

Database Management Systems

Madhavan Mukund

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

Sai University

Lecture 4, 25 August 2023

A query language — select and project

- List instructors from Physics department with salary above 90,000

σ dept_name = Physics
^ AND
salary > 90000

Instructor
V OR

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	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

A query language — select and project

- List instructors from Physics department with salary above 90,000
- List names of instructors

$\pi_{\text{name}}(\text{Instructor})$

$\bar{\pi}$ 1 or more columns


$\pi_{\text{ID, name}}(\text{Instructor})$

<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	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

A query language — select and project

- List instructors from Physics department with salary above 90,000
- List names of instructors
- List names of instructors from Physics department *with salary > 90000*

<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



A query language — join

- List details of courses offered by instructors

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	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

Teaches

ID	course_id	sec_id	semester	year
10101	CS-101	1	Fall	2017
10101	CS-315	1	Spring	2018
10101	CS-347	1	Fall	2017
12121	FIN-201	1	Spring	2018
15151	MU-199	1	Spring	2018
22222	PHY-101	1	Fall	2017
32343	HIS-351	1	Spring	2018
45565	CS-101	1	Spring	2018
45565	CS-319	1	Spring	2018
76766	BIO-101	1	Summer	2017
76766	BIO-301	1	Summer	2018
83821	CS-190	1	Spring	2017
83821	CS-190	2	Spring	2017
83821	CS-319	2	Spring	2018
98345	EE-181	1	Spring	2017

A query language — join

■ Instructor × Teaches

<i>Instructor.ID</i>	<i>name</i>	<i>dept_name</i>	<i>salary</i>	<i>teaches.ID</i>	<i>course_id</i>	<i>sec_id</i>	<i>semester</i>	<i>year</i>
10101	Srinivasan	Comp. Sci.	65000	10101	CS-101	1	Fall	2017
10101	Srinivasan	Comp. Sci.	65000	10101	CS-315	1	Spring	2018
10101	Srinivasan	Comp. Sci.	65000	10101	CS-347	1	Fall	2017
10101	Srinivasan	Comp. Sci.	65000	12121	FIN-201	1	Spring	2018
10101	Srinivasan	Comp. Sci.	65000	15151	MU-199	1	Spring	2018
10101	Srinivasan	Comp. Sci.	65000	22222	PHY-101	1	Fall	2017
...
...
12121	Wu	Finance	90000	10101	CS-101	1	Fall	2017
12121	Wu	Finance	90000	10101	CS-315	1	Spring	2018
12121	Wu	Finance	90000	10101	CS-347	1	Fall	2017
12121	Wu	Finance	90000	12121	FIN-201	1	Spring	2018
12121	Wu	Finance	90000	15151	MU-199	1	Spring	2018
12121	Wu	Finance	90000	22222	PHY-101	1	Fall	2017
...
...

A query language — join

- $\sigma_{\text{InstructorID} = \text{course_id}}(\text{Instructor} \times \text{Teaches})$

<i>Instructor.ID</i>	<i>name</i>	<i>dept_name</i>	<i>salary</i>	<i>teaches.ID</i>	<i>course_id</i>	<i>sec_id</i>	<i>semester</i>	<i>year</i>
10101	Srinivasan	Comp. Sci.	65000	10101	CS-101	1	Fall	2017
10101	Srinivasan	Comp. Sci.	65000	10101	CS-315	1	Spring	2018
10101	Srinivasan	Comp. Sci.	65000	10101	CS-347	1	Fall	2017
12121	Wu	Finance	90000	12121	FIN-201	1	Spring	2018
15151	Mozart	Music	40000	15151	MU-199	1	Spring	2018
22222	Einstein	Physics	95000	22222	PHY-101	1	Fall	2017
32343	El Said	History	60000	32343	HIS-351	1	Spring	2018
45565	Katz	Comp. Sci.	75000	45565	CS-101	1	Spring	2018
45565	Katz	Comp. Sci.	75000	45565	CS-319	1	Spring	2018
76766	Crick	Biology	72000	76766	BIO-101	1	Summer	2017
76766	Crick	Biology	72000	76766	BIO-301	1	Summer	2018
83821	Brandt	Comp. Sci.	92000	83821	CS-190	1	Spring	2017
83821	Brandt	Comp. Sci.	92000	83821	CS-190	2	Spring	2017
83821	Brandt	Comp. Sci.	92000	83821	CS-319	2	Spring	2018
98345	Kim	Elec. Eng.	80000	98345	EE-181	1	Spring	2017

