

Introduction to Database Systems

CSE 414

Lecture 3: SQL Basics

Review

- Relational data model
 - Schema + instance + query language
- Query language: SQL
 - Create tables
 - Retrieve records from tables
 - Declare keys and foreign keys

Discussion

- Tables are NOT ordered
 - they are sets or multisets (bags)
- Tables are FLAT
 - No nested attributes
- Tables DO NOT prescribe how they are implemented / stored on disk
 - This is called **physical data independence**

Table Implementation

- How would you implement this?

<u>cname</u>	country	no_employees	for_profit
GizmoWorks	USA	20000	True
Canon	Japan	50000	True
Hitachi	Japan	30000	True
HappyCam	Canada	500	False

Table Implementation

- How would you implement this?

<u>cname</u>	country	no_employees	for_profit
GizmoWorks	USA	20000	True
Canon	Japan	50000	True
Hitachi	Japan	30000	True
HappyCam	Canada	500	False

Row major: as an array of objects

GizmoWorks	Canon	Hitachi	HappyCam
USA	Japan	Japan	Canada
20000	50000	30000	500
True	True	True	False

Table Implementation

- How would you implement this?

<u>cname</u>	country	no_employees	for_profit
GizmoWorks	USA	20000	True
Canon	Japan	50000	True
Hitachi	Japan	30000	True
HappyCam	Canada	500	False

Column major: as one array per attribute

GizmoWorks	Canon	Hitachi	HappyCam
USA	Japan	Japan	Canada
20000	50000	30000	500
True	True	True	False

CSE 414 - Autumn 2018

Table Implementation

- How would you implement this?

<u>cname</u>	country	no_employees	for_profit
GizmoWorks	USA	20000	True
Canon	Japan	50000	True
Hitachi	Japan	30000	True
HappyCam	Canada	500	False

Physical data independence

The logical definition of the data remains unchanged, even when we make changes to the actual implementation

First Normal Form

<u>cname</u>	country	no_employees	for_profit
Canon	Japan	50000	Y
Hitachi	Japan	30000	Y

- All relations must be flat: we say that the relation is in *first normal form*

First Normal Form

<u>cname</u>	country	no_employees	for_profit
Canon	Japan	50000	Y
Hitachi	Japan	30000	Y

- All relations must be flat: we say that the relation is in *first normal form*
- E.g., we want to add products manufactured by each company:

First Normal Form

<u>cname</u>	country	no_employees	for_profit
Canon	Japan	50000	Y
Hitachi	Japan	30000	Y

- All relations must be flat: we say that the relation is in *first normal form*
- E.g., we want to add products manufactured by each company:

<u>cname</u>	country	no_employees	for_profit	products									
Canon	Japan	50000	Y	<table border="1"> <thead> <tr> <th><u>pname</u></th> <th>price</th> <th>category</th> </tr> </thead> <tbody> <tr> <td>SingleTouch</td> <td>149.99</td> <td>Photography</td> </tr> <tr> <td>Gadget</td> <td>200</td> <td>Toy</td> </tr> </tbody> </table>	<u>pname</u>	price	category	SingleTouch	149.99	Photography	Gadget	200	Toy
<u>pname</u>	price	category											
SingleTouch	149.99	Photography											
Gadget	200	Toy											
Hitachi	Japan	30000	Y	<table border="1"> <thead> <tr> <th><u>pname</u></th> <th>price</th> <th>category</th> </tr> </thead> <tbody> <tr> <td>AC</td> <td>300</td> <td>Appliance</td> </tr> </tbody> </table>	<u>pname</u>	price	category	AC	300	Appliance			
<u>pname</u>	price	category											
AC	300	Appliance											

First Normal Form

<u>cname</u>	country	no_employees	for_profit
Canon	Japan	50000	Y
Hitachi	Japan	30000	Y

- All relations must be flat: we say that the relation is in *first normal form*
- E.g., we want to add products manufactured by each company:

Non-1NF!

<u>cname</u>	country	no_employees	for_profit	products									
Canon	Japan	50000	Y	<table border="1"> <thead> <tr> <th><u>pname</u></th> <th>price</th> <th>category</th> </tr> </thead> <tbody> <tr> <td>SingleTouch</td> <td>149.99</td> <td>Photography</td> </tr> <tr> <td>Gadget</td> <td>200</td> <td>Toy</td> </tr> </tbody> </table>	<u>pname</u>	price	category	SingleTouch	149.99	Photography	Gadget	200	Toy
<u>pname</u>	price	category											
SingleTouch	149.99	Photography											
Gadget	200	Toy											
Hitachi	Japan	30000	Y	<table border="1"> <thead> <tr> <th><u>pname</u></th> <th>price</th> <th>category</th> </tr> </thead> <tbody> <tr> <td>AC</td> <td>300</td> <td>Appliance</td> </tr> </tbody> </table>	<u>pname</u>	price	category	AC	300	Appliance			
<u>pname</u>	price	category											
AC	300	Appliance											

First Normal Form

We will learn how different languages and data models handle this aspect

Now it's in 1NF

Company

<u>cname</u>	country	no_employees	for_profit
Canon	Japan	50000	Y
Hitachi	Japan	30000	Y

Products

<u>pname</u>	price	category	manufacturer
SingleTouch	149.99	Photography	Canon
AC	300	Appliance	Hitachi
Gadget	200	Toy	Canon

SQL

- **Structured Query Language**
- Most widely used language to query relational data
- One of the many languages for querying relational data
- **A declarative programming language**

Selections in SQL

```
SELECT *  
FROM Product  
WHERE price > 100.0
```

Demo 2

Product(pname, price, category, manufacturer)

