Principles of Database Systems
CSE 544p

Lecture #1

September 29, 2010

Staff

- Instructor: Dan Suciu
  - CSE 662, suciu@cs.washington.edu
    Office hours: by appointment

- TAs:
  - Jessica Leung joyleung@cs.washington.edu
  - TA: Paramjit Singh Sandhu paramsan@cs.washington.edu

- Technical support:
  - At UW: Fred Videon, fred@cs.washington.edu
  - At MS: Matt McGinley

Communications

- Web page: http://www.cs.washington.edu/p544
  - Lectures will be available here
  - Homework will be posted here
  - Announcements may be posted here
- Mailing list:
  - Announcements, group discussions
  - Please subscribe

Textbook(s)

Main textbook:

- *Database Management Systems*, Ramakrishnan and Gehrke

Second textbook:

- *Database Systems: The Complete Book*, Garcia-Molina, Ullman, Widom

Course Format

- Lectures Wednesdays, 6:30-9:20
- 7 Homework Assignments
- Take-home Final

Grading

- Homework: 70 %
- Take-home Final: 30%

Homework Assignments

1. SQL
2. Conceptual design
3. JAVA/SQL
4. Transactions
5. Database tuning
6. XML/XPath/XQuery
7. Pig Latin, on AWS

Due: Tuesdays', by 11:59pm. Three late days per person

Take-home Final

- Posted on December 8, at 11:59pm
- Due on December 9, by 11:59pm
- No late days/hours/minutes/seconds

December 9th ***is*** the day of your final

Software Tools

- SQL Server 2008
  - You have access to http://msdnaa.cs.washington.edu
  - Username is full @cs.washington.edu email address
  - Doesn't work ? Email ms-sw-admin@cs.washington.edu
  - Download the client, connect to IPROJSRV (may need tunneling)
  - OK to use you own server, just import IMDB (may need tunneling)
- Postgres: download from
  - download http://www.postgresql.org/download/
  - Is also installed on lab machines
- Xquery: download one interpreter from
  - Zorba: http://www.zorba-xquery.com/ (I use this one: ½ day installation)
  - Galax: http://galax.sourceforge.net/ (great in the past, seems less well maintained)
  - Saxon: http://saxon.sourceforge.net/ (from apache; very popular)
- Pig Latin: download from
  - http://hadoop.apache.org/pig/
  - We will also run it on Amazon Web Services

Rest of Today's Lecture

- Overview of DBMS
- SQL

Database

What is a database ?

Give examples of databases

Database

What is a database ?

- A collection of files storing related data

Give examples of databases

- Accounts database; payroll database; UW's students database; Amazon's products database; airline reservation database

Database Management System

What is a DBMS ?

Give examples of DBMS

Database Management System

What is a DBMS ?

- *A big C program written by someone else that allows us to manage efficiently a large database and allows it to persist over long periods of time*

Give examples of DBMS

- DB2 (IBM), SQL Server (MS), Oracle, Sybase
- MySQL, Postgres, …

SQL for Nerds, Greenspun, http://philip.greenspun.com/sql/ (Chap 1,2)

Market Shares

From 2006 Gartner report:

- IBM: 21% market with $3.2BN in sales
- Oracle: 47% market with $7.1BN in sales
- Microsoft: 17% market with $2.6BN in sales

An Example

The Internet Movie Database
http://www.imdb.com

- Entities:
  Actors (800k), Movies (400k), Directors, …

- Relationships:
  who played where, who directed what, …

Tables

Actor:

id	fName	lName	gender
195428	Tom	Hanks	M
645947	Amy	Hanks	F
. . .			

Cast:

pid	mid
195428	337166
. . .	

Movie:

id	Name	year
337166	Toy Story	1995
. . .	. . .	. ..

SQL

```
SELECT *
FROM  Actor
```

SQL

```
SELECT count(*)
FROM  Actor
```

This is an *aggregate query*

SQL

```
SELECT *
FROM  Actor
WHERE lName = 'Hanks'
```

This is a *selection query*

SQL

```
SELECT *
FROM  Actor, Casts, Movie
WHERE lname='Hanks' and Actor.id = Casts.pid
  and Casts.mid=Movie.id and Movie.year=1995
```

This query has *selections* and *joins*

817k actors, 3.5M casts, 380k movies;

How can it be so fast ?

How Can We Evaluate the Query ?

Actor:

id	**fName e**	**lName**	**gende r**
. . .		Hanks	
. . .			