A query language — join

- Instructor \bowtie InstructorID = course_id Teaches — cartesian project with select

Instructor.ID	name	dept_name	salary	teaches.ID	course_id	sec_id	semester	year
10101	Srinivasan	Comp. Sci.	65000	10101	CS-101	1	Fall	2017
10101	Srinivasan	Comp. Sci.	65000	10101	CS-315	1	Spring	2018
10101	Srinivasan	Comp. Sci.	65000	10101	CS-347	1	Fall	2017
12121	Wu	Finance	90000	12121	FIN-201	1	Spring	2018
15151	Mozart	Music	40000	15151	MU-199	1	Spring	2018
22222	Einstein	Physics	95000	22222	PHY-101	1	Fall	2017
32343	El Said	History	60000	32343	HIS-351	1	Spring	2018
45565	Katz	Comp. Sci.	75000	45565	CS-101	1	Spring	2018
45565	Katz	Comp. Sci.	75000	45565	CS-319	1	Spring	2018
76766	Crick	Biology	72000	76766	BIO-101	1	Summer	2017
76766	Crick	Biology	72000	76766	BIO-301	1	Summer	2018
83821	Brandt	Comp. Sci.	92000	83821	CS-190	1	Spring	2017
83821	Brandt	Comp. Sci.	92000	83821	CS-190	2	Spring	2017
83821	Brandt	Comp. Sci.	92000	83821	CS-319	2	Spring	2018
98345	Kim	Elec. Eng.	80000	98345	EE-181	1	Spring	2017

Select σ_{θ}

Project $\pi_{c_1, c_2, \dots, c_k}$

Cartesian Product \times

Join (Select + Cartesian Product)

$T_1 \bowtie_{\theta} T_2$

\sim theta

(arbitrary condition)

Revisit join in a bit

A query language — set operations

$R \subseteq S_1 \times S_2$ R is itself a set (of pairs (s_1, s_2))

$R_1 \subseteq S_1 \times S_2$

$R_2 \subseteq S_1 \times S_2$

→

$R_1 \cup R_2$

Set Union

$R_1 \cap R_2$

Set intersection

Instructors in Physics:
Comp. Sci

$\sigma_{dept=Physics}$ (Instructor)
 $\sigma_{dept=CompSci}$ (Instructor)

Enter Physics \cup CS

$\sigma \cup \sigma$

A query language — set operations

$\sigma_{\text{dept.name} = \text{Physics} \vee \text{dept.name} = \text{Comp.Sc}} (\text{Instructor})$

$\theta_1 \sigma_{\text{deptname} = \text{Physics}} (\text{Instructor})$

$\theta_2 \sigma_{\text{salary} > 90000} (\text{Instructor})$

$\sigma_{\theta_1 \wedge \theta_2}$

$\sigma \cdot \cap \sigma$

Third operation - Complementation

Squares $\subseteq \mathbb{N} \times \mathbb{N}$ $\{(i, i^2) \mid i \in \mathbb{N}\}$

$\{(0, 0), (1, 1), (2, 4), (3, 9), \dots\}$

NOT (Squares)

$\{(0, 1), (0, 2), \dots$

$(1, 0), (1, 2), \dots$



not(Square) does not contain (2, 3.5)

Instructor

ID	Name	Dept	Sex
101	ABC	X	M
202	DEF	Y	N

not(Instructor)?

