



CS425 – Summer 2016

Jason Arnold

Chapter 8: Relational Database Design

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan

See www.db-book.com for conditions on re-use



What is Good Design?

1) Easier: What is Bad Design?

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use



Combine Schemas?

- n Suppose we combine *instructor* and *department* into *inst_dept*
 - | (No connection to relationship set *inst_dept*)
- n Result is possible repetition of information

<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



Redundancy is Bad!

- n Update Physics Department
 - | multiple tuples to update
 - | Efficiency + potential for errors
- n Delete Physics Department
 - | update multiple tuples
 - | Efficiency + potential for errors
- n Departments without instructor or instructors without departments
 - | Need dummy department and dummy instructor
 - | Makes aggregation harder and error prone.

<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



A Combined Schema Without Repetition

- n Combining is not always bad!
- n Consider combining relations
 - | *sec_class(course_id, sec_id, building, room_number)* and
 - | *section(course_id, sec_id, semester, year)*into one relation
 - | *section(course_id, sec_id, semester, year, building, room_number)*
- n No repetition in this case

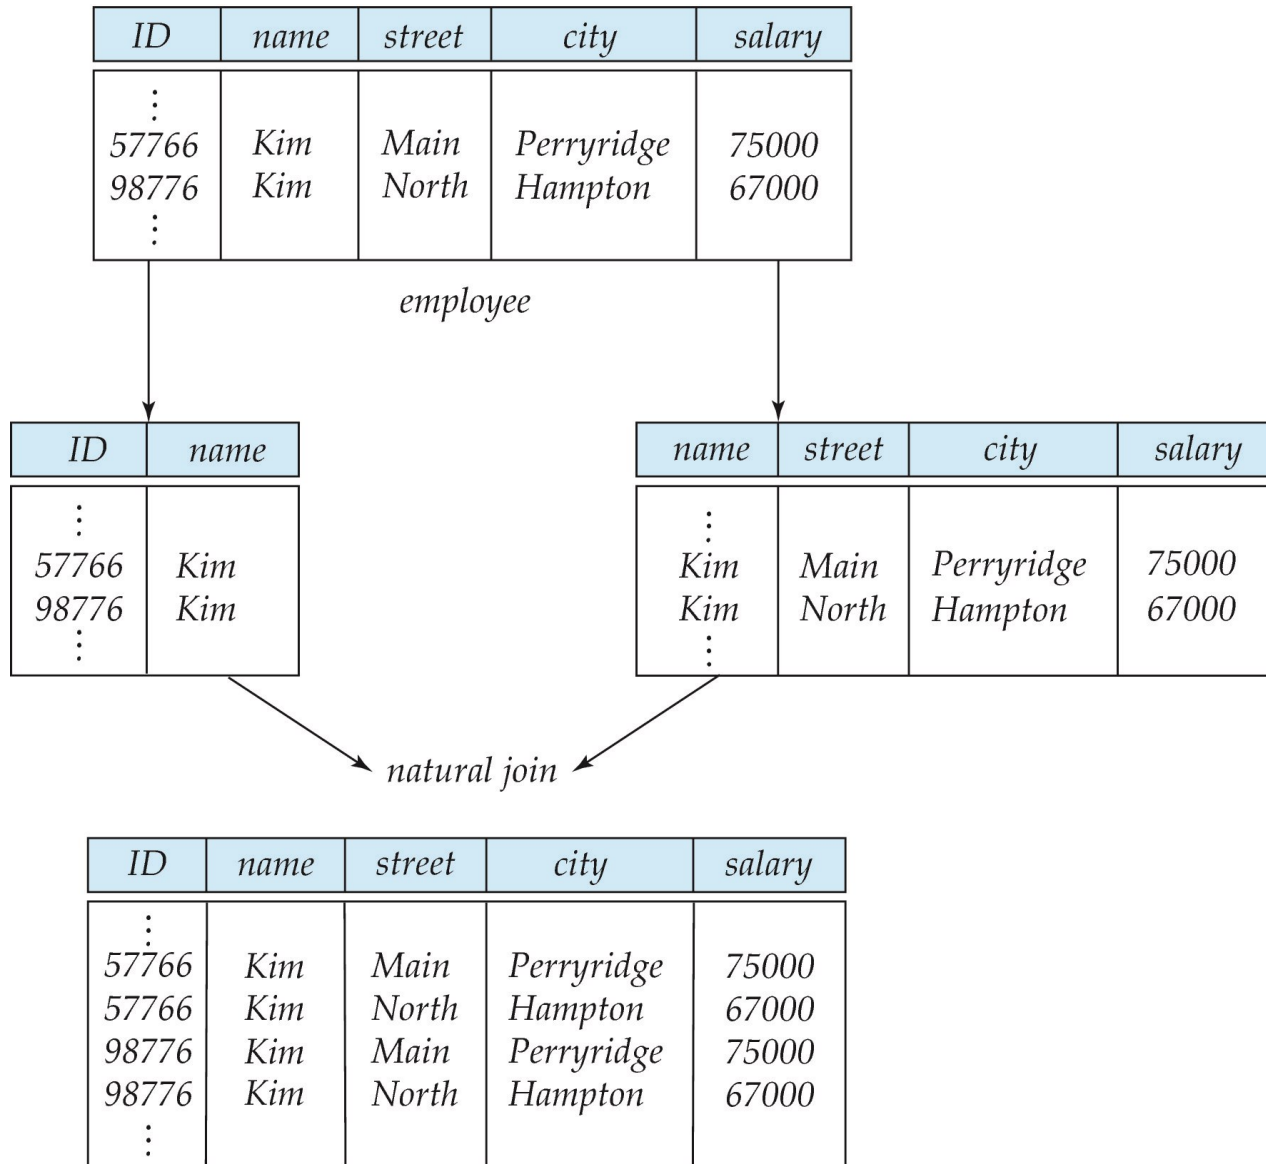


What About Smaller Schemas?

- n Suppose we had started with *inst_dept*. How would we know to split up (**decompose**) it into *instructor* and *department*?
- n Write a rule “if there were a schema (*dept_name*, *building*, *budget*), then *dept_name* would be a candidate key”
- n Denote as a **functional dependency**:
$$dept_name \rightarrow building, budget$$
- n In *inst_dept*, because *dept_name* is not a candidate key, the building and budget of a department may have to be repeated.