Cast:

pid	**mid**
. . .	
. . .	

Movie:

id	**Name**	**year**
. . .		1995
. . .		

Plan 1: . . . . [in class]

Plan 2: . . . . [in class]

Evaluating Tom Hanks

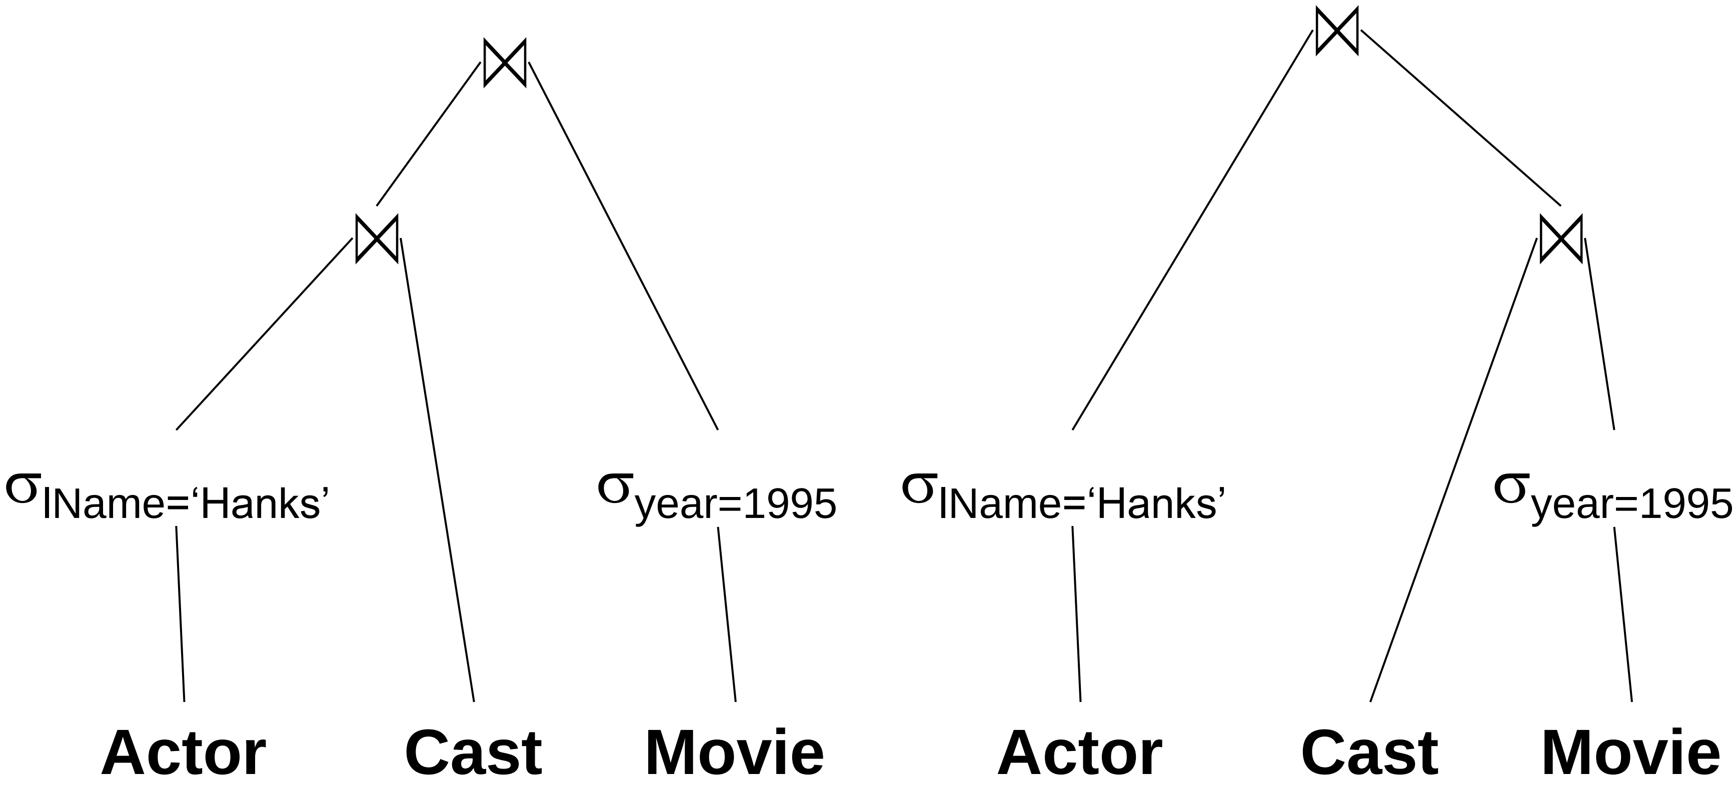

Optimization and Query Execution

- Indexes: on Actor.lName, on Movie.year
- Query optimization
  - Access path selection
  - Join order
- Statistics
- Multiple implementations of joins

