

Introduction to Data Management CSE 414

Unit 6: Conceptual Design
E/R Diagrams
Integrity Constraints
BCNF

(3 lectures)

2

Introduction to Data Management CSE 414

Design Theory and BCNF

CSE 414 – Autumn 2018

72

Relational Schema Design

Name	SSN	PhoneNumber	City
Fred	123-45-6789	206-555-1234	Seattle
Fred	123-45-6789	206-555-6543	Seattle
Joe	987-65-4321	908-555-2121	Westfield

One person may have multiple phones, but lives in only one city

Primary key is thus (SSN, PhoneNumber)

What is the problem with this schema?

CSE 414 – Autumn 2018

73

Relational Schema Design

Name	SSN	PhoneNumber	City
Fred	123-45-6789	206-555-1234	Seattle
Fred	123-45-6789	206-555-6543	Seattle
Joe	987-65-4321	908-555-2121	Westfield

Anomalies:

- **Redundancy** = repeat data
- **Update anomalies** = what if Fred moves to "Bellevue"?
- **Deletion anomalies** = what if Joe deletes his phone number?

CSE 414 – Autumn 2018

74

Relation Decomposition

Break the relation into two:

Name	SSN	PhoneNumber	City
Fred	123-45-6789	206-555-1234	Seattle
Fred	123-45-6789	206-555-6543	Seattle
Joe	987-65-4321	908-555-2121	Westfield

Name	SSN	City	SSN	PhoneNumber
Fred	123-45-6789	Seattle	123-45-6789	206-555-1234
Joe	987-65-4321	Westfield	123-45-6789	206-555-6543
			987-65-4321	908-555-2121

Anomalies have gone:

- No more repeated data
- Easy to move Fred to "Bellevue" (how ?)
- Easy to delete all Joe's phone numbers (how ?)

75

Relational Schema Design (or Logical Design)

How do we do this systematically?

- Start with some relational schema
- Find out its **functional dependencies** (FDs)
- Use FDs to **normalize** the relational schema

CSE 414 – Autumn 2018

76

Functional Dependencies (FDs)

Definition

If two tuples agree on the attributes

A_1, A_2, \dots, A_n

then they must also agree on the attributes

B_1, B_2, \dots, B_m

Formally:

$A_1, A_2, \dots, A_n \rightarrow B_1, B_2, \dots, B_m$

A_1, \dots, A_n determines B_1, \dots, B_m

CSE 414 – Autumn 2018 77

Functional Dependencies (FDs)

Definition $A_1, \dots, A_m \rightarrow B_1, \dots, B_n$ holds in R if:

$\forall t, t' \in R,$
 $(t.A_1 = t'.A_1 \wedge \dots \wedge t.A_m = t'.A_m \rightarrow t.B_1 = t'.B_1 \wedge \dots \wedge t.B_n = t'.B_n)$

R	A_1	...	A_m		B_1	...	B_n		
t									
t'									

if t, t' agree here then t, t' agree here 78

Example

An FD holds, or does not hold on an instance:

EmpID	Name	Phone	Position
E0045	Smith	1234	Clerk
E3542	Mike	9876	Salesrep
E1111	Smith	9876	Salesrep
E9999	Mary	1234	Lawyer

EmpID \rightarrow Name, Phone, Position
 Position \rightarrow Phone
 but not Phone \rightarrow Position

CSE 414 – Autumn 2018 79

Example

EmpID	Name	Phone	Position
E0045	Smith	1234	Clerk
E3542	Mike	9876	Salesrep
E1111	Smith	9876	Salesrep
E9999	Mary	1234	Lawyer

Position \rightarrow Phone

CSE 414 – Autumn 2018 80

Example

EmpID	Name	Phone	Position
E0045	Smith	1234	Clerk
E3542	Mike	9876	Salesrep
E1111	Smith	9876	Salesrep
E9999	Mary	1234	Lawyer

But not Phone \rightarrow Position

CSE 414 – Autumn 2018 81

Example

$name \rightarrow color$
 $category \rightarrow department$
 $color, category \rightarrow price$

name	category	color	department	price
Gizmo	Gadget	Green	Toys	49
Tweaker	Gadget	Green	Toys	99

Do all the FDs hold on this instance?

