

# RDBMS and SQL

Madhavan Mukund

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

Lecture 9, 22 October 2024

# 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

- Combine these into a single table?

# 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

Create new dept

Allocate building & budget

No family! — null values

# 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
- **Lossless decomposition**
  - Decompose  $R$  as  $R_1$  and  $R_2$
  - Want  $R = R_1 \bowtie R_2$

↑  
Natural join

N P1 E1  
N P2 E2

N P1  
N P2

N E1  
N E2

~ /  
N P1 E2

# 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

ID	name	salary	dept_name	building	budget
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

Assumption about data,  
given to us


# 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

- $R_1 \cap R_2 \rightarrow R_1$

- $R_1 \cap R_2 \rightarrow R_2$

Shared columns  
 $\Rightarrow$  natural join



# 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
  - $R_1 \cap R_2 \rightarrow R_1$
  - $R_1 \cap R_2 \rightarrow R_2$
- Decompose **Instructor-Department** as **Instructor** and **Department**
  - **Instructor**  $\cap$  **Department** is **dept\_name**
  - **dept\_name** is primary key for **Department**

$R_1$

$R_2$



# 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
  - $R_1 \cap R_2 \rightarrow R_1$
  - $R_1 \cap R_2 \rightarrow R_2$
- Decompose **Instructor-Department** as **Instructor** and **Department**
  - **Instructor**  $\cap$  **Department** is **dept\_name**
  - **dept\_name** is primary key for **Department**
- In general need to compute all implied dependencies
  - From  $A \rightarrow B$  and  $B \rightarrow C$ , conclude that  $A \rightarrow C$
- **Closure** of a set of dependencies  $F$  — denoted  $F^+$

$$\begin{array}{l} A \rightarrow B \\ A \rightarrow C \\ \hline A \rightarrow B, C \end{array}$$

Instructor - Dept

Instructor

Dept + Bldg

Dept + Budget

# Computing the closure of a set of attributes

- Given  $\mathcal{A} = \{A_1, A_2, \dots, A_k\}$  and  $B$ , does  $A_1, A_2, \dots, A_k \rightarrow B$ ?

$$A_1 \dots A_k \rightarrow B_1 \dots B_m$$

Sufficient to show

$$A_1 \dots A_k \rightarrow B_i \text{ for each } B_i$$

# Computing the closure of a set of attributes

- Given  $\mathcal{A} = \{A_1, A_2, \dots, A_k\}$  and  $B$ , does  $A_1, A_2, \dots, A_k \rightarrow B$ ?
- Iterative algorithm — check if  $B$  is in closure  $\mathcal{A}^+$

Initialize  $\mathcal{A}^+$  to  $\{A_1, A_2, \dots, A_k\}$

**repeat**

**for each**  $\beta \rightarrow \gamma$  in  $F$

    if  $\beta \subseteq \mathcal{A}^+$ , add  $\gamma$  to  $\mathcal{A}^+$

**end**

**until** no change in  $\mathcal{A}^+$

Original set

$A \rightarrow B$   
 $B \rightarrow C$   
 $A \rightarrow C$

$\alpha, \beta$  sets of attributes  
 $\alpha \rightarrow \beta \in F^+$  if  $\beta \subseteq \alpha^+$

# Normal forms

- Criteria to determine if the collection of tables is “good”

# Normal forms

- Criteria to determine if the collection of tables is “good”
- **Normalization** — decompose tables till they achieve a normal form

# Normal forms

- Criteria to determine if the collection of tables is “good”
- **Normalization** — decompose tables till they achieve a normal form
- Guided by functional dependencies

# Boyce-Codd Normal Form (BCNF)

- Relational schema  $R$ , set of functional dependencies  $F$



# Boyce-Codd Normal Form (BCNF)

- Relational schema  $R$ , set of functional dependencies  $F$
- Write  $\alpha, \beta$  to represent sequences of attributes  $A_1, A_2, \dots, A_k, B_1, B_2, \dots, B_m$