Recovery

- Transfer $100 from account #4662 to #7199:

```
X = Read(Account_1);
X.amount = X.amount - 100;
Write(Account_1, X);

Y = Read(Account_2);
Y.amount = Y.amount + 100;
Write(Account_2, Y);
```

Recovery

- Transfer $100 from account #4662 to #7199:

```
X = Read(Account_1);
X.amount = X.amount - 100;
Write(Account_1, X);

Y = Read(Account_2);
Y.amount = Y.amount + 100;
Write(Account_2, Y);
```

CRASH !

What is the problem ?

Concurrency Control

- How to overdraft your account:

User 1

User 2

```
X = Read(Account);
if (X.amount > 100)
   { dispense_money( );
      X.amount = X.amount – 100;
    }
else error(“Insufficient funds”);
```

```
X = Read(Account);
if (X.amount > 100)
   { dispense_money( );
      X.amount = X.amount – 100;
    }
else error(“Insufficient funds”);
```

What can go wrong ?

Transactions

- Recovery

- Concurrency control

ACID =
- Atomicity (= recovery)
- Consistency
- Isolation   (= concurrency control)
- Durability

Client/Server Database Architecture

- There is one single *server* that stores the database (called DBMS or RDBMS):
  - Usually a beefed-up system, e.g. IPROJSRV
  - But can be your own desktop…
  - … or a huge cluster running a parallel dbms
- Many *clients* running apps and connecting to DBMS
  - E.g. Microsoft's Management Studio
  - Or psql (for postgres)
  - Always: some else's big Java or C++ program
- The client "talks" to the server using JDBC protocol

Types of Usages for Databases

- OLTP (online-transaction-processing)
  – Many updates
  – Many “point queries”: retrieve the record with a given key.

- Decision-Support
  – Many aggregate/group-by queries.
  – Sometimes called *data warehouse*

SQL v.s. noSQL

- Reading for next time:
  SQL Databases v. NoSQL Databases, by Mike Stonebraker, CACM 53(4), 2010.

Data Management

- Data Management is more than databases !

Here is an example of a problem:

- Alice sends Bob in random order all the numbers 1, 2, 3, …, 100000000000000000000
- She does not repeat any number
- But she misses *exactly one*
- Help Bob find out which one is missing !

After you solve it, make it a bit harder:

- Alice misses *exactly ten* numbers
- Help Bob find out which ones are missing !

Accessing SQL Server

- SQL Server Management Studio
- Server Type = Database Engine
- Server Name = IPROJSRV
- Authentication = SQL Server Authentication
  - Login = your UW email address (not the CSE email)
  - Password = [in class]
- Must connect from within CSE, or must use tunneling
- Alternatively: install your own, get it from MSDNAA (see earlier slide)
- Then play with IMDB, start working on HW 1

Outline for Today

- Basics: we go quickly or skip slides, please read the slides at home
  - Datatypes in SQL
  - Simple Queries in SQL
  - Joins
- Subqueries: this is tough ! Please read the relational calculus and tuple calculus in the textbook (Chapter 4.3)
- Aggregates: separates pros from amateurs
- Nulls, Outer joins

SQL

- Data Definition Language (DDL)
  - Create/alter/delete tables and their attributes
  - Following lectures...
- Data Manipulation Language (DML)
  - Query one or more tables – discussed next !
  - Insert/delete/modify tuples in tables

Table name

Tables in SQL

Attribute names

Key

Product

PName	Price	Category	Manufacturer
Gizmo	$19.99	Gadgets	GizmoWorks
Powergizmo	$29.99	Gadgets	GizmoWorks
SingleTouch	$149.99	Photography	Canon
MultiTouch	$203.99	Household	Hitachi

Tuples or rows

Data Types in SQL

- Atomic types:
  - Characters: CHAR(20), VARCHAR(50)
  - Numbers: INT, BIGINT, SMALLINT, FLOAT
  - Others: MONEY, DATETIME, …
- Record (aka tuple)
  - Has atomic attributes
- Table (relation)
  - A set of tuples