CSE 414 – Autumn 2018 82

Example

$\text{name} \rightarrow \text{color}$
 $\text{category} \rightarrow \text{department}$
 $\text{color, category} \rightarrow \text{price}$

name	category	color	department	price
Gizmo	Gadget	Green	Toys	49
Tweaker	Gadget	Green	Toys	49
Gizmo	Stationary	Green	Office-suppl.	59

What about this one ?

CSE 414 – Autumn 2018

83

Buzzwords

- FD **holds** or **does not hold** on an instance
- If we can be sure that *every instance of R* will be one in which a given FD is true, then we say that **R satisfies the FD**
- If we say that R satisfies an FD, we are **stating a constraint on R**

CSE 414 – Autumn 2018

84

Why bother with FDs?

Name	SSN	PhoneNumber	City
Fred	123-45-6789	206-555-1234	Seattle
Fred	123-45-6789	206-555-6543	Seattle
Joe	987-65-4321	908-555-2121	Westfield

Anomalies:

- **Redundancy** = repeat data
- **Update anomalies** = what if Fred moves to "Bellevue"?
- **Deletion anomalies** = what if Joe deletes his phone number?

CSE 414 – Autumn 2018

85

An Interesting Observation

If all these FDs are true:

$\text{name} \rightarrow \text{color}$
 $\text{category} \rightarrow \text{department}$
 $\text{color, category} \rightarrow \text{price}$

Then this FD also holds:

$\text{name, category} \rightarrow \text{price}$

CSE 414 – Autumn 2018

86

An Interesting Observation

If all these FDs are true:

$\text{name} \rightarrow \text{color}$
 $\text{category} \rightarrow \text{department}$
 $\text{color, category} \rightarrow \text{price}$

Then this FD also holds:

$\text{name, category} \rightarrow \text{price}$

CSE 414 – Autumn 2018

87

An Interesting Observation

If all these FDs are true:

$\text{name} \rightarrow \text{color}$
 $\text{category} \rightarrow \text{department}$
 $\text{color, category} \rightarrow \text{price}$

Then this FD also holds:

$\text{name, category} \rightarrow \text{price}$

If we find out from application domain that a relation satisfies some FDs, it doesn't mean that we found all the FDs that it satisfies! There could be more FDs implied by the ones we have.

CSE 414 – Autumn 2018

88

Closure of a set of Attributes

Given a set of attributes A_1, \dots, A_n

The **closure** is the set of attributes B , notated $\{A_1, \dots, A_n\}^+$, s.t. $A_1, \dots, A_n \rightarrow B$

Example:
 1. $\text{name} \rightarrow \text{color}$
 2. $\text{category} \rightarrow \text{department}$
 3. $\text{color, category} \rightarrow \text{price}$

Closures:

$\text{name}^+ = \{\text{name, color}\}$
 $\text{color}^+ = \{\text{color}\}$

CSE 414 – Autumn 2018

89

Closure Algorithm

$X = \{A_1, \dots, A_n\}$.

Example:

Repeat until X doesn't change do:
if $B_1, \dots, B_n \rightarrow C$ is a FD **and**
 B_1, \dots, B_n are all in X
then add C to X.

1. $\text{name} \rightarrow \text{color}$
 2. $\text{category} \rightarrow \text{department}$
 3. $\text{color, category} \rightarrow \text{price}$

$\{\text{name, category}\}^+ =$
 $\{\text{name, category,}$ }

CSE 414 – Autumn 2018

90

Closure Algorithm

$X = \{A_1, \dots, A_n\}$.

Example:

Repeat until X doesn't change do:
if $B_1, \dots, B_n \rightarrow C$ is a FD **and**
 B_1, \dots, B_n are all in X
then add C to X.

