

Introduction to Data Management CSE 344

Lecture 7: SQL Wrap-up Relational Algebra

CSE 344 - Winter 2016

1

Announcements

- Webquiz 3 is open, due on Sunday
- Homework 3 is posted, due on Tuesday, 2/2
 - We are using Microsoft Azure Cloud services!
 - Use the promotion code you received

CSE 344 - Winter 2016

2

What We Learned Last Time

- Subqueries can occur in every clause:
 - SELECT
 - FROM
 - WHERE
- Monotone queries: SELECT-FROM-WHERE
 - Existential quantifier
- Non-monotone queries
 - Universal quantifier
 - Aggregation

CSE 344 - Winter 2016

3

Practice these queries in SQL

Likes(drinker, beer)
Frequents(drinker, bar)
Serves(bar, beer)

Ullman's drinkers-bars-beers example

Find drinkers that frequent some bar that serves some beer they like.

x: $\exists y. \exists z. \text{Frequents}(x, y) \wedge \text{Serves}(y, z) \wedge \text{Likes}(x, z)$

Find drinkers that frequent only bars that serves some beer they like.

x: $\forall y. \text{Frequents}(x, y) \Rightarrow (\exists z. \text{Serves}(y, z) \wedge \text{Likes}(x, z))$

Find drinkers that frequent some bar that serves only beers they like.

x: $\exists y. \text{Frequents}(x, y) \wedge \forall z. (\text{Serves}(y, z) \Rightarrow \text{Likes}(x, z))$

Find drinkers that frequent only bars that serves only beer they like.

x: $\forall y. \text{Frequents}(x, y) \Rightarrow \forall z. (\text{Serves}(y, z) \Rightarrow \text{Likes}(x, z))$ ⁴

Likes(drinker, beer)
Frequents(drinker, bar)
Serves(bar, beer)

Example 1

Find drinkers that frequent some bar that serves some beer they like.

```
SELECT DISTINCT X.drinker
FROM Frequents X, Serves Y, Likes Z
WHERE X.bar = Y.bar
AND Y.beer = Z.beer
AND X.drinker = Z.drinker
```

5

Likes(drinker, beer)
Frequents(drinker, bar)
Serves(bar, beer)

Example 2

Find drinkers that frequent some bar that serves only beers they don't like

```
SELECT DISTINCT Y.drinker
FROM Frequents Y
WHERE NOT EXISTS (SELECT *
FROM Serves Z, Likes U
WHERE Y.bar=Z.bar
AND Y.drinker=U.drinker
AND Z.beer = U.beer)
```

6

Likes(drinker, beer)
Frequents(drinker, bar)
Serves(bar, beer)

Example 3

Find drinkers that frequent only bars that serves some beer they like.

(Recall: In example 2, we found drinkers that frequent some bar that serves only beers they don't like)

```
SELECT X.drinker
FROM Frequents X
WHERE X.drinker
NOT IN (SELECT Y.drinker
        FROM Frequents Y
        WHERE NOT EXISTS ( SELECT *
                          FROM Serves Z, Likes U
                          WHERE Y.bar=Z.bar
                          AND Y.drinker=U.drinker
                          AND Z.beer = U.beer))
```

7

Product (pname, price, cid)
Company(cid, cname, city)

Unnesting Aggregates

Find the number of companies in each city

```
SELECT DISTINCT X.city, (SELECT count(*)
                        FROM Company Y
                        WHERE X.city = Y.city)
FROM Company X
```

```
SELECT city, count(*)
FROM Company
GROUP BY city
```

Equivalent queries

Note: no need for **DISTINCT**
(**DISTINCT** *is the same* as **GROUP BY**)

CSE 344 - Winter 2016

8

Product (pname, price, cid)
Company(cid, cname, city)

Unnesting Aggregates

Find the number of products made in each city

```
SELECT DISTINCT X.city, (SELECT count(*)
                        FROM Product Y, Company Z
                        WHERE Z.cid=Y.cid
                        AND Z.city = X.city)
FROM Company X
```

```
SELECT X.city, count(*)
FROM Company X, Product Y
WHERE X.cid=Y.cid
GROUP BY X.city
```

NOT equivalent !
You should know why!

CSE 344 - Winter 2016

9

Purchase(pid, product, quantity, price)

GROUP BY v.s. Nested Queries

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

```
SELECT DISTINCT x.product, (SELECT Sum(y.quantity)
                        FROM Purchase y
                        WHERE x.product = y.product
                        AND y.price > 1)
                        AS TotalSales
FROM Purchase x
WHERE x.price > 1
```

Why twice ?