Simple SQL Query

Product

PName	Price	Category	Manufacturer
Gizmo	$19.99	Gadgets	GizmoWorks
Powergizmo	$29.99	Gadgets	GizmoWorks
SingleTouch	$149.99	Photography	Canon
MultiTouch	$203.99	Household	Hitachi

```
SELECT   *
FROM     Product
WHERE    category='Gadgets'
```

PName	Price	Category	Manufacturer
Gizmo	$19.99	Gadgets	GizmoWorks
Powergizmo	$29.99	Gadgets	GizmoWorks

"selection"

Simple SQL Query

Product

PName	Price	Category	Manufacturer
Gizmo	$19.99	Gadgets	GizmoWorks
Powergizmo	$29.99	Gadgets	GizmoWorks
SingleTouch	$149.99	Photography	Canon
MultiTouch	$203.99	Household	Hitachi

```
SELECT   PName, Price, Manufacturer
FROM     Product
WHERE    Price > 100
```

PName	Price	Manufacturer
SingleTouch	$149.99	Canon
MultiTouch	$203.99	Hitachi

"selection" and "projection"

Details

- Case insensitive:

  SELECT = Select = select

  Product = product

  BUT: ‘Seattle’ ≠ ‘seattle’

- Constants:

  ‘abc’ - yes

  “abc” - no

Eliminating Duplicates

```
SELECT   DISTINCT category
FROM     Product
```

Category
Gadgets
Photography
Household

Compare to:

```
SELECT   category
FROM     Product
```

Category
Gadgets
Gadgets
Photography
Household

Ordering the Results

```
SELECT   pname, price, manufacturer
FROM     Product
WHERE    category='gizmo' AND price > 50
ORDER BY price, pname
```

Ties are broken by the second attribute on the ORDER BY list.

Ordering is ascending, unless you specify the DESC keyword.

PName	Price	Category	Manufacturer
Gizmo	$19.99	Gadgets	GizmoWorks
Powergizmo	$29.99	Gadgets	GizmoWorks
SingleTouch	$149.99	Photography	Canon
MultiTouch	$203.99	Household	Hitachi

```
SELECT   DISTINCT category
FROM     Product
ORDER BY category
```

⇨ ?

```
SELECT   Category
FROM     Product
ORDER BY PName
```

⇨ ?

```
SELECT   DISTINCT category
FROM     Product
ORDER BY PName
```

⇨ ?

Keys and Foreign Keys

Company

CName	StockPrice	Country
GizmoWorks	25	USA
Canon	65	Japan
Hitachi	15	Japan

Key

Product

PName	Price	Category	Manufacturer
Gizmo	$19.99	Gadgets	GizmoWorks
Powergizmo	$29.99	Gadgets	GizmoWorks
SingleTouch	$149.99	Photography	Canon
MultiTouch	$203.99	Household	Hitachi

Foreign key

Joins

Product (<u>PName</u>, Price, Category, Manufacturer)

Company (<u>CName</u>, stockPrice, Country)

Find all products under $200 manufactured in Japan;
return their names and prices.

```
SELECT   PName, Price
FROM     Product, Company
WHERE    Manufacturer=CName AND Country='Japan'
         AND Price <= 200
```

Join between Product and Company

Joins

Product

PName	Price	Category	Manufacturer
Gizmo	$19.99	Gadgets	GizmoWorks
Powergizmo	$29.99	Gadgets	GizmoWorks
SingleTouch	$149.99	Photography	Canon
MultiTouch	$203.99	Household	Hitachi

Company

Cname	StockPrice	Country
GizmoWorks	25	USA
Canon	65	Japan
Hitachi	15	Japan

```
SELECT   PName, Price
FROM     Product, Company
WHERE    Manufacturer=CName AND Country='Japan'
         AND Price <= 200
```

PName	Price
SingleTouch	$149.99

Tuple Variables

Person(pname, address, worksfor)
Company(cname, address)

Which address ?

```
SELECT   DISTINCT pname, address
FROM     Person, Company
WHERE    worksfor = cname
```

```
SELECT   DISTINCT Person.pname, Company.address
FROM     Person, Company
WHERE    Person.worksfor = Company.cname
```

```
SELECT   DISTINCT x.pname, y.address
FROM     Person AS x, Company AS y
WHERE    x.worksfor = y.cname
```

In Class

Product (pname, price, category, manufacturer)

Company (cname, stockPrice, country)