1. $\text{name} \rightarrow \text{color}$
 2. $\text{category} \rightarrow \text{department}$
 3. $\text{color, category} \rightarrow \text{price}$

$\{\text{name, category}\}^+ =$
 $\{\text{name, category, color,}$ }

CSE 414 – Autumn 2018

91

Closure Algorithm

$X = \{A_1, \dots, A_n\}$.

Example:

Repeat until X doesn't change do:
if $B_1, \dots, B_n \rightarrow C$ is a FD **and**
 B_1, \dots, B_n are all in X
then add C to X.

1. $\text{name} \rightarrow \text{color}$
 2. $\text{category} \rightarrow \text{department}$
 3. $\text{color, category} \rightarrow \text{price}$

$\{\text{name, category}\}^+ =$
 $\{\text{name, category, color, department}$ }

CSE 414 – Autumn 2018

92

Closure Algorithm

$X = \{A_1, \dots, A_n\}$.

Example:

Repeat until X doesn't change do:
if $B_1, \dots, B_n \rightarrow C$ is a FD **and**
 B_1, \dots, B_n are all in X
then add C to X.

1. $\text{name} \rightarrow \text{color}$
 2. $\text{category} \rightarrow \text{department}$
 3. $\text{color, category} \rightarrow \text{price}$

$\{\text{name, category}\}^+ =$
 $\{\text{name, category, color, department, price}\}$

CSE 414 – Autumn 2018

93

Closure Algorithm

$X = \{A_1, \dots, A_n\}$.

Example:

Repeat until X doesn't change do:
if $B_1, \dots, B_n \rightarrow C$ is a FD **and**
 B_1, \dots, B_n are all in X
then add C to X.

1. $\text{name} \rightarrow \text{color}$
 2. $\text{category} \rightarrow \text{department}$
 3. $\text{color, category} \rightarrow \text{price}$

$\{\text{name, category}\}^+ =$
 $\{\text{name, category, color, department, price}\}$

Hence: $\text{name, category} \rightarrow \text{color, department, price}$

CSE 414 – Autumn 2018

94

Example

In class:

$R(A,B,C,D,E,F)$

A, B	→	C
A, D	→	E
B	→	D
A, F	→	B

Compute $\{A,B\}^+$ $X = \{A, B, \quad \quad \quad \}$

Compute $\{A, F\}^+$ $X = \{A, F, \quad \quad \quad \}$

CSE 414 – Autumn 2018 95

Example

In class:

$R(A,B,C,D,E,F)$

A, B	→	C
A, D	→	E
B	→	D
A, F	→	B

Compute $\{A,B\}^+$ $X = \{A, B, C, D, E \quad \quad \}$

Compute $\{A, F\}^+$ $X = \{A, F, \quad \quad \quad \}$

CSE 414 – Autumn 2018 96

Example

In class:

$R(A,B,C,D,E,F)$

A, B	→	C
A, D	→	E
B	→	D
A, F	→	B

Compute $\{A,B\}^+$ $X = \{A, B, C, D, E \quad \quad \}$

Compute $\{A, F\}^+$ $X = \{A, F, B, C, D, E \quad \quad \}$

CSE 414 – Autumn 2018 97

Example

In class:

$R(A,B,C,D,E,F)$

A, B	→	C
A, D	→	E
B	→	D
A, F	→	B

Compute $\{A,B\}^+$ $X = \{A, B, C, D, E \quad \quad \}$

Compute $\{A, F\}^+$ $X = \{A, F, B, C, D, E \quad \quad \}$

CSE 414 – Autumn 2018 What is the key of R?

