

Design Theory for Relational Databases

Functional Dependencies

Decompositions

Normal Forms

Functional Dependencies

- ◆ $X \rightarrow Y$ is an assertion about a relation R that whenever two tuples of R agree on all the attributes of X , then they must also agree on all attributes in set Y .
 - ◆ Say " $X \rightarrow Y$ holds in R ."
 - ◆ **Convention:** ..., X, Y, Z represent sets of attributes; A, B, C, \dots represent single attributes.
 - ◆ **Convention:** no set formers in sets of attributes, just ABC , rather than $\{A, B, C\}$.

Splitting Right Sides of FD's

- ◆ $X \rightarrow A_1 A_2 \dots A_n$ holds for R exactly when each of $X \rightarrow A_1$, $X \rightarrow A_2, \dots, X \rightarrow A_n$ hold for R .
- ◆ Example: $A \rightarrow BC$ is equivalent to $A \rightarrow B$ and $A \rightarrow C$.
- ◆ There is no splitting rule for left sides.
- ◆ We'll generally express FD's with singleton right sides.

Example: FD's

Drinkers(name, addr, lemonsLiked,
manf, favLemonade)

◆ Reasonable FD's to assert:

1. name -> addr favLemonade

- ◆ Note this FD is the same as name -> addr and name -> favLemonade.

2. lemonsLiked -> manf

Example: Possible Data

name	addr	lemonadesLiked	manf	favLemonade
Janeway	Voyager	Bud	A.B.	WickedAle
Janeway	Voyager	WickedAle	Pete's	WickedAle
Spock	Enterprise	Bud	A.B.	Bud

Because name -> addr

Because name -> favLemonade

Because lemonadesLiked -> manf

Keys of Relations

- ◆ K is a *superkey* for relation R if K functionally determines all of R .
- ◆ K is a *key* for R if K is a superkey, but no proper subset of K is a superkey.

Example: Superkey

Drinkers(name, addr, lemonadesLiked,
manf, favLemonade)

- ◆ {name, lemonadesLiked} is a superkey because together these attributes determine all the other attributes.
 - ◆ name -> addr favLemonades
 - ◆ lemonadesLiked -> manf

Example: Key

- ◆ $\{\text{name}, \text{lemonadesLiked}\}$ is a **key** because neither $\{\text{name}\}$ nor $\{\text{lemonadesLiked}\}$ is a superkey.
 - ◆ name doesn't \rightarrow manf ; lemonadesLiked doesn't \rightarrow addr .
- ◆ There are no other keys, but lots of superkeys.
 - ◆ Any superset of $\{\text{name}, \text{lemonadesLiked}\}$.

Where Do Keys Come From?

1. Just assert a key K .
 - ◆ The only FD's are $K \rightarrow A$ for all attributes A .
2. Assert FD's and deduce the keys by systematic exploration.

More FD's From "Physics"

- ◆ **Example:** "no two courses can meet in the same room at the same time" tells us: **hour room -> course.**

Inferring FD's

- ◆ We are given FD's $X_1 \rightarrow A_1, X_2 \rightarrow A_2, \dots, X_n \rightarrow A_n$, and we want to know whether an FD $Y \rightarrow B$ must hold in any relation that satisfies the given FD's.
 - ◆ Example: If $A \rightarrow B$ and $B \rightarrow C$ hold, surely $A \rightarrow C$ holds, even if we don't say so.
- ◆ Important for design of good relation schemas.

Inference Test

- ◆ To test if $Y \rightarrow B$, start by assuming two tuples agree in all attributes of Y .

$\leftarrow Y \rightarrow$

0000000. . . 0

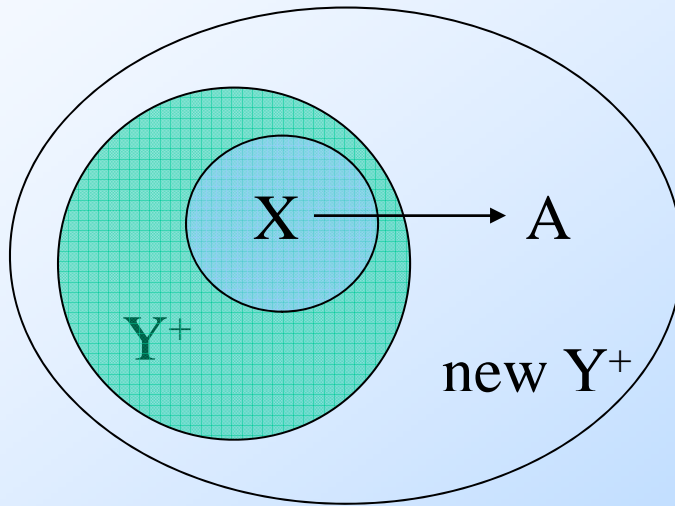
00000?? . . . ?