# Boyce-Codd Normal Form (BCNF)

- Relational schema  $R$ , set of functional dependencies  $F$
- Write  $\alpha, \beta$  to represent sequences of attributes  $A_1, A_2, \dots, A_k, B_1, B_2, \dots, B_m$
- $R$  is in BCNF if, for every  $\alpha \rightarrow \beta \in F^+$ , one of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$


$F^+$  has  $\alpha \rightarrow \beta$  if

$\alpha$  closure contains  $\beta$

$$R = (C_1, C_2, \dots, C_n)$$

$$A = (C_1, C_3, C_5) \quad A^+$$

# Instruction



Key

# Boyce-Codd Normal Form (BCNF)

- Relational schema  $R$ , set of functional dependencies  $F$
- Write  $\alpha, \beta$  to represent sequences of attributes  $A_1, A_2, \dots, A_k, B_1, B_2, \dots, B_m$
- $R$  is in BCNF if, for every  $\alpha \rightarrow \beta \in F^+$ , one of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$
- `InstructorDepartment (ID, name, salary, dept_name, building, budget)` not in BCNF



# Boyce-Codd Normal Form (BCNF)

- Relational schema  $R$ , set of functional dependencies  $F$
- Write  $\alpha, \beta$  to represent sequences of attributes  $A_1, A_2, \dots, A_k, B_1, B_2, \dots, B_m$
- $R$  is in BCNF if, for every  $\alpha \rightarrow \beta \in F^+$ , one of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$
- `InstructorDepartment(ID, name, salary, dept_name, building, budget)` not in BCNF
- `Instructor(ID, name, dept_name, salary)` and `Department(dept_name, building, budget)` are in BCNF

# Achieving BCNF

- $\alpha \rightarrow \beta \in F^+$  is a BCNF violation for  $R$  if neither of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$

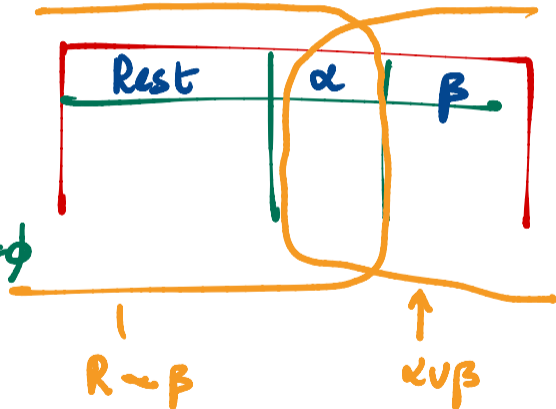
$\alpha$  is not a superkey for  $R$

# Achieving BCNF

- $\alpha \rightarrow \beta \in F^+$  is a BCNF violation for  $R$  if neither of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$
- To fix this, decompose  $R$  as
  - $\alpha \cup \beta$
  - $R \setminus (\beta \setminus \alpha)$

$$\alpha = A_1, A_2$$
$$\beta = B_1, B_2, A_1$$

$$R \setminus \beta \quad \text{if} \quad \alpha \cap \beta = \emptyset$$



# Achieving BCNF

- $\alpha \rightarrow \beta \in F^+$  is a BCNF violation for  $R$  if neither of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$
- To fix this, decompose  $R$  as
  - $\alpha \cup \beta$
  - $R \setminus (\beta \setminus \alpha)$
- Example:  $\text{dept\_name} \rightarrow \text{building, budget}$  is a BCNF violation for  $\text{InstructorDepartment}(\text{ID}, \text{name}, \text{salary}, \text{dept\_name}, \text{building}, \text{budget})$