Practice at Home

Find all FD's implied by:

A, B	→	C
A, D	→	B
B	→	D

CSE 414 – Autumn 2018 99

Practice at Home

Find all FD's implied by:

A, B	→	C
A, D	→	B
B	→	D

Step 1: Compute X^+ , for every X:

$A^+ = A, B^+ = BD, C^+ = C, D^+ = D$
 $AB^+ = ABCD, AC^+ = AC, AD^+ = ABCD,$
 $BC^+ = BCD, BD^+ = BD, CD^+ = CD$
 $ABC^+ = ABD^+ = ACD^+ = ABCD$ (no need to compute– why ?)
 $BCD^+ = BCD, ABCD^+ = ABCD$

Step 2: Enumerate all FD's $X \rightarrow Y$, s.t. $Y \subseteq X^+$ and $X \cap Y = \emptyset$:

$AB \rightarrow CD, AD \rightarrow BC, ABC \rightarrow D, ABD \rightarrow C, ACD \rightarrow B$

100

Keys

- A **superkey** is a set of attributes A_1, \dots, A_n s.t. for any other attribute B , we have $A_1, \dots, A_n \rightarrow B$
- A **key** is a minimal superkey
 - A superkey and for which no subset is a superkey

CSE 414 – Autumn 2018

101

Computing (Super)Keys

- For all sets X , compute X^+
- If $X^+ = [\text{all attributes}]$, then X is a superkey
- Try reducing to the minimal X 's to get the key

CSE 414 – Autumn 2018

102

Example

Product(name, price, category, color)

name, category \rightarrow price
category \rightarrow color

What is the key ?

CSE 414 – Autumn 2018

103

Example

Product(name, price, category, color)

name, category \rightarrow price
category \rightarrow color

What is the key ?

$(\text{name, category})^+ = \{\text{name, category, price, color}\}$

Hence (name, category) is a key

CSE 414 – Autumn 2018

104

Example

Product(name, price, category, color)

name, category \rightarrow price
category \rightarrow color

What is the key ?

$(\text{name, category})^+ = \{\text{name, category, price, color}\}$

CSE 414 – Autumn 2018

105

Key or Keys ?

Can we have more than one key ?

Given $R(A,B,C)$ define FD's s.t. there are two or more distinct keys


CSE 414 – Autumn 2018

106

Key or Keys ?

Can we have more than one key ?

Given $R(A,B,C)$ define FD's s.t. there are two or more distinct keys



what are the keys here ?

CSE 414 – Autumn 2018

107

Eliminating Anomalies

Main idea:

- $X \rightarrow A$ is OK if X is a (super)key
- $X \rightarrow A$ is not OK otherwise
 - Need to decompose the table, but how?

Boyce-Codd Normal Form

CSE 414 – Autumn 2018

108

Boyce-Codd Normal Form


Dr. Raymond F. Boyce

CSE 414 – Autumn 2018

109

Edgar Frank "Ted" Codd

"A Relational Model of Data for Large Shared Data Banks"



CSE 414 – Autumn 2018

110

Boyce-Codd Normal Form

If there are no "bad" FDs:

Definition. A relation R is in BCNF if:
Whenever $X \rightarrow B$ is a non-trivial dependency, then X is a superkey.

Equivalently:


Definition. A relation R is in BCNF if:
 $\forall X$, either $X^+ = X$ (i.e., X is not in any FDs)
or $X^+ = [\text{all attributes}]$ (computed using FDs)

CSE 414 – Autumn 2018

111

BCNF Decomposition Algorithm

Normalize(R)
find X s.t.: $X \neq X^+$ and $X^+ \neq [\text{all attributes}]$
if (not found) **then** "R is in BCNF"
let $Y = X^+ - X$; $Z = [\text{all attributes}] - X^+$
decompose R into $R_1(X \cup Y)$ and $R_2(X \cup Z)$
Normalize(R_1); Normalize(R_2);



CSE 414 – Autumn 2018

112

Example

Name	SSN	PhoneNumber	City
Fred	123-45-6789	206-555-1234	Seattle
Fred	123-45-6789	206-555-6543	Seattle
Joe	987-65-4321	908-555-2121	Westfield
Joe	987-65-4321	908-555-1234	Westfield

$SSN \rightarrow Name, City$



The only key is: {SSN, PhoneNumber}

Hence $SSN \rightarrow Name, City$ is a "bad" dependency

In other words:

$SSN+ = SSN, Name, City$ and is neither SSN nor All Attributes

Example BCNF Decomposition

Name	SSN	City
Fred	123-45-6789	Seattle
Joe	987-65-4321	Westfield

$SSN \rightarrow Name, City$



Let's check anomalies:

- Redundancy ?
- Update ?
- Delete ?