Find all Chinese companies that manufacture products both in the 'toy' category

```
SELECT   cname

FROM

WHERE
```

In Class

Product (pname, price, category, manufacturer)

Company (cname, stockPrice, country)

Find all Chinese companies that manufacture products both in the ‘electronic’ and ‘toy’ categories

```
SELECT   cname

FROM

WHERE
```

Meaning (Semantics) of SQL Queries

```
SELECT a1, a2, …, ak
FROM   R1 AS x1, R2 AS x2, …, Rn AS xn
WHERE  Conditions
```

$\text{SELECT } a_1, a_2, \ldots, a_k$
$\text{FROM } R_1 \text{ AS } x_1, R_2 \text{ AS } x_2, \ldots, R_n \text{ AS } x_n$
WHERE Conditions

Answer = {}
for x_1 **in** R_1 **do**
    for x_2 **in** R_2 **do**
        …..
            for x_n **in** R_n **do**
                if Conditions
                    then Answer = Answer $\cup$ $\{(a_1,\ldots,a_k)\}$
return Answer

Using the Formal Semantics

What do these queries compute ?

```
SELECT  DISTINCT R.A
FROM   R, S
WHERE  R.A=S.A
```

Returns $R \cap S$

```
SELECT  DISTINCT R.A
FROM   R, S, T
WHERE  R.A=S.A   OR   R.A=T.A
```

If $S \neq \emptyset$ and $T \neq \emptyset$
then returns $R \cap (S \cup T)$
else returns $\emptyset$

Joins Introduce Duplicates

Product (pname, price, category, manufacturer)

Company (cname, stockPrice, country)

Find all countries that manufacture some product in the ‘Gadgets’ category.

```
SELECT   Country
FROM     Product, Company
WHERE    Manufacturer=CName AND Category=‘Gadgets’
```

Joins Introduce Duplicates

Product

Name	Price	Category	Manufacturer
Gizmo	$19.99	Gadgets	GizmoWorks
Powergizmo	$29.99	Gadgets	GizmoWorks
SingleTouch	$149.99	Photography	Canon
MultiTouch	$203.99	Household	Hitachi

Company

Cname	StockPrice	Country
GizmoWorks	25	USA
Canon	65	Japan
Hitachi	15	Japan

```
SELECT   Country
FROM     Product, Company
WHERE    Manufacturer=CName AND Category='Gadgets'
```

Country
USA
USA

Duplicates !
Remember to add DISTINCT

Subqueries

- A *subquery* is another SQL query nested inside a larger query
- Such inner-outer queries are called *nested queries*
- A subquery may occur in:
  1. A SELECT clause
  2. A FROM clause
  3. A WHERE clause

Rule of thumb: avoid writing nested queries when possible; keep in mind that sometimes it's impossible

1. Subqueries in SELECT

Product (<u>pname</u>, price, company)

Company(<u>cname</u>, city)

For each product return the city where it is manufactured

```
SELECT X.pname, (SELECT Y.city
                 FROM Company Y
                 WHERE Y.cname=X.company)
FROM  Product X
```

What happens if the subquery returns more than one city ?

1. Subqueries in SELECT

Product (pname, price, company)

Company(cname, city)

Whenever possible, don't use a nested queries:

```
SELECT pname, (SELECT city FROM Company WHERE cname=company)
FROM  Product
```

=

```
SELECT pname, city
FROM   Product, Company
WHERE cname=company
```

We have "unnested" the query

1. Subqueries in SELECT

Product (pname, price, company)

Company(cname, city)

Compute the number of products made in each city

```
SELECT DISTINCT city, (SELECT count(*)
                       FROM Product
                       WHERE cname=company)
FROM  Company
```

Better: we can unnest by using a GROUP BY (next lecture)

2. Subqueries in FROM

Product (<u>pname</u>, price, company)

Company(<u>cname</u>, city)

Find all products whose prices is > 20 and < 30

```
SELECT X.city
FROM (SELECT * FROM Product AS Y WHERE Y.price > 20) AS  X
WHERE X.price < 30
```

Unnest this query !