Inference Test – (2)

- ◆ Use the given FD's to infer that these tuples must also agree in certain other attributes.
 - ◆ If B is one of these attributes, then $Y \rightarrow B$ is true.
 - ◆ Otherwise, the two tuples, with any forced equalities, form a two-tuple relation that proves $Y \rightarrow B$ does not follow from the given FD's.

Closure Test

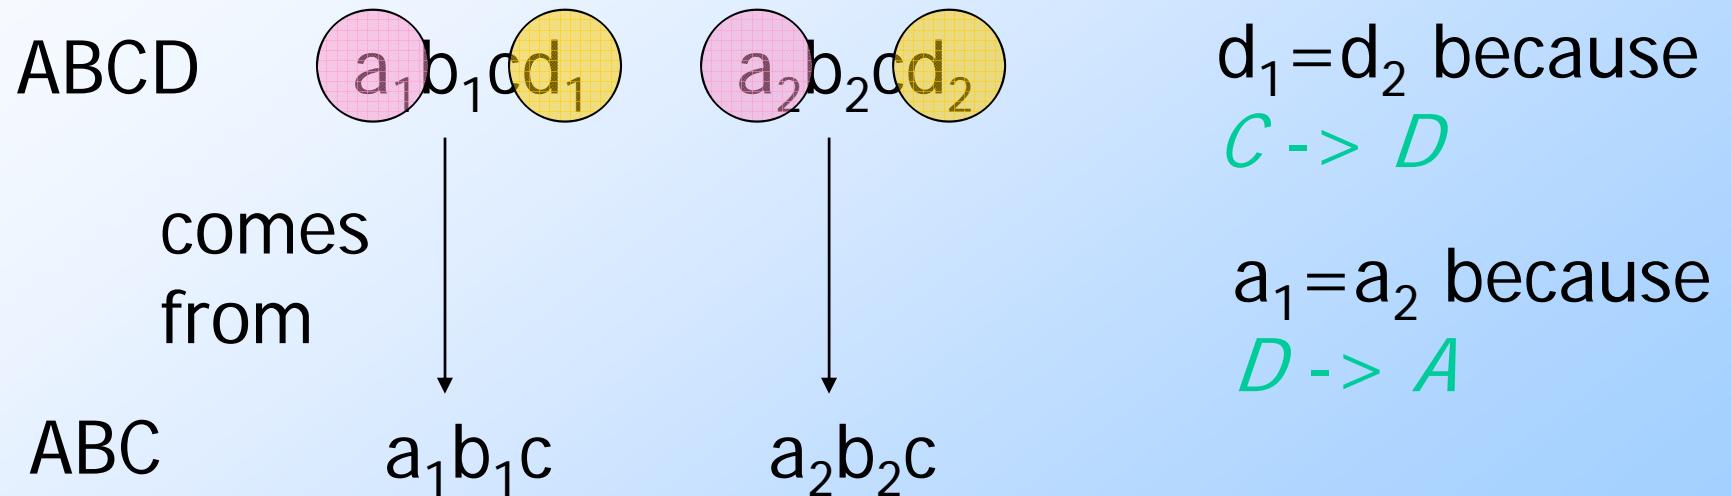
- ◆ An easier way to test is to compute the *closure* of Y , denoted Y^+ .
- ◆ **Basis:** $Y^+ = Y$.
- ◆ **Induction:** Look for an FD's left side X that is a subset of the current Y^+ . If the FD is $X \rightarrow A$, add A to Y^+ .



Finding All Implied FD's

- ◆ **Motivation:** “normalization,” the process where we break a relation schema into two or more schemas.
- ◆ Example: $ABCD$ with FD's $AB \rightarrow C$, $C \rightarrow D$, and $D \rightarrow A$.
 - ◆ Decompose into ABC , AD . What FD's hold in ABC ?
 - ◆ Not only $AB \rightarrow C$, but also $C \rightarrow A$!

Why?



Thus, tuples in the projection with equal C's have equal A's;

$C \rightarrow A$.

Basic Idea

1. Start with given FD's and find all *nontrivial* FD's that follow from the given FD's.
 - ◆ Nontrivial = right side not contained in the left.
2. Restrict to those FD's that involve only attributes of the projected schema.

Simple, Exponential Algorithm

1. For each set of attributes X , compute X^+ .
2. Add $X \rightarrow A$ for all A in $X^+ - X$.
3. However, drop $XY \rightarrow A$ whenever we discover $X \rightarrow A$.
 - ◆ Because $XY \rightarrow A$ follows from $X \rightarrow A$ in any projection.
4. Finally, use only FD's involving projected attributes.