 - | This indicates the need to decompose *inst_dept*
- n Not all decompositions are good. Suppose we decompose *employee*(*ID*, *name*, *street*, *city*, *salary*) into
employee1 (*ID*, *name*)
employee2 (*name*, *street*, *city*, *salary*)
- n The next slide shows how we lose information -- we cannot reconstruct the original *employee* relation -- and so, this is a **lossy decomposition**.



A Lossy Decomposition





Goals of Lossless-Join Decomposition

- n Lossless-Join decomposition means splitting a table in a way so that we do not lose information
 - | That means we should be able to reconstruct the original table from the decomposed table using joins

A	B	C
α	1	A
β	2	B

r

A	B
α	1
β	2

$\Pi_{A,B}(r)$

B	C
1	A
2	B

$\Pi_{B,C}(r)$

$\Pi_A(r) \bowtie \Pi_B(r)$

A	B	C
α	1	A
β	2	B



Goal — Devise a Theory for the Following

- n Decide whether a particular relation R is in “good” form.
- n In the case that a relation R is not in “good” form, decompose it into a set of relations $\{R_1, R_2, \dots, R_n\}$ such that
 - | each relation is in good form
 - | the decomposition is a lossless-join decomposition
- n Our theory is based on:
 - | **1) Models of dependency between attribute values**
 - ▶ **functional dependencies**
 - ▶ multivalued dependencies
 - | **2) Concept of lossless decomposition**
 - | **3) Normal Forms Based On**
 - ▶ Atomicity of values
 - ▶ Avoidance of redundancy
 - ▶ Lossless decomposition



Functional Dependencies

- n Constraints on the set of legal instances for a relation schema.
- n Require that the value for a certain set of attributes determines uniquely the value for another set of attributes.
- n A functional dependency is a generalization of the notion of a *key*.
 - | *Thus, every key is a functional dependency*



Functional Dependencies (Cont.)