CSE 344 - Winter 2016

10

Author(login, name)
Wrote(login, url)

More Unnesting

Find authors who wrote ≥ 10 documents:

CSE 344 - Winter 2016

11

Author(login, name)
Wrote(login, url)

More Unnesting

Find authors who wrote ≥ 10 documents:

Attempt 1: with nested queries

```
SELECT DISTINCT Author.name
FROM Author
WHERE (SELECT count(Wrote.url)
      FROM Wrote
      WHERE Author.login=Wrote.login)
      >= 10
```

This is
SQL by
a novice

CSE 344 - Winter 2016

12

Author(login, name)
Wrote(login, url)

More Unnesting

Find authors who wrote ≥ 10 documents:

Attempt 1: with nested queries

Attempt 2: using GROUP BY and HAVING

```
SELECT Author.name
FROM Author, Wrote
WHERE Author.login=Wrote.login
GROUP BY Author.name
HAVING count(wrote.url) >= 10
```

This is
SQL by
an expert

CSE 344 - Winter 2016

13

Product(pname, price, cid)
Company(cid, cname, city)

Finding Witnesses

For each city, find the most expensive product made in that city

CSE 344 - Winter 2016

14

Product(pname, price, cid)
Company(cid, cname, city)

Finding Witnesses

For each city, find the most expensive product made in that city

Finding the maximum price is easy...

```
SELECT x.city, max(y.price)
FROM Company x, Product y
WHERE x.cid = y.cid
GROUP BY x.city;
```

But we need the *witnesses*, i.e. the products with max price

CSE 344 - Winter 2016

15

Product(pname, price, cid)
Company(cid, cname, city)

Finding Witnesses

To find the witnesses, compute the maximum price in a subquery

```
SELECT DISTINCT u.city, v.pname, v.price
FROM Company u, Product v,
(SELECT x.city, max(y.price) as maxprice
FROM Company x, Product y
WHERE x.cid = y.cid
GROUP BY x.city) w
WHERE u.cid = v.cid
and u.city = w.city
and v.price=w.maxprice;
```

CSE 344 - Winter 2016

16

Product(pname, price, cid)
Company(cid, cname, city)

Finding Witnesses

Or we can use a subquery in where clause

```
SELECT u.city, v.pname, v.price
FROM Company u, Product v
WHERE u.cid = v.cid
and v.price >= ALL (SELECT y.price
FROM Company x, Product y
WHERE u.city=x.city
and x.cid=y.cid);
```

CSE 344 - Winter 2016

17

Product(pname, price, cid)
Company(cid, cname, city)

Finding Witnesses

There is a more concise solution here:

```
SELECT u.city, v.pname, v.price
FROM Company u, Product v, Company x, Product y
WHERE u.cid = v.cid and u.city = x.city and x.cid = y.cid
GROUP BY u.city, v.pname, v.price
HAVING v.price = max(y.price);
```

CSE 344 - Winter 2016

18

Where We Are

- Motivation for using a DBMS for managing data
- SQL, SQL, SQL
 - Declaring the schema for our data (CREATE TABLE)
 - Inserting data one row at a time or in bulk (INSERT/import)
 - Modifying the schema and updating the data (ALTER/UPDATE)
 - Querying the data (SELECT)
- Next step: More knowledge of how DBMSs work
 - Client-server architecture
 - Relational algebra, query execution, and physical tuning

CSE 344 - Winter 2016

19

Where We Are

- Motivation for using a DBMS for managing data
- SQL, SQL, SQL
 - Declaring the schema for our data (CREATE TABLE)
 - Inserting data one row at a time or in bulk (INSERT/import)
 - Modifying the schema and updating the data (ALTER/UPDATE)
 - Querying the data (SELECT)
- Next step: More knowledge of how DBMSs work
 - Client-server architecture
 - Relational algebra and query execution

CSE 344 - Winter 2016

20

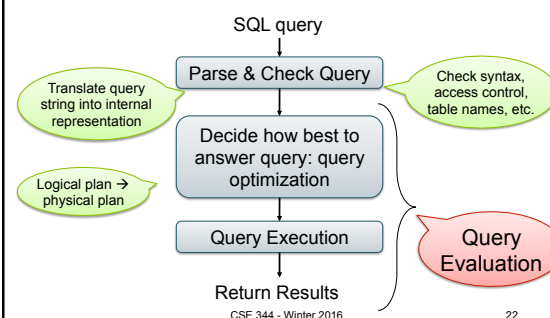
Where We Are

- Motivation for using a DBMS for managing data
- SQL, SQL, SQL
 - Declaring the schema for our data (CREATE TABLE)
 - Inserting data one row at a time or in bulk (INSERT/import)
 - Modifying the schema and updating the data (ALTER/UPDATE)
 - Querying the data (SELECT)
- Next step: More knowledge of how DBMSs work
 - Client-server architecture
 - Relational algebra and query execution

CSE 344 - Winter 2016

21

Query Evaluation Steps



CSE 344 - Winter 2016

22

The WHAT and the HOW

- SQL = **WHAT** we want to get from the data
- Relational Algebra = **HOW** to get the data we want
- The passage from **WHAT** to **HOW** is called **query optimization**
 - SQL -> Relational Algebra -> Physical Plan
 - Relational Algebra = Logical Plan