A Few Tricks

- ◆ No need to compute the closure of the empty set or of the set of all attributes.
- ◆ If we find $X^+ = \text{all attributes}$, so is the closure of any superset of X .

Example: Projecting FD's

◆ ABC with FD's $A \rightarrow B$ and $B \rightarrow C$.
Project onto AC .

- ◆ $A^+ = ABC$; yields $A \rightarrow B, A \rightarrow C$.
 - We do not need to compute AB^+ or AC^+ .
- ◆ $B^+ = BC$; yields $B \rightarrow C$.
- ◆ $C^+ = C$; yields nothing.
- ◆ $BC^+ = BC$; yields nothing.


Example -- Continued

- ◆ Resulting FD's: $A \rightarrow B$, $A \rightarrow C$, and $B \rightarrow C$.
- ◆ Projection onto AC : $A \rightarrow C$.
 - ◆ Only FD that involves a subset of $\{A, C\}$.

A Geometric View of FD's

- ◆ Imagine the set of all *instances* of a particular relation.
- ◆ That is, all finite sets of tuples that have the proper number of components.
- ◆ Each instance is a point in this space.

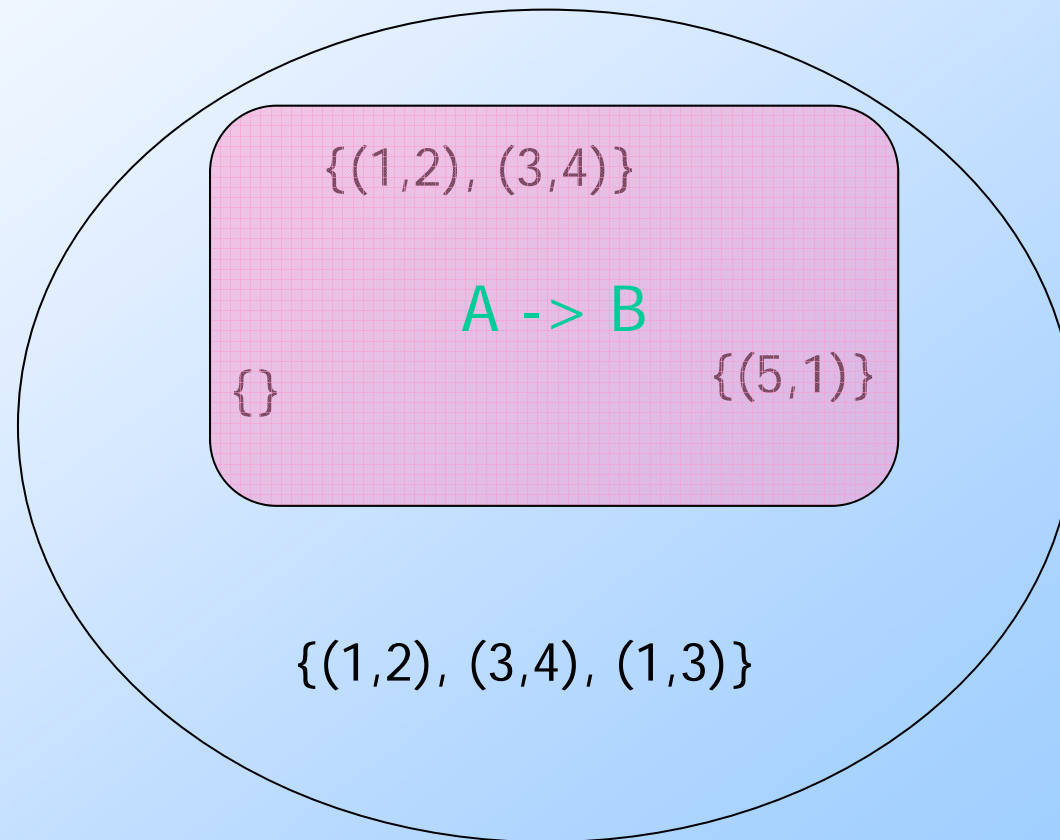
Example: $R(A,B)$



An FD is a Subset of Instances

- ◆ For each FD $X \rightarrow A$ there is a subset of all instances that satisfy the FD.
- ◆ We can represent an FD by a region in the space.
- ◆ Trivial FD = an FD that is represented by the entire space.
 - ◆ Example: $A \rightarrow A$.