- n Let R be a relation schema

$$\alpha \subseteq R \text{ and } \beta \subseteq R$$

- n The **functional dependency**

$$\alpha \rightarrow \beta$$

holds on R if and only if for any legal relations $r(R)$, whenever any two tuples t_1 and t_2 of r agree on the attributes α , they also agree on the attributes β . That is,

$$t_1[\alpha] = t_2[\alpha] \Rightarrow t_1[\beta] = t_2[\beta]$$

- n Example: Consider $r(A,B)$ with the following instance of r .

1	4
1	5
3	7

- n On this instance, $A \rightarrow B$ does **NOT** hold, but $B \rightarrow A$ does hold.



Functional Dependencies (Cont.)

- n Let R be a relation schema

$$\alpha \subseteq R \text{ and } \beta \subseteq R$$

- n The **functional dependency**

$$\alpha \rightarrow \beta$$

holds on R if and only if for any legal relations $r(R)$, whenever any two tuples t_1 and t_2 of r agree on the attributes α , they also agree on the attributes β . That is,

$$t_1[\alpha] = t_2[\alpha] \Rightarrow t_1[\beta] = t_2[\beta]$$

- n Example: Consider $r(A,B)$ with the following instance of r .

1	4
1	5
3	7

A = 1 and B = 4
A = 1 and B = 5

- n On this instance, $A \rightarrow B$ does **NOT** hold, but $B \rightarrow A$ does hold.



Functional Dependencies (Cont.)

- n K is a superkey for relation schema R if and only if $K \rightarrow R$
- n K is a candidate key for R if and only if
 - | $K \rightarrow R$, and
 - | for no $\alpha \subset K$, $\alpha \rightarrow R$
- n Functional dependencies allow us to express constraints that cannot be expressed using superkeys. Consider the schema:
inst_dept (*ID*, *name*, *salary*, *dept_name*, *building*, *budget*).

We expect these functional dependencies to hold:

dept_name \rightarrow *building*

and $ID \rightarrow building$

but would not expect the following to hold:

dept_name $\rightarrow salary$



Functional Dependencies (Cont.)

- n A functional dependency is **trivial** if it is satisfied by all instances of a relation
 - | Example:
 - ▶ $ID, name \rightarrow ID$
 - ▶ $name \rightarrow name$
 - | In general, $\alpha \rightarrow \beta$ is trivial if $\beta \subseteq \alpha$



Closure of a Set of Functional Dependencies

- n Given a set F of functional dependencies, there are certain other functional dependencies that are logically implied by F .
 - | For example: If $A \rightarrow B$ and $B \rightarrow C$, then we can infer that $A \rightarrow C$
- n The set of **all** functional dependencies logically implied by F is the **closure** of F .
- n We denote the *closure* of F by F^+ .
- n F^+ is a superset of F .



Functional-Dependency Theory

- n We now consider the formal theory that tells us which functional dependencies are implied logically by a given set of functional dependencies.
- n How do we get the initial set of FDs?
 - | Semantics of the domain we are modelling
 - | Has to be provided by a human (the designer)
- n Example:
 - | Relation Citizen(SSN, FirstName, LastName, Address)
 - | We know that SSN is unique and a person has a unique SSN
 - | Thus, $SSN \rightarrow FirstName, LastName$



Closure of a Set of Functional Dependencies

- n We can find F^+ , the closure of F , by repeatedly applying **Armstrong's Axioms**:
 - | if $\beta \subseteq \alpha$, then $\alpha \rightarrow \beta$ (**reflexivity**)
 - | if $\alpha \rightarrow \beta$, then $\gamma \alpha \rightarrow \gamma \beta$ (**augmentation**)
 - | if $\alpha \rightarrow \beta$, and $\beta \rightarrow \gamma$, then $\alpha \rightarrow \gamma$ (**transitivity**)
- n These rules are
 - | **sound** (generate only functional dependencies that actually hold), and
 - | **complete** (generate all functional dependencies that hold).