101 ABC Y M

101 ABC X N

Unmanageable list of entries

Therefore No "not" in relational algebra

Courses offered in Fall 2017 that are not
offered in Spring 2018



So far
 \cup, \cap

$\nabla_{\text{semester}} = \text{Fall 1 year} = 2017$ (Teacher)


A query language — set operations

Set Difference

$X - Y$

$X \setminus Y$

$X \setminus Y$ All elements in X that are not in Y



Set Difference

Fall 2017 \ Spring 2018

Select

Project

Cartesian Product

Join

Union

Intersection

Set Difference

(No set complement)

Intermediate Names - Assignment

Fall 2017 ← ∇ Semester = Fall \wedge Year = 2017 (Teacher)

Spring 2018 ← . . .

Fall 2017 \ Spring 2018

Renaming

Rename a relation

Instructor \rightsquigarrow I

ρ I (Instructor)
↑
rho

Identify all teachers teaching 2 courses

for each row $r1$ in *Teachers*

for each row $r2$ in *Teachers*

$r1$.teacher.id = $r2$.teacher.id

$r1$.course.id \neq $r2$.course.id

If so, record $r1$.teacher.id

A query language — assignment, renaming

Same column name in T1 & T2

T2.column name
T1.column name } to distinguish

For my nested loop

Teaches R1 Teaches
R1.inst = R2.inst
How to distinguish?

A query language — assignment, renaming

$\rho_{T1}(\text{Teaches}) \bowtie_{T1.\text{teachid} = T2.\text{teachid}}$ $\rho_{T2}(\text{Teaches})$

\wedge
 $T1.\text{courseid} \neq$
 $T2.\text{courseid}$

Without
renaming

$\text{Teaches}.\text{teachid} = \text{Teaches}.\text{teachid}$

More generally, can also rename columns

ρ_{T_1} (Teaches)

TeachId	ConcId	Sec	Yr
---------	--------	-----	----


$\rho_{T_1}(id, c, s, y)$ (Teaches)

↓

id	c	s	y
----	---	---	---

A query language — assignment, renaming

Select
Project
Cartesian Product



Join

Union
Intersection
Set Diff

Assignment
Renaming

A query language — assignment, renaming

Very common version of Join

— Same Column Name matches across tables

Personal

ID	Name	DoB

Accounts

ID	Salary

Personal \bowtie Accounts
Personal.ID
= Accounts.ID

"NATURAL JOIN"

Personal \bowtie Accounts

Sets – ordering

<i>ID</i>	<i>name</i>	<i>dept_name</i>	<i>salary</i>
10101	<u>Srinivasan</u>	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	<u>Katz</u>	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	<u>Brandt</u>	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

Instructor

$\{A, B, C\}$

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

Instructor, unsorted

$\{C, B, A\}$

Sets – duplicates

~~{A, A, B, C}~~ X

TU dept_name (Instructor)

CompSci

Finance

Music

Physics

History

Biology

↓ Salary

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

Instructor, unsorted

Tax query People from Chennai with income > 5cr

∇ Dist = Chennai ∧ (TN Data × IT Data)
Income > 5cr ∧
TN.name = IT.name

Unnecessarily large

Instead

Much smaller table

∇ TN.name = ITName (∇ Dist = Chennai (TN Data) × ∇ Income > 5cr (IT Data))

Second form is more efficient

Declarative programming - Say what you want to compute,
not how to compute it

```
fact(n) {  
  ans = 1  
  for i = 1 to n  
    ans = ans * i  
  return (ans)  
}
```

```
fact(0) = 1  
fact(n) = n * fact(n-1)  
n > 0
```

- “Programming language” for relational algebra

Structured Query Language
IBM's Project R

- “Programming language” for relational algebra
- Querying tables — select, project, join, ...

- “Programming language” for relational algebra
- Querying tables — select, project, join, ...
- Data definition — describing relational schema

- “Programming language” for relational algebra
- Querying tables — select, project, join, ...
- Data definition — describing relational schema
- Data manipulation — populating and modifying rows in tables