# Achieving BCNF

- $\alpha \rightarrow \beta \in F^+$  is a BCNF violation for  $R$  if neither of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$
- To fix this, decompose  $R$  as
  - $\alpha \cup \beta$
  - $R \setminus (\beta \setminus \alpha)$
- Example:  $\text{dept\_name} \rightarrow \text{building, budget}$  is a BCNF violation for  $\text{InstructorDepartment}(\text{ID}, \text{name}, \text{salary}, \text{dept\_name}, \text{building}, \text{budget})$
- Decompose as
  - $\text{Department}(\text{dept\_name}, \text{building}, \text{budget})$
  - $\text{Instructor}(\text{ID}, \text{name}, \text{dept\_name}, \text{salary})$

$\alpha \cup \beta$   
 $R \setminus \beta$

# Dependency preservation

- `Advisor(student_id, faculty_id, dept_name)`
- Each faculty member is in only one department
- Students can be across multiple departments
- Each student has at most one advisor in each department

✓  $fac\_id \rightarrow dept$   
stud-id w/ a key

✓  $stud\_id, dept \rightarrow fac\_id$   
NOT A PROBLEM

S1 FI DI  
S2 FI DI

Split

(Stud, Fac)

(Fac, Dept)

No FD  
either way

$R \rightarrow \beta$

$\alpha \vee \beta$

# Dependency preservation

- `Advisor(student_id, faculty_id, dept_name)`
- Each faculty member is in only one department
- Students can be across multiple departments
- Each student has at most one advisor in each department
- BCNF decomposition is `(student_id, faculty_id)`, `(faculty_id, dept_name)`

Fac-id  $\rightarrow$  Dept

Studid, Dept  $\rightarrow$  Fac-id

Fac-id  $\rightarrow$  Dept

Need join to  
check second FD

# Dependency preservation

- `Advisor(student_id, faculty_id, dept_name)`
- Each faculty member is in only one department
- Students can be across multiple departments
- Each student has at most one advisor in each department
- BCNF decomposition is `(student_id, faculty_id)`, `(faculty_id, dept_name)`
- Functional dependencies
  - `faculty_id → dept_name`
  - `student_id, dept_name → faculty_id`

# Dependency preservation

- `Advisor(student_id, faculty_id, dept_name)`
- Each faculty member is in only one department
- Students can be across multiple departments
- Each student has at most one advisor in each department
- BCNF decomposition is `(student_id, faculty_id), (faculty_id, dept_name)`
- Functional dependencies
  - `faculty_id → dept_name`
  - `student_id, dept_name → faculty_id`
- Need join to check second dependency

All dependencies can  
be checked without  
joins

# Dependency preservation, formally


- Given a set of dependencies  $F$  and a decomposition of  $R$  as  $R_1, R_2, \dots, R_k$

# Dependency preservation, formally

- Given a set of dependencies  $F$  and a decomposition of  $R$  as  $R_1, R_2, \dots, R_k$
- Can **locally** check a dependency  $\alpha \rightarrow \beta$  in  $R_i$  if  $\alpha \cup \beta \subseteq R_i$

# Dependency preservation, formally

- Given a set of dependencies  $F$  and a decomposition of  $R$  as  $R_1, R_2, \dots, R_k$
- Can **locally** check a dependency  $\alpha \rightarrow \beta$  in  $R_i$  if  $\alpha \cup \beta \subseteq R_i$
- Let  $F_i$  be set of dependencies in  $F^+$  locally checkable in  $R_i$





# Dependency preservation, formally

- Given a set of dependencies  $F$  and a decomposition of  $R$  as  $R_1, R_2, \dots, R_k$
- Can **locally** check a dependency  $\alpha \rightarrow \beta$  in  $R_i$  if  $\alpha \cup \beta \subseteq R_i$
- Let  $F_i$  be set of dependencies in  $F^+$  locally checkable in  $R_i$
- Let  $G = F_1 \cup F_2 \cup \dots \cup F_k$ . Is  $G^+ = F^+$