Example

n $R = (A, B, C, G, H, I)$

$F = \{ A \rightarrow B$

$A \rightarrow C$

$CG \rightarrow H$

$CG \rightarrow I$

$B \rightarrow H\}$

n some members of F^+

| $A \rightarrow H$

▶ by transitivity from $A \rightarrow B$ and $B \rightarrow H$

| $AG \rightarrow I$

▶ by augmenting $A \rightarrow C$ with G , to get $AG \rightarrow CG$
and then transitivity with $CG \rightarrow I$

| $CG \rightarrow HI$

▶ by augmenting $CG \rightarrow I$ to infer $CG \rightarrow CGI$,
and augmenting of $CG \rightarrow H$ to infer $CGI \rightarrow HI$,
and then transitivity



Procedure for Computing F^+

- n To compute the closure of a set of functional dependencies F :

$F^+ = F$

repeat

for each functional dependency f in F^+

 apply reflexivity and augmentation rules on f

 add the resulting functional dependencies to F^+

for each pair of functional dependencies f_1 and f_2 in F^+

if f_1 and f_2 can be combined using transitivity

then add the resulting functional dependency to F^+

until F^+ does not change any further

NOTE: We shall see an alternative more efficient procedure for this task later



Closure of Functional Dependencies (Cont.)

- n Additional rules:
 - | If $\alpha \rightarrow \beta$ holds and $\alpha \rightarrow \gamma$ holds, then $\alpha \rightarrow \beta\gamma$ holds (**union**)
 - | If $\alpha \rightarrow \beta\gamma$ holds, then $\alpha \rightarrow \beta$ holds and $\alpha \rightarrow \gamma$ holds (**decomposition**)
 - | If $\alpha \rightarrow \beta$ holds and $\gamma\beta \rightarrow \delta$ holds, then $\alpha\gamma \rightarrow \delta$ holds (**pseudotransitivity**)

The above rules can be inferred from Armstrong's axioms.



Closure of Attribute Sets

- n Given a set of attributes α , define the **closure** of α **under** F (denoted by α^+) as the set of attributes that are functionally determined by α under F
- n Algorithm to compute α^+ , the closure of α under F

```
result :=  $\alpha$ ;  
while (changes to result) do  
  for each  $\beta \rightarrow \gamma$  in  $F$  do  
    begin  
      if  $\beta \subseteq \textit{result}$  then result := result  $\cup$   $\gamma$   
    end
```



Example of Attribute Set Closure

n $R = (A, B, C, G, H, I)$

n $F = \{A \rightarrow B$
 $A \rightarrow C$
 $CG \rightarrow H$
 $CG \rightarrow I$
 $B \rightarrow H\}$

n $(AG)^+$

1. $result = AG$

2. $result = ABCG$ ($A \rightarrow C$ and $A \rightarrow B$)

3. $result = ABCGH$ ($CG \rightarrow H$ and $CG \subseteq AGBC$)

4. $result = ABCGHI$ ($CG \rightarrow I$ and $CG \subseteq AGBCH$)

n Is AG a candidate key?

1. Is AG a super key?

1. Does $AG \rightarrow R$? == Is $(AG)^+ \supseteq R$

2. Is any subset of AG a superkey?

1. Does $A \rightarrow R$? == Is $(A)^+ \supseteq R$

2. Does $G \rightarrow R$? == Is $(G)^+ \supseteq R$



Uses of Attribute Closure

There are several uses of the attribute closure algorithm:

- n Testing for superkey:
 - | To test if α is a superkey, we compute α^+ , and check if α^+ contains all attributes of R .
- n Testing functional dependencies
 - | To check if a functional dependency $\alpha \rightarrow \beta$ holds (or, in other words, is in F^+), just check if $\beta \subseteq \alpha^+$.
 - | That is, we compute α^+ by using attribute closure, and then check if it contains β .
 - | Is a simple and cheap test, and very useful
- n Computing closure of F
 - | For each $\gamma \subseteq R$, we find the closure γ^+ , and for each $S \subseteq \gamma^+$, we output a functional dependency $\gamma \rightarrow S$.