SSN	PhoneNumber
123-45-6789	206-555-1234
123-45-6789	206-555-6543
987-65-4321	908-555-2121
987-65-4321	908-555-1234

CSE 414 – Autumn 2018

114

Find X s.t.: $X \neq X^+$ and $X^+ \neq [all\ attributes]$

Example BCNF Decomposition

Person(name, SSN, age, hairColor, phoneNumber)

$SSN \rightarrow name, age$

$age \rightarrow hairColor$



CSE 414 – Autumn 2018

115

Find X s.t.: $X \neq X^+$ and $X^+ \neq [all\ attributes]$

Example BCNF Decomposition

Person(name, SSN, age, hairColor, phoneNumber)

$SSN \rightarrow name, age$

$age \rightarrow hairColor$

Iteration 1: Person: $SSN+ = SSN, name, age, hairColor$
Decompose into: P(SSN, name, age, hairColor)
Phone(SSN, phoneNumber)



CSE 414 – Autumn 2018

116

Find X s.t.: $X \neq X^+$ and $X^+ \neq [all\ attributes]$

Example BCNF Decomposition

Person(name, SSN, age, hairColor, phoneNumber)

$SSN \rightarrow name, age$

$age \rightarrow hairColor$

What are the keys?

Iteration 1: Person: $SSN+ = SSN, name, age, hairColor$
Decompose into: P(SSN, name, age, hairColor)
Phone(SSN, phoneNumber)

Iteration 2: P: $age+ = age, hairColor$

Decompose: People(SSN, name, age)
Hair(age, hairColor)
Phone(SSN, phoneNumber)

CSE 414 – Autumn 2018

117

Find X s.t.: $X \neq X^+$ and $X^+ \neq [all\ attributes]$

Example BCNF Decomposition

Person(name, SSN, age, hairColor, phoneNumber)

$SSN \rightarrow name, age$

$age \rightarrow hairColor$

Note the keys!

Iteration 1: Person: $SSN+ = SSN, name, age, hairColor$
Decompose into: P(SSN, name, age, hairColor)
Phone(SSN, phoneNumber)

Iteration 2: P: $age+ = age, hairColor$

Decompose: People(SSN, name, age)
Hair(age, hairColor)
Phone(SSN, phoneNumber)

CSE 414 – Autumn 2018

118

R(A,B,C,D)

Example: BCNF

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

R(A,B,C,D)

CSE 414 – Autumn 2018 119

R(A,B,C,D)

Example: BCNF

Recall: find X s.t.
 $X \subsetneq X^+ \subseteq [\text{all-attrs}]$

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

R(A,B,C,D)

CSE 414 – Autumn 2018 120

R(A,B,C,D)

Example: BCNF

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

R(A,B,C,D)
 $A^+ = ABC \neq ABCD$

CSE 414 – Autumn 2018 121

R(A,B,C,D)

Example: BCNF

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

R(A,B,C,D)
 $A^+ = ABC \neq ABCD$

$R_1(A,B,C)$ $R_2(A,D)$

CSE 414 – Autumn 2018 122

R(A,B,C,D)

Example: BCNF

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

R(A,B,C,D)
 $A^+ = ABC \neq ABCD$

$R_1(A,B,C)$
 $B^+ = BC \neq ABC$

$R_2(A,D)$

CSE 414 – Autumn 2018 123

R(A,B,C,D)

Example: BCNF

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

R(A,B,C,D)
 $A^+ = ABC \neq ABCD$

$R_1(A,B,C)$
 $B^+ = BC \neq ABC$

$R_{11}(B,C)$ $R_{12}(A,B)$

$R_2(A,D)$

What are the keys?