Example: $A \rightarrow B$ for $R(A,B)$

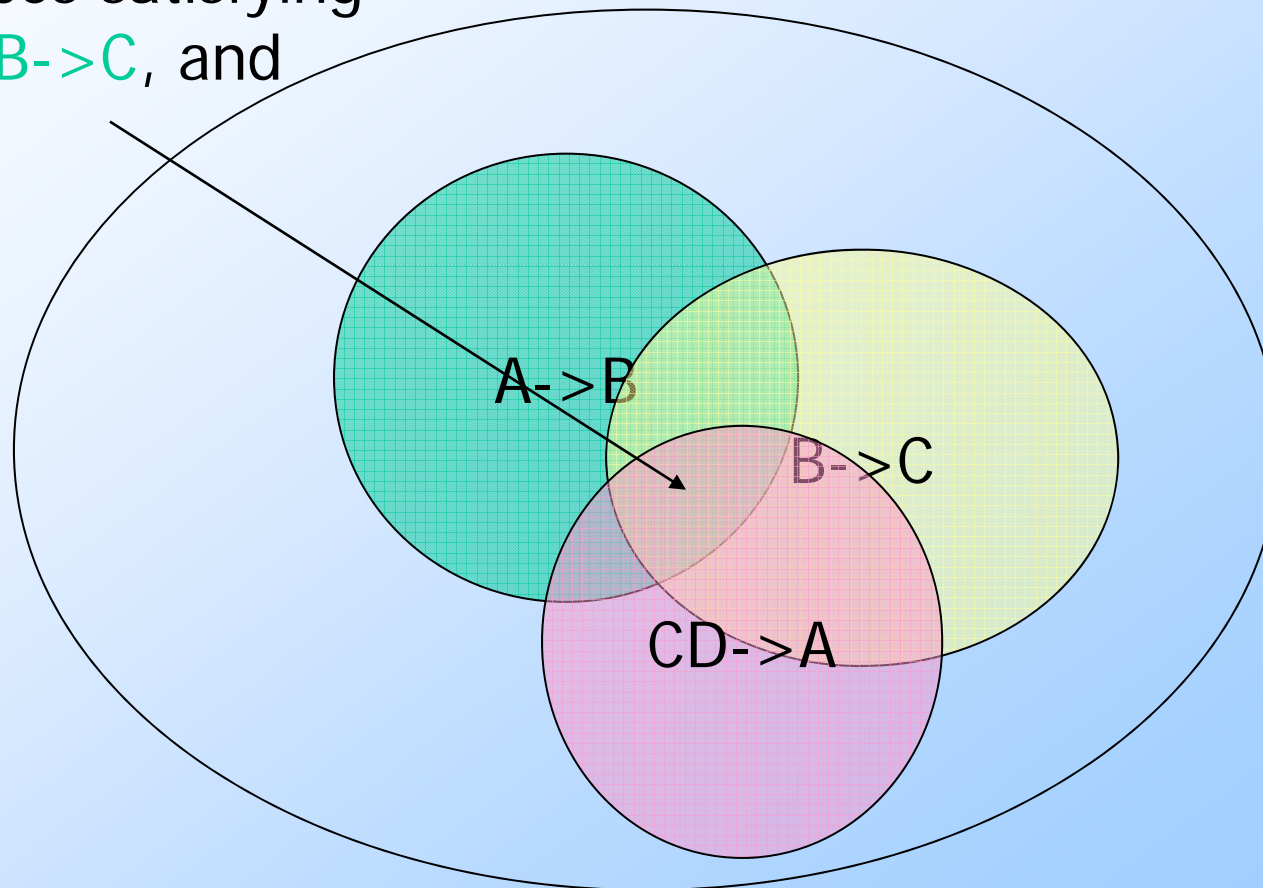


Representing Sets of FD's

- ◆ If each FD is a set of relation instances, then a collection of FD's corresponds to the intersection of those sets.
 - ◆ Intersection = all instances that satisfy all of the FD's.