Canonical Cover

- n Sets of functional dependencies may have redundant dependencies that can be inferred from the others
 - | For example: $A \rightarrow C$ is redundant in: $\{A \rightarrow B, B \rightarrow C, A \rightarrow C\}$
 - | Parts of a functional dependency may be redundant
 - ▶ E.g.: on RHS: $\{A \rightarrow B, B \rightarrow C, A \rightarrow CD\}$ can be simplified to
$$\{A \rightarrow B, B \rightarrow C, A \rightarrow D\}$$
 - ▶ E.g.: on LHS: $\{A \rightarrow B, B \rightarrow C, AC \rightarrow D\}$ can be simplified to
$$\{A \rightarrow B, B \rightarrow C, A \rightarrow D\}$$
- n Intuitively, a **canonical cover** of F is a “minimal” set of functional dependencies equivalent to F , having no redundant dependencies or redundant parts of dependencies



Extraneous Attributes

- n Consider a set F of functional dependencies and the functional dependency $\alpha \rightarrow \beta$ in F .
 - | Attribute A is **extraneous** in α if $A \in \alpha$ and F logically implies $(F - \{\alpha \rightarrow \beta\}) \cup \{(\alpha - A) \rightarrow \beta\}$.
 - | Attribute A is **extraneous** in β if $A \in \beta$ and the set of functional dependencies $(F - \{\alpha \rightarrow \beta\}) \cup \{\alpha \rightarrow (\beta - A)\}$ logically implies F .
- n *Note:* implication in the opposite direction is trivial in each of the cases above, since a “stronger” functional dependency always implies a weaker one
- n Example: Given $F = \{A \rightarrow C, AB \rightarrow C\}$
 - | B is extraneous in $AB \rightarrow C$ because $\{A \rightarrow C, AB \rightarrow C\}$ logically implies $A \rightarrow C$ (i.e. the result of dropping B from $AB \rightarrow C$).
- n Example: Given $F = \{A \rightarrow C, AB \rightarrow CD\}$
 - | C is extraneous in $AB \rightarrow CD$ since $AB \rightarrow C$ can be inferred even after deleting C



Lossless Join-Decomposition Dependency Preservation

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use



So Far

- n **Theory of dependencies**
- n **What is missing?**
 - | When is a decomposition loss-less
 - ▶ Lossless-join decomposition
 - ▶ Dependencies on the input are preserved
- n **What else is missing?**
 - | Define what constitutes a good relation
 - ▶ Normal forms
 - | How to check for a good relation
 - ▶ Test normal forms
 - | How to achieve a good relation
 - ▶ Translate into normal form
 - ▶ Involves decomposition



Lossless-join Decomposition

- n For the case of $R = (R_1, R_2)$, we require that for all possible relation instances r on schema R

$$r = \Pi_{R_1}(r) \bowtie \Pi_{R_2}(r)$$

- n A decomposition of R into R_1 and R_2 is lossless join if at least one of the following dependencies is in F^+ :

- | $R_1 \cap R_2 \rightarrow R_1$

- | $R_1 \cap R_2 \rightarrow R_2$

- n The above functional dependencies are a **sufficient** condition for lossless join decomposition; the dependencies are a **necessary** condition only if all constraints are functional dependencies



Example

- n $R = (A, B, C)$
 $F = \{A \rightarrow B, B \rightarrow C\}$
 - | Can be decomposed in two different ways

