

Lecture 04: SQL

Wednesday, January 11, 2006

1

Outline

- Two Examples
- Nulls (6.1.6)
- Outer joins (6.3.8)
- Database Modifications (6.5)

2

Two Examples

Store(sid, sname)
Product(pid, pname, price, sid)

Find all stores that sell *only* products with price > 100

same as:

Find all stores s.t. all their products have price > 100)

3

```
SELECT Store.name
FROM Store, Product
WHERE Store.sid = Product.sid
GROUP BY Store.sid, Store.name
HAVING 100 < min(Product.price)
```

Why both ?

Almost equivalent...

```
SELECT Store.name
FROM Store
WHERE
  100 < ALL (SELECT Product.price
            FROM product
            WHERE Store.sid = Product.sid)
```

```
SELECT Store.name
FROM Store
WHERE Store.sid NOT IN
  (SELECT Product.sid
   FROM Product
   WHERE Product.price <= 100)
```

4

Two Examples

Store(sid, sname)
Product(pid, pname, price, sid)

For each store,
find its most expensive product

5

Two Examples

This is easy but doesn't do what we want:

```
SELECT Store.sname, max(Product.price)
FROM   Store, Product
WHERE  Store.sid = Product.sid
GROUP BY Store.sid, Store.sname
```

Better:

But may
return
multiple
product names
per store

```
SELECT Store.sname, x.pname
FROM   Store, Product x
WHERE  Store.sid = x.sid and
       x.price >=
       ALL (SELECT y.price
            FROM Product y
            WHERE Store.sid = y.sid)
```

Two Examples

Finally, choose some pid arbitrarily, if there are many with highest price:

```
SELECT Store.sname, max(x.pname)
FROM Store, Product x
WHERE Store.sid = x.sid and
      x.price >=
      ALL (SELECT y.price
           FROM Product y
           WHERE Store.sid = y.sid)
GROUP BY Store.sname
```

7

NULLS in SQL

- Whenever we don't have a value, we can put a NULL
- Can mean many things:
 - Value does not exist
 - Value exists but is unknown
 - Value not applicable
 - Etc.
- The schema specifies for each attribute if can be null (*nullable* attribute) or not
- How does SQL cope with tables that have NULLs ?

8

Null Values

- If $x = \text{NULL}$ then $4 \cdot (3 - x) / 7$ is still NULL
- If $x = \text{NULL}$ then $x = \text{"Joe"}$ is UNKNOWN
- In SQL there are three boolean values:

FALSE	$=$	0
UNKNOWN	$=$	0.5
TRUE	$=$	1

9

Null Values

- $C1 \text{ AND } C2 = \min(C1, C2)$
- $C1 \text{ OR } C2 = \max(C1, C2)$
- $\text{NOT } C1 = 1 - C1$

```
SELECT *  
FROM Person  
WHERE (age < 25) AND  
      (height > 6 OR weight > 190)
```

E.g.
age=20
height=NULL
weight=200

Rule in SQL: include only tuples that yield TRUE

10

Null Values

Unexpected behavior:

```
SELECT *  
FROM Person  
WHERE age < 25 OR age >= 25
```

Some Persons are not included !

11

Null Values

Can test for NULL explicitly:

- x IS NULL
- x IS NOT NULL

```
SELECT *  
FROM Person  
WHERE age < 25 OR age >= 25 OR age IS NULL
```

Now it includes all Persons

12

Outerjoins

Explicit joins in SQL = "inner joins":

Product(name, category)
Purchase(prodName, store)

```
SELECT Product.name, Purchase.store  
FROM Product JOIN Purchase ON  
Product.name = Purchase.prodName
```

Same as:

```
SELECT Product.name, Purchase.store  
FROM Product, Purchase  
WHERE Product.name = Purchase.prodName
```

But Products that never sold will be lost !

13

Outerjoins

Left outer joins in SQL:

Product(name, category)
Purchase(prodName, store)

```
SELECT Product.name, Purchase.store  
FROM Product LEFT OUTER JOIN Purchase ON  
Product.name = Purchase.prodName
```

14

Product

Name	Category
Gizmo	gadget
Camera	Photo
OneClick	Photo

Purchase

ProdName	Store
Gizmo	Wiz
Camera	Ritz
Camera	Wiz

Name	Store
Gizmo	Wiz
Camera	Ritz
Camera	Wiz
OneClick	NULL

15

Application

Compute, for each product, the total number of sales in 'September'

Product(name, category)

Purchase(prodName, month, store)

```
SELECT Product.name, count(*)
FROM Product, Purchase
WHERE Product.name = Purchase.prodName
and Purchase.month = 'September'
GROUP BY Product.name
```

What's wrong ?

16

Application

Compute, for each product, the total number of sales in 'September'

Product(name, category)

Purchase(prodName, month, store)

```
SELECT Product.name, count(*)
FROM   Product LEFT OUTER JOIN Purchase ON
        Product.name = Purchase.prodName
        and Purchase.month = 'September'
GROUP BY Product.name
```

Now we also get the products who sold in 0 quantity

17

Outer Joins

- Left outer join:
 - Include the left tuple even if there's no match
- Right outer join:
 - Include the right tuple even if there's no match
- Full outer join:
 - Include the both left and right tuples even if there's no match

18

Modifying the Database

Three kinds of modifications

- Insertions
- Deletions
- Updates

Sometimes they are all called “updates”

19

Insertions

General form:

```
INSERT INTO R(A1,....., An) VALUES (v1,....., vn)
```

Example: Insert a new purchase to the database:

```
INSERT INTO Purchase(buyer, seller, product, store)  
VALUES ('Joe', 'Fred', 'wakeup-clock-espresso-machine',  
       'The Sharper Image')
```

Missing attribute → NULL.

May drop attribute names if give them in order.

20

Insertions

```
INSERT INTO PRODUCT(name)
SELECT DISTINCT Purchase.product
FROM Purchase
WHERE Purchase.date > "10/26/01"
```

The query replaces the VALUES keyword.
Here we insert *many* tuples into PRODUCT

21

Insertion: an Example

```
Product(name, listPrice, category)
Purchase(prodName, buyerName, price)
```

`prodName` is foreign key in `Product.name`

Suppose database got corrupted and we need to fix it:

Product

name	listPrice	category
gizmo	100	gadgets

Purchase

prodName	buyerName	price
camera	John	200
gizmo	Smith	80
camera	Smith	225

Task: insert in `Product` all `prodNames` from `Purchase`

22

Insertion: an Example

```
INSERT INTO Product(name)
SELECT DISTINCT prodName
FROM Purchase
WHERE prodName NOT IN (SELECT name FROM Product)
```

name	listPrice	category
gizmo	100	Gadgets
camera	-	-

23

Insertion: an Example

```
INSERT INTO Product(name, listPrice)
SELECT DISTINCT prodName, price
FROM Purchase
WHERE prodName NOT IN (SELECT name FROM Product)
```

name	listPrice	category
gizmo	100	Gadgets
camera	200	-
camera ??	225 ??	-

← Depends on the implementation ₂₄

Deletions

Example:

```
DELETE FROM PURCHASE
WHERE seller = 'Joe' AND
      product = 'Brooklyn Bridge'
```

Factoid about SQL: there is no way to delete only a single occurrence of a tuple that appears twice in a relation.

25

Updates

Example:

```
UPDATE PRODUCT
SET price = price/2
WHERE Product.name IN
      (SELECT product
       FROM Purchase
       WHERE Date = 'Oct, 25, 1999');
```

26