

Introduction to Database Systems CSE 414

Lecture 3: SQL Basics

CSE 414 - Autumn 2018

1

Review

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

CSE 414 - Autumn 2018

2

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**

CSE 414 - Autumn 2018

3

Table Implementation

- How would you implement this?

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

CSE 414 - Autumn 2018

4

Table Implementation

- How would you implement this?

cname	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

CSE 414 - Autumn 2018

5

Table Implementation

- How would you implement this?

cname	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

6

Table Implementation

- How would you implement this?

cname	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

7

First Normal Form

cname	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*

CSE 414 - Autumn 2018

8

First Normal Form

cname	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:

CSE 414 - Autumn 2018

9

First Normal Form

cname	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:

cname	country	no_employees	for_profit	products		
				name	price	category
Canon	Japan	50000	Y	SingleTouch	149.99	Photography
				Gadget	200	Toy
Hitachi	Japan	30000	Y	AC	300	Appliance

First Normal Form

cname	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!

cname	country	no_employees	for_profit	products		
				name	price	category
Canon	Japan	50000	Y	SingleTouch	149.99	Photography
				Gadget	200	Toy
Hitachi	Japan	30000	Y	AC	300	Appliance

First Normal Form

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

Now it's in 1NF

Company

cname	country	no_employees	for_profit
Canon	Japan	50000	Y
Hitachi	Japan	30000	Y

Products

name	price	category	manufacturer
SingleTouch	149.99	Photography	Canon
AC	300	Appliance	Hitachi
Gadget	200	Toy	Canon

CSE 414 - Autumn 2018

12

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

CSE 414 - Autumn 2018

13

Selections in SQL

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

CSE 414 - Autumn 2018

14

Demo 2

CSE 414 - Autumn 2018

15

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

Joins in SQL

pname	price	category	manufacturer	cname	country
MultiTouch	199.99	gadget	Canon	GizmoWorks	USA
SingleTouch	49.99	photography	Canon	Canon	Japan
Gizom	50	gadget	GizmoWorks	Hitachi	Japan
SuperGizmo	250.00	gadget	GizmoWorks		

Retrieve all Japanese products that cost < \$150

CSE 414 - Autumn 2018

16

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

Joins in SQL

pname	price	category	manufacturer	cname	country
MultiTouch	199.99	gadget	Canon	GizmoWorks	USA
SingleTouch	49.99	photography	Canon	Canon	Japan
Gizom	50	gadget	GizmoWorks	Hitachi	Japan
SuperGizmo	250.00	gadget	GizmoWorks		

Retrieve all Japanese products that cost < \$150

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

CSE 414 - Autumn 2018

17

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

Joins in SQL

pname	price	category	manufacturer	cname	country
MultiTouch	199.99	gadget	Canon	GizmoWorks	USA
SingleTouch	49.99	photography	Canon	Canon	Japan
Gizom	50	gadget	GizmoWorks	Hitachi	Japan
SuperGizmo	250.00	gadget	GizmoWorks		

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
```

CSE 414 - Autumn 2018

Join Predicate
Selection predicates

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

Joins in SQL

pname	price	category	manufacturer	cname	country
MultiTouch	199.99	gadget	Canon	GizmoWorks	USA
SingleTouch	49.99	photography	Canon	Canon	Japan
Gizom	50	gadget	GizmoWorks	Hitachi	Japan
SuperGizmo	250.00	gadget	GizmoWorks		

Retrieve all USA companies that manufacture "gadget" products

CSE 414 - Autumn 2018 19

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

Joins in SQL

pname	price	category	manufacturer	cname	country
MultiTouch	199.99	gadget	Canon	GizmoWorks	USA
SingleTouch	49.99	photography	Canon	Canon	Japan
Gizom	50	gadget	GizmoWorks	Hitachi	Japan
SuperGizmo	250.00	gadget	GizmoWorks		

Retrieve all USA companies that manufacture "gadget" products

Why DISTINCT?

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

CSE 414 - Autumn 2018 20

Demo 2 – cont'd

CSE 414 - Autumn 2018 21

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**

CSE 414 - Autumn 2018 22

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

Inner Join

Employee		Sales	
id	name	employeeID	productID
1	Joe	1	344
2	Jack	1	355
3	Jill	2	544

Retrieve employees and their sales

CSE 414 - Autumn 2018 23

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

Inner Join

Employee		Sales	
id	name	employeeID	productID
1	Joe	1	344
2	Jack	1	355
3	Jill	2	544

Retrieve employees and their sales

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

CSE 414 - Autumn 2018 24

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

Inner Join

Employee		Sales	
id	name	employeeID	productID
1	Joe	1	344
2	Jack	1	355
3	Jill	2	544

Retrieve employees and their sales

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

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

CSE 414 - Autumn 2018 25

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

Inner Join

Employee		Sales	
id	name	employeeID	productID
1	Joe	1	344
2	Jack	1	355
3	Jill	2	544

Retrieve employees and their sales

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

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

Jill is missing

CSE 414 - Autumn 2018 26

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

Inner Join

Employee		Sales	
id	name	employeeID	productID
1	Joe	1	344
2	Jack	1	355
3	Jill	2	544

Retrieve employees and their sales

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

Alternative syntax

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

Jill is missing

CSE 414 - Autumn 2018 27

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

Outer Join

Employee		Sales	
id	name	employeeID	productID
1	Joe	1	344
2	Jack	1	355
3	Jill	2	544

Retrieve employees and their sales

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

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

Jill is present

CSE 414 - Autumn 2018 28