- n $R_1 = (A, B), R_2 = (B, C)$

- | Lossless-join decomposition:

$$R_1 \cap R_2 = \{B\} \text{ and } B \rightarrow BC$$

- | Dependency preserving

- n $R_1 = (A, B), R_2 = (A, C)$

- | Lossless-join decomposition:

$$R_1 \cap R_2 = \{A\} \text{ and } A \rightarrow AB$$

- | Not dependency preserving
(cannot check $B \rightarrow C$ without computing $R_1 \bowtie R_2$)



Dependency Preservation

- n Let F_i be the set of dependencies F^+ that include only attributes in R_i .
 - ▶ A decomposition is **dependency preserving**, if
$$(F_1 \cup F_2 \cup \dots \cup F_n)^+ = F^+$$
 - ▶ If it is not, then checking updates for violation of functional dependencies may require computing joins, which is expensive.



Example

- n $R = (A, B, C)$
 $F = \{A \rightarrow B$
 $B \rightarrow C\}$
Key = $\{A\}$
- n Decomposition $R_1 = (A, B), R_2 = (B, C)$
 - | Lossless-join decomposition
 - | Dependency preserving



Normal Forms

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use



So Far

- n **Theory of dependencies**
- n **Decompositions and ways to check whether they are “good”**
 - | Lossless
 - | Dependency preserving
- n **What is missing?**
 - | Define what constitutes a good relation
 - ▶ Normal forms
 - | How to check for a good relation
 - ▶ Test normal forms
 - | How to achieve a good relation
 - ▶ Translate into normal form
 - ▶ Involves decomposition



Goals of Normalization

- n Let R be a relation scheme with a set F of functional dependencies.
- n Decide whether a relation scheme R is in “good” form.
- n In the case that a relation scheme R is not in “good” form, decompose it into a set of relation scheme $\{R_1, R_2, \dots, R_n\}$ such that
 - | each relation scheme is in good form
 - | the decomposition is a lossless-join decomposition
 - | Preferably, the decomposition should be dependency preserving.



First Normal Form

- n A domain is **atomic** if its elements are considered to be indivisible units
 - | Examples of non-atomic domains:
 - ▶ Set of names, composite attributes
 - ▶ Identification numbers like CS101 that can be broken up into parts
- n A relational schema R is in **first normal form** if the domains of all attributes of R are atomic
- n Non-atomic values complicate storage and encourage redundant (repeated) storage of data
 - | Example: Set of accounts stored with each customer, and set of owners stored with each account
 - | We assume all relations are in first normal form
 - | (revisited in Chapter 22 of the textbook: Object Based Databases)



First Normal Form (Cont' d)

- n Atomicity is actually a property of how the elements of the domain are used.
 - | Example: Strings would normally be considered indivisible
 - | Suppose that students are given roll numbers which are strings of the form *CS0012* or *EE1127*
 - | If the first two characters are extracted to find the department, the domain of roll numbers is not atomic.
 - | Doing so is a bad idea: leads to encoding of information in application program rather than in the database.



Second Normal Form

- n A relation schema R in **1NF** is in **second normal form (2NF)** iff
 - | No non-prime attribute depends on parts of a candidate key
 - | An attribute is non-prime if it does not belong to any candidate key for R



Second Normal Form Example

- n $R(A,B,C,D)$
 - | $A,B \rightarrow C,D$
 - | $A \rightarrow C$
 - | $B \rightarrow D$
- n $\{A,B\}$ is the only candidate key
- n R is not in 2NF, because $A \rightarrow C$ where A is part of a candidate key and C is not part of a candidate key
- n Interpretation $\mathbf{R}(A,B,C,D)$ is **Advisor**(InstrSSN, StudentCWID, InstrName, StudentName)
 - | Indication that we are putting stuff together that does not belong together



Second Normal Form Interpretation

- n Why is a dependency on parts of a candidate key bad?
 - | That is why is a relation that is not in 2NF bad?
- n 1) A dependency on part of a candidate key indicates potential for redundancy
 - | **Advisor**(InstrSSN, StudentCWID, InstrName, StudentName)
 - | StudentCWID \rightarrow StudentName
 - | If a student is advised by multiple instructors we record his name several times
- n 2) A dependency on parts of a candidate key shows that some attributes are unrelated to other parts of a candidate key
 - | That means the table should be split



2NF is What We Want?