Example

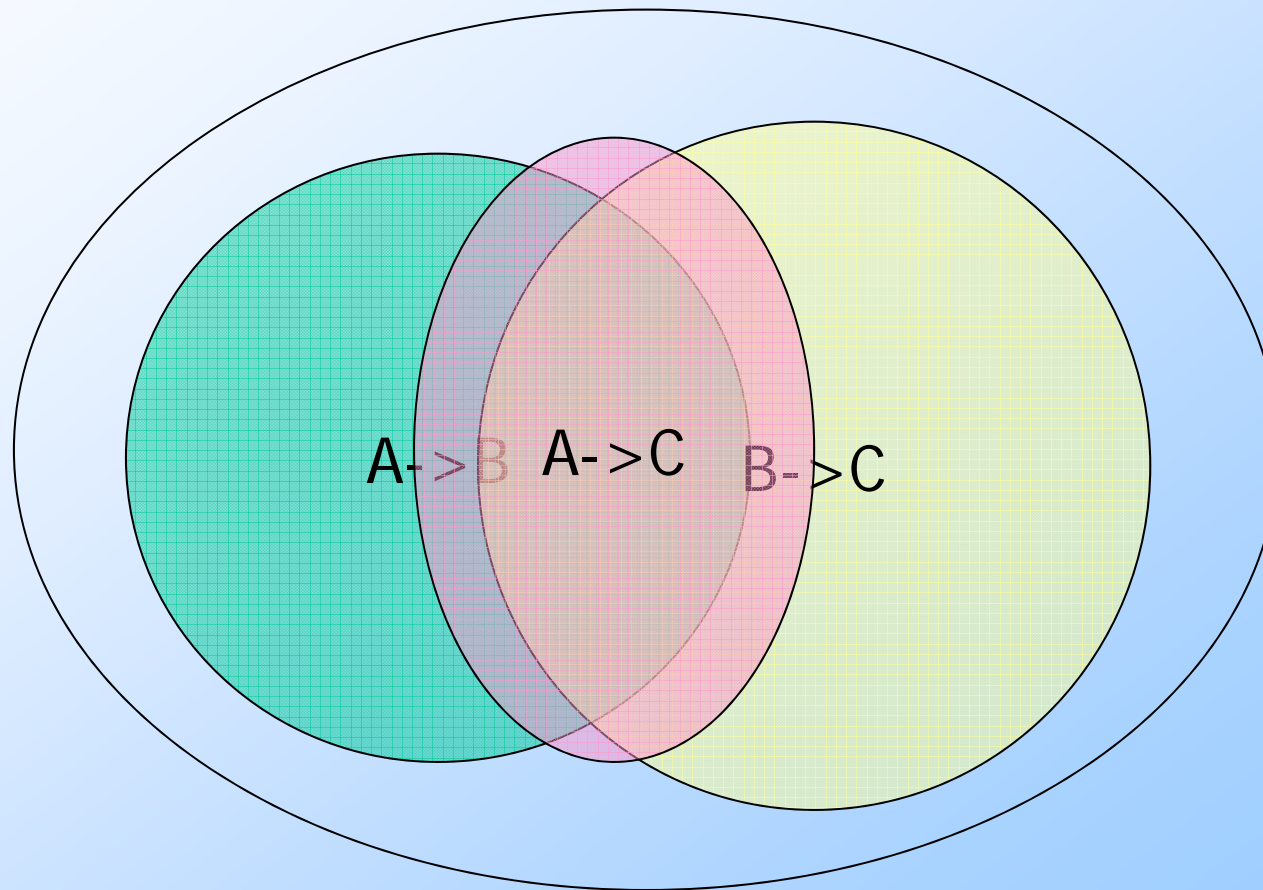
Instances satisfying
 $A \rightarrow B$, $B \rightarrow C$, and
 $CD \rightarrow A$



Implication of FD's

- ◆ If an FD $Y \rightarrow B$ follows from FD's $X_1 \rightarrow A_1, \dots, X_n \rightarrow A_n$, then the region in the space of instances for $Y \rightarrow B$ must include the intersection of the regions for the FD's $X_i \rightarrow A_i$.
 - ◆ That is, every instance satisfying all the FD's $X_i \rightarrow A_i$ surely satisfies $Y \rightarrow B$.
 - ◆ But an instance could satisfy $Y \rightarrow B$, yet not be in this intersection.

Example



Relational Schema Design

- ◆ Goal of relational schema design is to avoid anomalies and redundancy.
 - ◆ *Update anomaly* : one occurrence of a fact is changed, but not all occurrences.
 - ◆ *Deletion anomaly* : valid fact is lost when a tuple is deleted.

Example of Bad Design

Drinkers(name, addr, lemonadesLiked, manf, favLemonade)

<u>name</u>	addr	<u>lemonadesLiked</u>	manf	<u>favLemonade</u>
Janeway	Voyager	Bud	A.B.	WickedAle
Janeway	???	WickedAle	Pete's	???
Spock	Enterprise	Bud	???	Bud

Data is redundant, because each of the ???'s can be figured out by using the FD's name -> addr favLemonade and lemonadesLiked -> manf.

This Bad Design Also Exhibits Anomalies

name	addr	lemonadesLiked	manf	favLemonade
Janeway	Voyager	Bud	A.B.	WickedAle
Janeway	Voyager	WickedAle	Pete's	WickedAle
Spock	Enterprise	Bud	A.B.	Bud

- **Update anomaly:** if Janeway is transferred to *Intrepid*, will we remember to change each of her tuples?
- **Deletion anomaly:** If nobody likes Bud, we lose track of the fact that Anheuser-Busch manufactures Bud.

