Introduction to Database Systems CSE 414

Lectures 4 and 5: Aggregates in SQL

Announcements

- Homework 1 is due on Wednesday
- Quiz 2 will be out today and due on Friday

Outline

- Outer joins (6.3.8)
- Aggregations (6.4.3 6.4.6)
- Examples, examples, examples...

Outerjoins

Product(<u>name</u>, category) Purchase(prodName, store) -- prodName is foreign key

An "inner join": SELECT Product.name, Purchase.store FROM Product, Purchase WHERE Product.name = Purchase.prodName

Same as:

SELECT Product.name, Purchase.store FROM Product JOIN Purchase ON

Product.name = Purchase.prodName

But some Products are not listed! Why?

Outerjoins

Product(<u>name</u>, category) Purchase(prodName, store) -- prodName is foreign key

If we want to include products that never sold, then we need an "outerjoin":

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

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

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 both left and right tuples even if there's no match



Comment about SQLite

- One cannot load NULL values such that they are actually loaded as null values
- So we need to use two steps:
 - Load null values using some type of special value
 - Update the special values to actual null values

```
update Purchase
  set price = null
  where price = 'null'
```

Simple Aggregations

Five basic aggregate operations in SQL

select count(*) from Purchase
select sum(quantity) from Purchase
select avg(price) from Purchase
select max(quantity) from Purchase
select min(quantity) from Purchase

Except count, all aggregations apply to a single attribute

Aggregates and NULL Values

```
Null values are not used in aggregates
   insert into Purchase
   values(12, 'gadget', NULL, NULL, 'april')
Let's try the following
    select count(*) from Purchase
    select count(quantity) from Purchase
    select sum(quantity) from Purchase
    select sum(quantity)
    from Purchase
    where quantity is not null;
                                               11
```

Counting Duplicates

COUNT applies to duplicates, unless otherwise stated:

SELECT	Count(product)
FROM	Purchase
WHERE	price > 4.99

same as Count(*)

We probably want:

SELECTCount(DISTINCT product)FROMPurchaseWHEREprice> 4.99

More Examples

SELECTSum(price * quantity)FROMPurchase

SELECTSum(price * quantity)FROMPurchaseWHEREproduct = 'bagel'



Simple Aggregations

Purchase	Product	Price	Quantity
	Bagel	3	20
	Bagel	1.50	20
	Banana	0.5	50
	Banana	2	10
	Banana	4	10
SELECT S FROM P WHERE p	90		

CSE 414 - Spring 2013

(= 60+30)

Grouping and Aggregation

Purchase(product, price, quantity)

Find total quantities for all sales over \$1, by product.

SELECT	product, Sum(quantity) AS TotalSales
FROM	Purchase
WHERE	price > 1
GROUP BY	product

Let's see what this means...

Grouping and Aggregation

- 1. Compute the FROM and WHERE clauses.
- 2. Group by the attributes in the GROUPBY
- 3. Compute the SELECT clause: grouped attributes and aggregates.

1&2. FROM-WHERE-GROUPBY



3. SELECT

Product	Price	Quantity										
Bagel	3	20	ς.	Product	TotalSales							
Bagel	1.50	20		Bagel	40							
Banana	0.5	50									Banana	20
Banana	2	10										
Banana	4	10										

SELECTproduct, Sum(quantity) AS TotalSalesFROMPurchaseWHEREprice > 1GROUP BYproduct



Need to be Careful...

SELECT product, max(quantity) FROM Purchase GROUP BY product		Product	Price	Quantity
		Bagel	3	20
		Bagel	1.50	20
SELECT product, quantity FROM Purchase		Banana	0.5	50
GROUP BY product		Banana	2	10
		Banana	4	10
sqlite is WRONG on this query.		Advanced DI Server) giv	BMS (e.g. So ves an error	QL

Ordering Results

SELECT product, sum(price*quantity) as rev FROM purchase GROUP BY product ORDER BY rev desc

HAVING Clause

Same query as earlier, except that we consider only products that had at least 30 sales.

SELECT	product, sum(price*quantity)		
FROM	Purchase		
WHERE	price > 1		
GROUP BY product			
HAVING	Sum(quantity) > 30		

HAVING clause contains conditions on aggregates.

WHERE vs HAVING

- WHERE condition is applied to individual rows
 - The rows may or may not contribute to the aggregate
 - No aggregates allowed here
- HAVING condition is applied to the entire group
 - Entire group is returned, or not al all
 - May use aggregate functions in the group

Aggregates and Joins

```
create table Product
 (pid int primary key,
    pname varchar(15),
    manufacturer varchar(15));
```

```
insert into product values(1, 'bagel', 'Sunshine Co.');
insert into product values(2, 'banana', 'BusyHands');
insert into product values(3, 'gizmo', 'GizmoWorks');
insert into product values(4, 'gadget', 'BusyHands');
insert into product values(5, 'powerGizmo', 'PowerWorks');
```

Aggregate + Join Example

SELECT x.manufacturer, count(*) FROM Product x, Purchase y WHERE x.pname = y.product GROUP BY x.manufacturer

> SELECT x.manufacturer, y.month, count(*) FROM Product x, Purchase y WHERE x.pname = y.product GROUP BY x.manufacturer, y.month

What do these

query mean?

General form of Grouping and Aggregation

SELECT	S
FROM	R_1, \ldots, R_n
WHERE	C1
GROUP BY	a ₁ ,,a _k
HAVING	C2

S = may contain attributes a₁,...,a_k and/or any aggregates but NO OTHER ATTRIBUTES
C1 = is any condition on the attributes in R₁,...,R_n
C2 = is any condition on aggregate expressions and on attributes a₁,...,a_k

CSE 414 - Spring 2013

Why?

Semantics of SQL With Group-By



Evaluation steps:

- 1. Evaluate FROM-WHERE using Nested Loop Semantics
- 2. Group by the attributes a_1, \ldots, a_k
- 3. Apply condition C2 to each group (may have aggregates)
- 4. Compute aggregates in S and return the result

Empty Groups

- In the result of a group by query, there is one row per group in the result
- No group can be empty!
- In particular, count(*) is never 0

SELECT x.manufacturer, count(*) FROM Product x, Purchase y WHERE x.pname = y.product GROUP BY x.manufacturer What if there are no purchases for a manufacturer

Empty Groups: Example



Empty Group Problem

SELECT x.manufacturer, count(*) FROM Product x, Purchase y WHERE x.pname = y.product GROUP BY x.manufacturer What if there are no purchases for a manufacturer

Empty Group Solution: Outer Join

SELECT x.manufacturer, count(y.pid) FROM Product x LEFT OUTER JOIN Purchase y ON x.pname = y.product GROUP BY x.manufacturer