CSE 344 - Winter 2016

23

Overview: SQL = WHAT

Product(pid, name, price)
 Purchase(pid, cid, store)
 Customer(cid, name, city)

```

SELECT DISTINCT x.name, z.name
FROM Product x, Purchase y, Customer z
WHERE x.pid = y.pid and y.cid = z.cid and
      x.price > 100 and z.city = 'Seattle'
  
```

It's clear **WHAT** we want, unclear **HOW** to get it

CSE 344 - Winter 2016

24

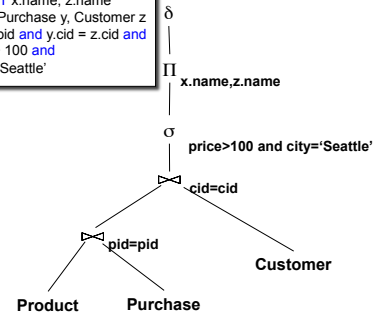
Overview: Relational Algebra = HOW

```
SELECT DISTINCT x.name, z.name
FROM Product x, Purchase y, Customer z
WHERE x.pid = y.pid and y.cid = z.cid and
      x.price > 100 and
      z.city = 'Seattle'
```

25

Overview: Relational Algebra = HOW

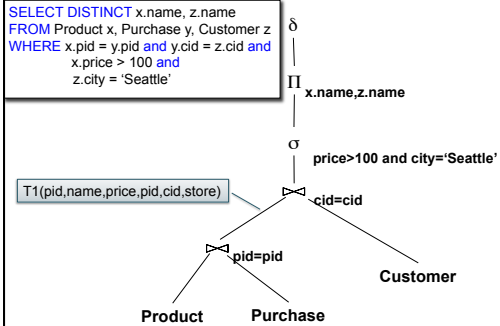
```
SELECT DISTINCT x.name, z.name
FROM Product x, Purchase y, Customer z
WHERE x.pid = y.pid and y.cid = z.cid and
      x.price > 100 and
      z.city = 'Seattle'
```



26

Overview: Relational Algebra = HOW

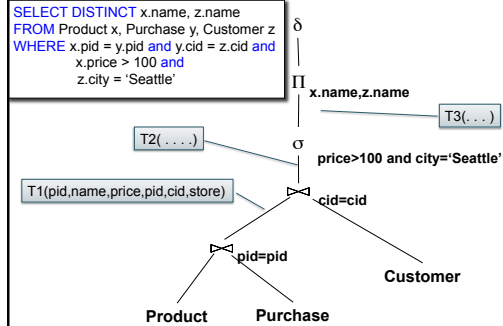
```
SELECT DISTINCT x.name, z.name
FROM Product x, Purchase y, Customer z
WHERE x.pid = y.pid and y.cid = z.cid and
      x.price > 100 and
      z.city = 'Seattle'
```



27

Overview: Relational Algebra = HOW

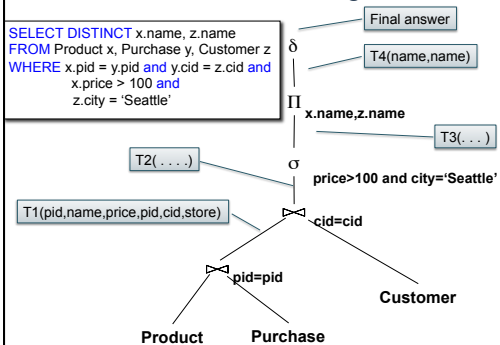
```
SELECT DISTINCT x.name, z.name
FROM Product x, Purchase y, Customer z
WHERE x.pid = y.pid and y.cid = z.cid and
      x.price > 100 and
      z.city = 'Seattle'
```



28

Overview: Relational Algebra = HOW

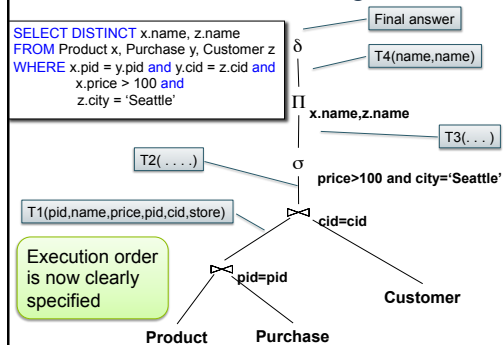
```
SELECT DISTINCT x.name, z.name
FROM Product x, Purchase y, Customer z
WHERE x.pid = y.pid and y.cid = z.cid and
      x.price > 100 and
      z.city = 'Seattle'
```



29

Overview: Relational Algebra = HOW

```
SELECT DISTINCT x.name, z.name
FROM Product x, Purchase y, Customer z
WHERE x.pid = y.pid and y.cid = z.cid and
      x.price > 100 and
      z.city = 'Seattle'
```



30