Boyce-Codd Normal Form

- ◆ We say a relation R is in *BCNF* if whenever $X \rightarrow Y$ is a nontrivial FD that holds in R , X is a superkey.
 - ◆ Remember: *nontrivial* means Y is not contained in X .
 - ◆ Remember, a *superkey* is any superset of a key (not necessarily a proper superset).

Example

Drinkers(name, addr, lemonadesLiked, manf, favLemonade)

FD's: name->addr favLemonade, lemonadesLiked->manf

- ◆ Only key is {name, lemonadesLiked}.
- ◆ In each FD, the left side is *not* a superkey.
- ◆ Any one of these FD's shows *Drinkers* is not in BCNF

Another Example

Lemonades(name, manf, manfAddr)

FD's: name->manf, manf->manfAddr

- ◆ Only key is {name} .
- ◆ name->manf does not violate BCNF, but manf->manfAddr does.

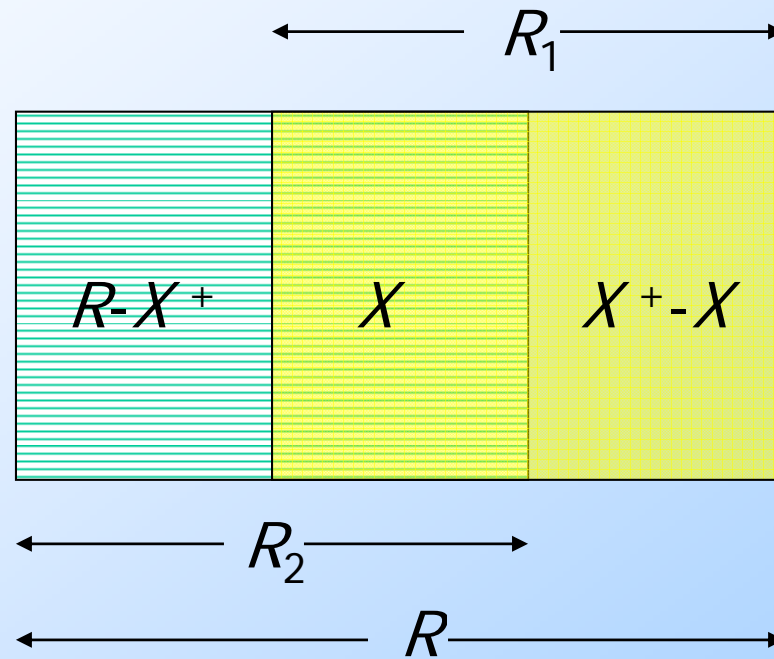
Decomposition into BCNF

- ◆ Given: relation R with FD's F .
- ◆ Look among the given FD's for a BCNF violation $X \rightarrow Y$.
 - ◆ If any FD following from F violates BCNF, then there will surely be an FD in F itself that violates BCNF.
- ◆ Compute X^+ .
 - ◆ Not all attributes, or else X is a superkey.

Decompose R Using $X \rightarrow Y$

- ◆ Replace R by relations with schemas:
 1. $R_1 = X^+$.
 2. $R_2 = R - (X^+ - X)$.
- ◆ *Project* given FD's F onto the two new relations.

Decomposition Picture



Example: BCNF Decomposition

Drinkers(name, addr, lemonadesLiked, manf, favLemonade)

$F = \text{name} \rightarrow \text{addr}, \quad \text{name} \rightarrow \text{favLemonade},$
 $\quad \text{lemonadesLiked} \rightarrow \text{manf}$

- ◆ Pick BCNF violation $\text{name} \rightarrow \text{addr}$.
- ◆ Close the left side: $\{\text{name}\}^+ = \{\text{name}, \text{addr}, \text{favLemonade}\}$.
- ◆ Decomposed relations:
 1. Drinkers1(name, addr, favLemonade)
 2. Drinkers2(name, lemonadesLiked, manf)

Example -- Continued

- ◆ We are not done; we need to check Drinkers1 and Drinkers2 for BCNF.
- ◆ Projecting FD's is easy here.
- ◆ For `Drinkers1(name, addr, favLemonade)`, relevant FD's are `name->addr` and `name->favLemonade`.
 - ◆ Thus, `{name}` is the only key and Drinkers1 is in BCNF.