Company(cname, country)

Joins in SQL

pname	price	category	manufacturer
MultiTouch	199.99	gadget	Canon
SingleTouch	49.99	photography	Canon
Gizom	50	gadget	GizmoWorks
SuperGizmo	250.00	gadget	GizmoWorks

cname	country
GizmoWorks	USA
Canon	Japan
Hitachi	Japan

Retrieve all Japanese products that cost < \$150

Product(pname, price, category, manufacturer)

Company(cname, country)

Joins in SQL

pname	price	category	manufacturer
MultiTouch	199.99	gadget	Canon
SingleTouch	49.99	photography	Canon
Gizom	50	gadget	GizmoWorks
SuperGizmo	250.00	gadget	GizmoWorks

cname	country
GizmoWorks	USA
Canon	Japan
Hitachi	Japan

Retrieve all Japanese products that cost < \$150

```
SELECT pname, price
FROM Product, Company
WHERE ...
```

Product(pname, price, category, manufacturer)
Company(cname, country)

Joins in SQL

pname	price	category	manufacturer
MultiTouch	199.99	gadget	Canon
SingleTouch	49.99	photography	Canon
Gizom	50	gadget	GizmoWorks
SuperGizmo	250.00	gadget	GizmoWorks

cname	country
GizmoWorks	USA
Canon	Japan
Hitachi	Japan

Retrieve all Japanese products that cost < \$150

```
SELECT P.pname, C.price
FROM Product as P, Company as C
WHERE P.manufacturer=C.cname AND
      P.country='Japan' AND C.price < 150
```

Join Predicate

Selection predicates

Product(pname, price, category, manufacturer)

Company(cname, country)

Joins in SQL

pname	price	category	manufacturer
MultiTouch	199.99	gadget	Canon
SingleTouch	49.99	photography	Canon
Gizom	50	gadget	GizmoWorks
SuperGizmo	250.00	gadget	GizmoWorks

cname	country
GizmoWorks	USA
Canon	Japan
Hitachi	Japan

Retrieve all USA companies that manufacture “gadget” products

Product(pname, price, category, manufacturer)

Company(cname, country)

Joins in SQL

pname	price	category	manufacturer
MultiTouch	199.99	gadget	Canon
SingleTouch	49.99	photography	Canon
Gizom	50	gadget	GizmoWorks
SuperGizmo	250.00	gadget	GizmoWorks

cname	country
GizmoWorks	USA
Canon	Japan
Hitachi	Japan

Retrieve all USA companies
that manufacture “gadget” products

```
SELECT DISTINCT cname  
FROM Product, Company  
WHERE country='USA' AND category = 'gadget'  
AND manufacturer = cname
```

Why
DISTINCT?

Demo 2 – cont'd

Joins in SQL

- The standard join in SQL is sometimes called an **inner join**
 - Each row in the result **must come from both tables in the join**
- Sometimes we want to include rows from only one of the two table: **outer join**

Employee(id, name)
Sales(employeeID, productID)

Inner Join

Employee

<u>id</u>	name
1	Joe
2	Jack
3	Jill

Sales

<u>employeeID</u>	productID
1	344
1	355
2	544

Retrieve employees and their sales

Employee(id, name)
Sales(employeeID, productID)

Inner Join

Employee

<u>id</u>	name
1	Joe
2	Jack
3	Jill

Sales

<u>employeeID</u>	productID
1	344
1	355
2	544

Retrieve employees and their sales

```
SELECT *  
FROM Employee E, Sales S  
WHERE E.id = S.employeeID
```

Employee(id, name)
Sales(employeeID, productID)

Inner Join

Employee

<u>id</u>	name
1	Joe
2	Jack
3	Jill

Sales

<u>employeeID</u>	productID
1	344
1	355
2	544

Retrieve employees and their sales

```
SELECT *  
FROM Employee E, Sales S  
WHERE E.id = S.employeeID
```

id	name	empolyeeID	productID
1	Joe	1	344
1	Joe	1	355
2	Jack	2	544

Employee(id, name)
Sales(employeeID, productID)

Inner Join

Employee

<u>id</u>	name
1	Joe
2	Jack
3	Jill

Sales

<u>employeeID</u>	productID
1	344
1	355
2	544

Retrieve employees and their sales

```
SELECT *  
FROM Employee E, Sales S  
WHERE E.id = S.employeeID
```

id	name	empolyeeID	productID
1	Joe	1	344
1	Joe	1	355
2	Jack	2	544

Jill is missing

Employee(id, name)
Sales(employeeID, productID)

Inner Join

Employee

<u>id</u>	name
1	Joe
2	Jack
3	Jill

Sales

<u>employeeID</u>	productID
1	344
1	355
2	544

Retrieve employees and their sales

```
SELECT *  
FROM Employee E  
INNER JOIN  
Sales S  
ON E.id = S.employeeID
```

Alternative
syntax

id	name	empolyeeID	productID
1	Joe	1	344
1	Joe	1	355
2	Jack	2	544

Jill is
missing

Employee(id, name)
 Sales(employeeID, productID)

Outer Join

Employee

<u>id</u>	name
1	Joe
2	Jack
3	Jill

Sales

<u>employeeID</u>	productID
1	344
1	355
2	544

Retrieve employees and their sales

Jill is present

```
SELECT *
FROM Employee E
LEFT OUTER JOIN
Sales S
ON E.id = S.employeeID
```

id	name	empolyeeID	productID
1	Joe	1	344
1	Joe	1	355
2	Jack	2	544
3	Jill	NULL	NULL