# Dependency preservation, formally

- Given a set of dependencies  $F$  and a decomposition of  $R$  as  $R_1, R_2, \dots, R_k$
- Can **locally** check a dependency  $\alpha \rightarrow \beta$  in  $R_i$  if  $\alpha \cup \beta \subseteq R_i$
- Let  $F_i$  be set of dependencies in  $F^+$  locally checkable in  $R_i$
- Let  $G = F_1 \cup F_2 \cup \dots \cup F_k$ . Is  $G^+ = F^+$
- How do we compute  $F_i$  for each  $R_i$ ?
  - Let  $R_i$  have attributes  $A_1, A_2, \dots, A_m$
  - For each subset  $\alpha$  of  $A_1, A_2, \dots, A_m$ , compute  $\alpha^+$  with respect to  $F^+$
  - For each  $B \in \alpha^+ \cap \{A_1, A_2, \dots, A_m\}$ , add  $\alpha \rightarrow B$  to  $R_i$


# Third normal form (3NF)

■  $R$  is in 3NF if, for every  $\alpha \rightarrow \beta \in F^+$ , one of the following holds

- 1 ■  $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ ) ✓
- 2 ■  $\alpha$  is a superkey for  $R$  ✓
- 3 ■ Each attribute  $A$  in  $\beta \setminus \alpha$  is contained in some candidate key for  $R$

FacId  $\rightarrow$  Dept  
Stud, Dept  $\rightarrow$  Fac

3NF  
Dept



BCNF  
1 or 2

3NF  
1 or 2 or 3

Violation

$B \rightarrow C, D$

A, C is a key

E, D is a key

# Third normal form (3NF)

- $R$  is in 3NF if, for every  $\alpha \rightarrow \beta \in F^+$ , one of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$
  - Each attribute  $A$  in  $\beta \setminus \alpha$  is contained in some candidate key for  $R$
- BCNF is a stricter condition than 3NF

# Third normal form (3NF)

- $R$  is in 3NF if, for every  $\alpha \rightarrow \beta \in F^+$ , one of the following holds
  - $\alpha \rightarrow \beta$  is **trivial** (i.e.,  $\beta \subseteq \alpha$ )
  - $\alpha$  is a superkey for  $R$
  - Each attribute  $A$  in  $\beta \setminus \alpha$  is contained in some candidate key for  $R$
- BCNF is a stricter condition than 3NF
- Priorities
  - Lossless decomposition
  - BCNF
  - Dependency preservation

1NF - all attributes are "simple"

2NF  $\approx$  BCNF

# Beyond functional dependencies

- Suppose we collect emergency contact details for each students — phone and email
  - At least two emergency contacts of each type

# Beyond functional dependencies

- Suppose we collect emergency contact details for each students — phone and email
  - At least two emergency contacts of each type
- Consider a table `Emergency(student_id, phone, email)`
  - Two phone numbers and two emails will generate four rows



# Beyond functional dependencies

- Suppose we collect emergency contact details for each students — phone and email
  - At least two emergency contacts of each type
- Consider a table `Emergency(student_id, phone, email)`
  - Two phone numbers and two emails will generate four rows
- This redundancy cannot be explained in terms of functional dependencies

# Beyond functional dependencies

- Suppose we collect emergency contact details for each students — phone and email
  - At least two emergency contacts of each type
- Consider a table `Emergency(student_id, phone, email)`
  - Two phone numbers and two emails will generate four rows
- This redundancy cannot be explained in terms of functional dependencies
- **Multivalued dependency** — closure under swaps

# Beyond functional dependencies

- Suppose we collect emergency contact details for each students — phone and email
  - At least two emergency contacts of each type
- Consider a table `Emergency(student_id, phone, email)`
  - Two phone numbers and two emails will generate four rows
- This redundancy cannot be explained in terms of functional dependencies
- **Multivalued dependency** — closure under swaps
- 4NF

# Practical matters

- Validating functional dependencies

# Practical matters

- Validating functional dependencies
- Redundancy vs computing joins — materialized views