Example -- Continued

- ◆ For $Drinkers2(\underline{name}, \underline{lemonadesLiked}, manf)$, the only FD is $lemonadesLiked \rightarrow manf$, and the only key is $\{name, lemonadesLiked\}$.
 - ◆ Violation of BCNF.
- ◆ $lemonadesLiked^+ = \{lemonadesLiked, manf\}$, so we decompose *Drinkers2* into:
 1. $Drinkers3(\underline{lemonadesLiked}, manf)$
 2. $Drinkers4(\underline{name}, \underline{lemonadesLiked})$

Example -- Concluded

- ◆ The resulting decomposition of *Drinkers* :
 1. *Drinkers1*(name, addr, favLemonade)
 2. *Drinkers3*(lemonadesLiked, manf)
 3. *Drinkers4*(name, lemonadesLiked)
- ◆ Notice: *Drinkers1* tells us about drinkers, *Drinkers3* tells us about lemonades, and *Drinkers4* tells us the relationship between drinkers and the lemonades they like.

Third Normal Form -- Motivation

- ◆ There is one structure of FD's that causes trouble when we decompose.
- ◆ $AB \rightarrow C$ and $C \rightarrow B$.
 - ◆ Example: A = street address, B = city, C = zip code.
- ◆ There are two keys, $\{A, B\}$ and $\{A, C\}$.
- ◆ $C \rightarrow B$ is a BCNF violation, so we must decompose into AC , BC .

We Cannot Enforce FD's

- ◆ The problem is that if we use AC and BC as our database schema, we cannot enforce the FD $AB \rightarrow C$ by checking FD's in these decomposed relations.
- ◆ Example with $A = \text{street}$, $B = \text{city}$, and $C = \text{zip}$ on the next slide.

An Unenforceable FD

street	zip
545 Tech Sq.	02138
545 Tech Sq.	02139

city	zip
Cambridge	02138
Cambridge	02139

Join tuples with equal zip codes.

street	city	zip
545 Tech Sq.	Cambridge	02138
545 Tech Sq.	Cambridge	02139

Although no FD's were violated in the decomposed relations, FD **street city -> zip** is violated by the database as a whole.

3NF Let's Us Avoid This Problem

- ◆ 3rd Normal Form (3NF) modifies the BCNF condition so we do not have to decompose in this problem situation.
- ◆ An attribute is *prime* if it is a member of any key.
- ◆ $X \rightarrow A$ violates 3NF if and only if X is not a superkey, and also A is not prime.

Example: 3NF

- ◆ In our problem situation with FD's $AB \rightarrow C$ and $C \rightarrow B$, we have keys AB and AC .
- ◆ Thus A , B , and C are each prime.
- ◆ Although $C \rightarrow B$ violates BCNF, it does not violate 3NF.

What 3NF and BCNF Give You

- ◆ There are two important properties of a decomposition:
 1. *Lossless Join* : it should be possible to project the original relations onto the decomposed schema, and then reconstruct the original.
 2. *Dependency Preservation* : it should be possible to check in the projected relations whether all the given FD's are satisfied.

3NF and BCNF -- Continued

- ◆ We can get (1) with a BCNF decomposition.
- ◆ We can get both (1) and (2) with a 3NF decomposition.
- ◆ But we can't always get (1) and (2) with a BCNF decomposition.
 - ◆ street-city-zip is an example.

Testing for a Lossless Join

- ◆ If we project R onto R_1, R_2, \dots, R_k , can we recover R by rejoining?
- ◆ Any tuple in R can be recovered from its projected fragments.
- ◆ So the only question is: **when we rejoin, do we ever get back something we didn't have originally?**

The Chase Test

- ◆ Suppose tuple t comes back in the join.
- ◆ Then t is the join of projections of some tuples of R , one for each R_i of the decomposition.
- ◆ Can we use the given FD's to show that one of these tuples must be t ?

The Chase – (2)

- ◆ Start by assuming $t = abc\dots$.
- ◆ For each i , there is a tuple s_i of R that has a, b, c, \dots in the attributes of R_i .
- ◆ s_i can have any values in other attributes.
- ◆ We'll use the same letter as in t , but with a subscript, for these components.

Example: The Chase

- ◆ Let $R = ABCD$, and the decomposition be AB , BC , and CD .
- ◆ Let the given FD's be $C \rightarrow D$ and $B \rightarrow A$.
- ◆ Suppose the tuple $t = abcd$ is the join of tuples projected onto AB , BC , CD .

The tuples
of R pro-
jected onto
AB, BC, CD.

The *Tableau*

<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>
<i>a</i>	<i>b</i>	<i>c</i> ₁	<i>d</i> ₁
<i>a</i>₂ <i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>₂ <i>d</i>
<i>a</i> ₃	<i>b</i> ₃	<i>c</i>	<i>d</i>

Use *B* > *A*

Use *C* > *D*

We've proved the
second tuple must be *t*.

