

<i>ID</i>	<i>name</i>	<i>salary</i>	<i>dept_name</i>	<i>building</i>	<i>budget</i>
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

# Functional dependencies

- $A_1, A_2, \dots, A_k \rightarrow B_1, B_2, \dots, B_m$ 
  - LHS attributes uniquely fix RHS attributes
  - Must hold for **every instance** — semantic property of attributes
- Need not correspond to superkeys
  - $dept\_name \rightarrow building$
  - $dept\_name \rightarrow budget$
- Use to identify sources of redundancy, guide decomposition

<i>ID</i>	<i>name</i>	<i>salary</i>	<i>dept_name</i>	<i>building</i>	<i>budget</i>
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

# Lossless decomposition and functional dependencies

- Decompose  $R$  as  $R_1$  and  $R_2$

# Lossless decomposition and functional dependencies

- Decompose  $R$  as  $R_1$  and  $R_2$
- Decomposition is lossless if at least one of the following functional dependencies hold
  - $R_1 \cap R_2 \rightarrow R_1$
  - $R_1 \cap R_2 \rightarrow R_2$

Use this to  
guide  
decomposition

Name	Phone	Name	Email
S	P1	S	E1
S	P2	<del>S</del>	<del>E2</del>

S	P1	E1
S	P2	E1