What happens if in R we first pick B^+ ? Or AB^+ ?

CSE 414 – Autumn 2018 124

Getting Practical

How to implement normalization in SQL

CSE 414 – Autumn 2018

140

Motivation

- We learned about how to normalize tables to avoid anomalies
- How can we implement normalization in SQL if we can't modify existing tables?
 - This might be due to legacy applications that rely on previous schemas to run

CSE 414 – Autumn 2018

141

Views

- A **view** in SQL =
 - A table computed from other tables, s.t., whenever the base tables are updated, the view is updated too
- More generally:
 - A **view** is derived data that keeps track of changes in the original data
- Compare:
 - A **function** computes a value from other values, but does not keep track of changes to the inputs

CSE 414 – Autumn 2018

142

Purchase(customer, product, store)
Product(pname, price)

StorePrice(store, price)

A Simple View

Create a view that returns for each store the prices of products purchased at that store

```
CREATE VIEW StorePrice AS
SELECT DISTINCT x.store, y.price
FROM Purchase x, Product y
WHERE x.product = y.pname
```

This is like a new table
StorePrice(store, price)

CSE 414 – Autumn 2018

143

Purchase(customer, product, store)
Product(pname, price)

StorePrice(store, price)

We Use a View Like Any Table

- A "high end" store is a store that sell some products over 1000.
- For each customer, return all the high end stores that they visit.

```
SELECT DISTINCT u.customer, u.store
FROM Purchase u, StorePrice v
WHERE u.store = v.store
AND v.price > 1000
```

CSE 414 – Autumn 2018

144

Types of Views

- **Virtual views**
 - Computed only on-demand – slow at runtime
 - Always up to date
- **Materialized views**
 - Pre-computed offline – fast at runtime
 - May have stale data (must recompute or update)
 - Indexes are materialized views
- A key component of physical tuning of databases is the selection of materialized views and indexes

CSE 414 – Autumn 2018

145

Vertical Partitioning

Resumes

SSN	Name	Address	Resume	Picture
234234	Mary	Houston	Doc1...	JPG1...
345345	Sue	Seattle	Doc2...	JPG2...
345343	Joan	Seattle	Doc3...	JPG3...
432432	Ann	Portland	Doc4...	JPG4...

T1

SSN	Name	Address
234234	Mary	Houston
345345	Sue	Seattle
...		

T2

SSN	Resume
234234	Doc1...
345345	Doc2...

T3

SSN	Picture
234234	JPG1...
345345	JPG2...

T2.SSN is a key and a foreign key to T1.SSN. Same for T3.SSN ¹⁴⁶

T1(ssn,name,address) Resumes(ssn,name,address,resume,picture)
T2(ssn,resume)
T3(ssn,picture)

Vertical Partitioning

```
CREATE VIEW Resumes AS
SELECT T1.ssn, T1.name, T1.address,
       T2.resume, T3.picture
FROM   T1,T2,T3
WHERE  T1.ssn=T2.ssn AND T1.ssn=T3.ssn
```

CSE 414 – Autumn 2018 147

T1(ssn,name,address) Resumes(ssn,name,address,resume,picture)
T2(ssn,resume)
T3(ssn,picture)

Vertical Partitioning

```
CREATE VIEW Resumes AS
SELECT T1.ssn, T1.name, T1.address,
       T2.resume, T3.picture
FROM   T1,T2,T3
WHERE  T1.ssn=T2.ssn AND T1.ssn=T3.ssn
```

```
SELECT address
FROM   Resumes
WHERE  name = 'Sue'
```

CSE 414 – Autumn 2018 148

T1(ssn,name,address) Resumes(ssn,name,address,resume,picture)
T2(ssn,resume)
T3(ssn,picture)

Vertical Partitioning

```
CREATE VIEW Resumes AS
SELECT T1.ssn, T1.name, T1.address,
       T2.resume, T3.picture
FROM   T1,T2,T3
WHERE  T1.ssn=T2.ssn AND T1.ssn=T3.ssn
```

```
SELECT address
FROM   Resumes
WHERE  name = 'Sue'
```