- n **Instructor**(Name, Salary, DepName, DepBudget) = I(A,B,C,D)
 - | $A \rightarrow B, C, D$
 - | $C \rightarrow D$
- n {Name} is the only candidate key
- n I is in 2NF
- n However, as we have seen before I still has update redundancy that can cause update anomalies
 - | We repeat the budget of a department if there is more than one instructor working for that department



Third Normal Form

- n A relation schema R is in **third normal form (3NF)** if for all:

$$\alpha \rightarrow \beta \text{ in } F^+$$

at least one of the following holds:

- | $\alpha \rightarrow \beta$ is trivial (i.e., $\beta \in \alpha$)
- | α is a superkey for R
- | Each attribute A in $\beta - \alpha$ is contained in a candidate key for R .
(**NOTE:** each attribute may be in a different candidate key)

Alternatively,

- | Every attribute depends directly on a candidate key, i.e., for every attribute A there is a dependency $X \rightarrow A$, but no dependency $Y \rightarrow A$ where Y is not a candidate key



3NF Example

- n **Instructor**(Name, Salary, DepName, DepBudget) = I(A,B,C,D)
 - | $A \rightarrow B, C, D$
 - | $C \rightarrow D$
- n {Name} is the only candidate key
- n I is in 2NF
- n I is not in 3NF



Testing for 3NF

- n Optimization: Need to check only FDs in F , need not check all FDs in F^+ .
- n Use attribute closure to check for each dependency $\alpha \rightarrow \beta$, if α is a superkey.
- n If α is not a superkey, we have to verify if each attribute in β is contained in a candidate key of R
 - | this test is rather more expensive, since it involve finding candidate keys
 - | testing for 3NF has been shown to be NP-hard
 - | Interestingly, decomposition into third normal form (described shortly) can be done in polynomial time



2NF/3NF Decomposition: An Example

n Relation schema:

$cust_banker_branch = (\underline{customer_id}, \underline{employee_id}, branch_name, type)$

n The functional dependencies for this relation schema are:

1. $customer_id, employee_id \rightarrow branch_name, type$
2. $employee_id \rightarrow branch_name$
3. $customer_id, branch_name \rightarrow employee_id$

n We first compute a canonical cover

| $branch_name$ is extraneous in the r.h.s. of the 1st dependency

| No other attribute is extraneous, so we get $F_C =$

$customer_id, employee_id \rightarrow type$

$employee_id \rightarrow branch_name$

$customer_id, branch_name \rightarrow employee_id$



Another 3NF Example

n Relation *dept_advisor*:

| ***dept_advisor*** (*s_ID*, *i_ID*, *dept_name*)

$F = \{s_ID, dept_name \rightarrow i_ID,$

$i_ID \rightarrow dept_name\}$

| Two candidate keys: *s_ID*, *dept_name*, and *i_ID*, *s_ID*

| *R* is in 3NF



Redundancy in 3NF

- n There is some redundancy in this schema **dept_advisor** (*s_ID*, *i_ID*, *dept_name*)
- n Example of problems due to redundancy in 3NF

| $R = (J, K, L)$
 $F = \{JK \rightarrow L, L \rightarrow K\}$

<i>J</i>	<i>L</i>	<i>K</i>
<i>j</i> ₁	<i>l</i> ₁	<i>k</i> ₁
<i>j</i> ₂	<i>l</i> ₁	<i>k</i> ₁
<i>j</i> ₃	<i>l</i> ₁	<i>k</i> ₁
<i>null</i>	<i>l</i> ₂	<i>k</i> ₂

- n repetition of information (e.g., the relationship *l*₁, *k*₁)
 - (*i_ID*, *dept_name*)
- n need to use null values (e.g., to represent the relationship *l*₂, *k*₂ where there is no corresponding value for *J*).
 - (*i_ID*, *dept_name*) if there is no separate relation mapping instructors to departments