Summary of the Chase

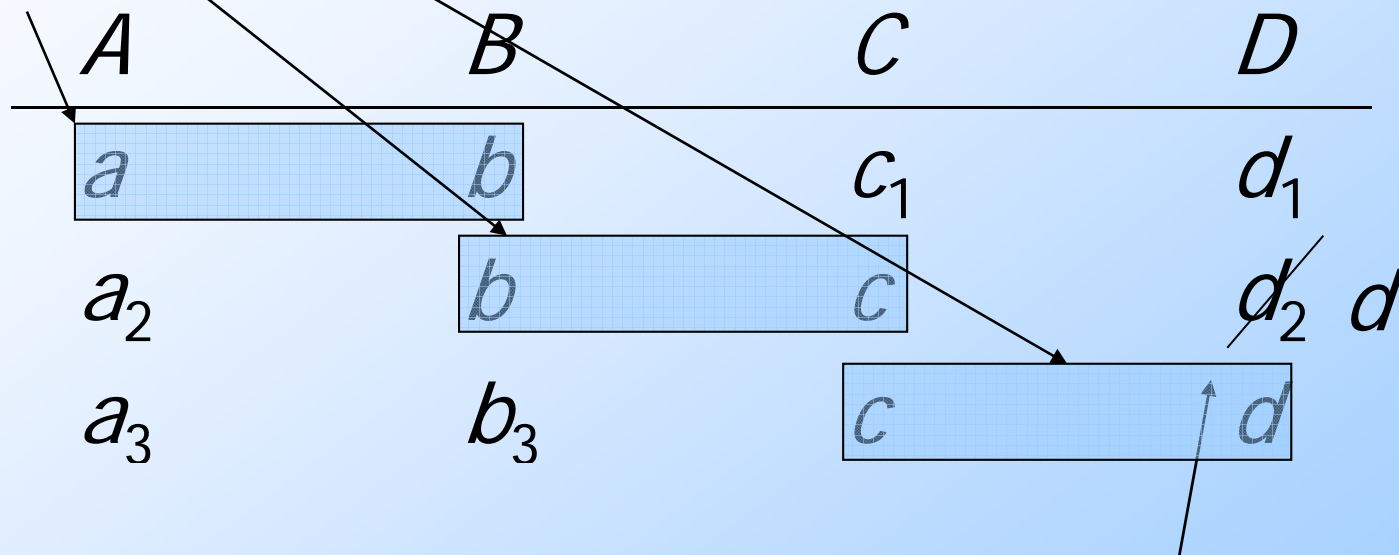
1. If two rows agree in the left side of a FD, make their right sides agree too.
2. Always replace a subscripted symbol by the corresponding unsubscripted one, if possible.
3. If we ever get an unsubscripted row, we know any tuple in the project-join is in the original (the join is lossless).
4. Otherwise, the final tableau is a counterexample.

Example: Lossy Join

- ◆ Same relation $R = ABCD$ and same decomposition.
- ◆ But with only the FD $C \rightarrow D$.

These projections
rejoin to form
abcd.

The *Tableau*



These three tuples are an example *R* that shows the join lossy. *abcd* is not in *R*, but we can project and rejoin to get *abcd*.

Use *C* → *D*

3NF Synthesis Algorithm

- ◆ We can always construct a decomposition into 3NF relations with a lossless join and dependency preservation.
- ◆ Need *minimal basis* for the FD's:
 1. Right sides are single attributes.
 2. No FD can be removed.
 3. No attribute can be removed from a left side.

Constructing a Minimal Basis

1. Split right sides.
2. Repeatedly try to remove an FD and see if the remaining FD's are equivalent to the original.
3. Repeatedly try to remove an attribute from a left side and see if the resulting FD's are equivalent to the original.

3NF Synthesis – (2)

- ◆ One relation for each FD in the minimal basis.
 - ◆ Schema is the union of the left and right sides.
- ◆ If no key is contained in an FD, then add one relation whose schema is some key.

Example: 3NF Synthesis

- ◆ Relation $R = ABCD$.
- ◆ FD's $A \rightarrow B$ and $A \rightarrow C$.
- ◆ Decomposition: AB and AC from the FD's, plus AD for a key.

Why It Works

- ◆ **Preserves dependencies**: each FD from a minimal basis is contained in a relation, thus preserved.
- ◆ **Lossless Join**: use the chase to show that the row for the relation that contains a key can be made all-unsubscripted variables.
- ◆ **3NF**: hard part – a property of minimal bases.