Original query:

```
SELECT T1.address
FROM   T1, T2, T3
WHERE  T1.name = 'Sue'
       AND T1.SSN=T2.SSN
       AND T1.SSN = T3.SSN
```

CSE 414 – Autumn 2018 149

T1(ssn,name,address) Resumes(ssn,name,address,resume,picture)
T2(ssn,resume)
T3(ssn,picture)

Vertical Partitioning

```
CREATE VIEW Resumes AS
SELECT T1.ssn, T1.name, T1.address,
       T2.resume, T3.picture
FROM   T1,T2,T3
WHERE  T1.ssn=T2.ssn AND T1.ssn=T3.ssn
```

```
SELECT address
FROM   Resumes
WHERE  name = 'Sue'
```

Modified query:

```
SELECT T1.address
FROM   T1, T2, T3
WHERE  T1.name = 'Sue'
       AND T1.SSN=T2.SSN
       AND T1.SSN = T3.SSN
```

Final query:

```
SELECT T1.address
FROM   T1
WHERE  T1.name = 'Sue'
```

150

Vertical Partitioning Applications

- **Advantages**
 - Speeds up queries that touch only a small fraction of columns
 - Single column can be compressed effectively, reducing disk I/O
- **Disadvantages**
 - Updates are expensive!
 - Need many joins to access many columns
 - Repeated key columns add overhead

CSE 414 – Autumn 2018 151

Horizontal Partitioning

Customers

SSN	Name	City
234234	Mary	Houston
345345	Sue	Seattle
345343	Joan	Seattle
234234	Ann	Portland
--	Frank	Calgary
--	Jean	Montreal

.....

CustomersInHouston

SSN	Name	City
234234	Mary	Houston

CustomersInSeattle

SSN	Name	City
345345	Sue	Seattle
345343	Joan	Seattle

.....

CSE 414 – Autumn 2018 152

CustomersInHouston(ssn,name,city) Customers(ssn,name,city)
 CustomersInSeattle(ssn,name,city)

Horizontal Partitioning

```
CREATE VIEW Customers AS
  CustomersInHouston
  UNION ALL
  CustomersInSeattle
  UNION ALL
  ...
```

CSE 414 – Autumn 2018 153

CustomersInHouston(ssn,name,city) Customers(ssn,name,city)
 CustomersInSeattle(ssn,name,city)

Horizontal Partitioning

```
SELECT name
FROM Customers
WHERE city = 'Seattle'
```

Which tables are inspected by the system ?

CSE 414 – Autumn 2018 154

CustomersInHouston(ssn,name,city) Customers(ssn,name,city)
 CustomersInSeattle(ssn,name,city)

Horizontal Partitioning

Better: remove CustomerInHouston.city etc

```
CREATE VIEW Customers AS
  (SELECT SSN, name, 'Houston' as city
   FROM CustomersInHouston)
  UNION ALL
  (SELECT SSN, name, 'Seattle' as city
   FROM CustomersInSeattle)
  UNION ALL
  ...
```

CSE 414 – Autumn 2018 155

CustomersInHouston(ssn,name,city) Customers(ssn,name,city)
 CustomersInSeattle(ssn,name,city)

Horizontal Partitioning

```
SELECT name
FROM Customers
WHERE city = 'Seattle'
```

```
SELECT name
FROM CustomersInSeattle
```

CSE 414 – Autumn 2018 156

Horizontal Partitioning Applications

- Performance optimization
 - Especially for data warehousing
 - E.g., one partition per month
 - E.g., archived applications and active applications
- Distributed and parallel databases
- Data integration

CSE 414 – Autumn 2018 157

Conclusion

- Poor schemas can lead to performance inefficiencies
- E/R diagrams are means to structurally visualize and design relational schemas
- Normalization is a principled way of converting schemas into a form that avoid such problems
- BCNF is one of the most widely used normalized form in practice

158