Boyce-Codd Normal Form (BCNF)

n A relation schema R is in **BCNF** if for all:

$$\alpha \rightarrow \beta \text{ in } F^+$$

at least one of the following holds:

| $\alpha \rightarrow \beta$ is trivial (i.e., $\beta \in \alpha$)

| α is a superkey for R

~~| Each attribute A in $\beta - \alpha$ is contained in a candidate key for R .~~

~~— (**NOTE:** each attribute may be in a different candidate key)~~



BCNF and Dependency Preservation

- n If a relation is in BCNF it is in 3NF
- n Constraints, including functional dependencies, are costly to check in practice unless they pertain to only one relation
- n Because it is **not always** possible to achieve **both BCNF and dependency preservation**, we usually consider normally *third normal form*.



Comparison of BCNF and 3NF

- n It is always possible to decompose a relation into a set of relations that are in 3NF such that:
 - | the decomposition is lossless
 - | the dependencies are preserved
- n It is always possible to decompose a relation into a set of relations that are in BCNF such that:
 - | the decomposition is lossless
 - | it may not be possible to preserve dependencies.



Summary Normal Forms

n BCNF \rightarrow 3NF \rightarrow 2NF \rightarrow 1NF

n **1NF**

| atomic attributes

n **2NF**

| no non-trivial dependencies of non-prime attributes on parts of the key

n **3NF**

| no transitive non-trivial dependencies on the key

n **4NF and 5NF**



Final Thoughts on Design Process

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use



Overall Database Design Process

- n We have assumed schema R is given
 - | R could have been generated when converting an ER diagram to a set of tables.
 - | R could have been a single relation containing *all* attributes that are of interest (called **universal relation**).
 - | Normalization breaks R into smaller relations.
 - | R could have been the result of some ad hoc design of relations, which we then test/convert to normal form.



ER Model and Normalization

- n When an ER diagram is carefully designed, identifying all entities correctly, the tables generated from the ER diagram should not need further normalization.
- n However, in a real (imperfect) design, there can be functional dependencies from non-key attributes of an entity to other attributes of the entity
 - | Example: an *employee* entity with attributes *department_name* and *building*, and a functional dependency $department_name \rightarrow building$
 - | Good design would have made department an entity
- n Functional dependencies from non-key attributes of a relationship set possible, but rare --- most relationships are binary



Denormalization for Performance

- n May want to use non-normalized schema for performance
- n For example, displaying *prereqs* along with *course_id*, and *title* requires join of *course* with *prereq*
- n Alternative 1: Use denormalized relation containing attributes of *course* as well as *prereq* with all above attributes
 - | faster lookup
 - | extra space and extra execution time for updates
 - | extra coding work for programmer and possibility of error in extra code
- n Alternative 2: use a materialized view defined as
course prereq
 - | Benefits and drawbacks same as above, except no extra coding work for programmer and avoids possible errors



Other Design Issues

- n Some aspects of database design are not caught by normalization
- n Examples of bad database design, to be avoided:

Instead of *earnings* (*company_id*, *year*, *amount*), use

- | *earnings_2004*, *earnings_2005*, *earnings_2006*, etc., all on the schema (*company_id*, *earnings*).
 - ▶ Above are in BCNF, but make querying across years difficult and needs new table each year
- | *company_year* (*company_id*, *earnings_2004*, *earnings_2005*, *earnings_2006*)
 - ▶ Also in BCNF, but also makes querying across years difficult and requires new attribute each year.
 - ▶ Is an example of a **crosstab**, where values for one attribute become column names
 - ▶ Used in spreadsheets, and in data analysis tools



Recap

- n Functional and Multi-valued Dependencies
 - | Axioms
 - | Closure
 - | Minimal Cover
 - | Attribute Closure
- n Redundancy and lossless decomposition
- n Normal-Forms
 - | 1NF, 2NF, 3NF
 - | BCNF
 - | 4NF, 5NF



End of Chapter

modified from:

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use