3. Subqueries in WHERE

Product (pname, price, company)

Company(cname, city)

Existential quantifiers

Find all cities that make some products with price < 100

Using EXISTS:

```
SELECT DISTINCT  Company.city
FROM     Company
WHERE  EXISTS (SELECT *
                           FROM Product
                          WHERE company = cname and Produc.price < 100)
```

3. Subqueries in WHERE

Product (<u>pname</u>, price, company)

Company(<u>cname</u>, city)

Existential quantifiers

Find all cities that make <u>some</u> products with price < 100

Predicate Calculus (a.k.a. First Order Logic)

$$\{ y \mid \exists x.\ \mathrm{Company}(x,y) \wedge (\exists z.\ \exists p.\ \mathrm{Product}(z,p,x) \wedge p < 100) \}$$

3. Subqueries in WHERE

Product (pname, price, company)

Company(cname, city)

Existential quantifiers

Find all cities that make some products with price < 100

Using IN

```
SELECT DISTINCT  Company.city
FROM     Company
WHERE Company.cname IN (SELECT Product.company
                        FROM Product
                        WHERE Produc.price < 100)
```

3. Subqueries in WHERE

Product (pname, price, company)

Company(cname, city)

Existential quantifiers

Find all cities that make some products with price < 100

Using ANY:

```
SELECT DISTINCT  Company.city
FROM     Company
WHERE 100 > ANY (SELECT price
                        FROM Product
                        WHERE company = cname)
```

3. Subqueries in WHERE

Product (<u>pname</u>, price, company)

Existential quantifiers

Company(<u>cname</u>, city)

Find all cities that make <u>some</u> products with price < 100

Now let's unnest it:

```
SELECT DISTINCT  Company.cname
FROM     Company, Product
WHERE  Company.cname = Product.company and Product.price < 100
```

Existential quantifiers are easy ! ☺

3. Subqueries in WHERE

Universal quantifiers

Product (<u>pname</u>, price, company)

Company(<u>cname</u>, city)

Find all cities with companies
that make <u>only</u> products with price < 100

Universal quantifiers are hard ! ☹

3. Subqueries in WHERE

Universal quantifiers

Product (pname, price, company)

Company(cname, city)

Find all cities with companies
that make <u>only</u> products with price < 100

Predicate Calculus (a.k.a. First Order Logic)

$$\{ y \mid \exists x.\ \mathrm{Company}(x,y) \land (\forall z.\ \forall p.\ \mathrm{Product}(z,p,x) \rightarrow p < 100) \}$$

3. Subqueries in WHERE

De Morgan's Laws:

$$\neg(A \wedge B) = \neg A \vee \neg B$$

$$\neg(A \vee B) = \neg A \wedge \neg B$$

$$\neg\forall x.\ P(x) = \exists x.\ \neg\ P(x)$$

$$\neg\exists x.\ P(x) = \forall x.\ \neg\ P(x)$$

$$\neg(A \Rightarrow B) = A \wedge \neg B$$

$$\{\ y \mid \exists x.\ \text{Company}(x,y)\ \wedge\ (\forall z.\ \forall p.\ \text{Product}(z,p,x) \Rightarrow p < 100)\ \}$$

$$=$$

$$\{\ y \mid \exists x.\ \text{Company}(x,y)\ \wedge\ \neg\ (\exists z \exists p.\ \text{Product}(z,p,x)\ \wedge\ p \geq 100)\ \}$$

$$=$$

$$\{\ y \mid \exists x.\ \text{Company}(x,y))\ \}\ \ -$$
$$\{\ y \mid \exists x.\ \text{Company}(x,y)\ \wedge\ (\exists z \exists p.\ \text{Product}(z,p,x)\ \wedge\ p \geq 100)\ \}$$

3. Subqueries in WHERE

1. Find *the other* companies: i.e. s.t. some product ≥ 100

```
SELECT DISTINCT  Company.city
FROM     Company
WHERE  Company.cname IN (SELECT Product.company
                          FROM Product
                          WHERE Produc.price >= 100
```

2. Find all companies s.t. all their products have price < 100

```
SELECT DISTINCT  Company.city
FROM     Company
WHERE  Company.cname NOT IN (SELECT Product.company
                              FROM Product
                              WHERE Produc.price >= 100
```

3. Subqueries in WHERE

Product (pname, price, company)

Universal quantifiers

Company(cname, city)

Find all cities with companies
that make only products with price < 100

Using EXISTS:

```
SELECT DISTINCT  Company.city
FROM      Company
WHERE NOT EXISTS (SELECT *
                  FROM Product
                  WHERE company = cname and Produc.price >= 100)
```

3. Subqueries in WHERE

Product (pname, price, company)

Company(cname, city)

Universal quantifiers

Find all cities with companies
that make <u>only</u> products with price < 100

Using ALL:

```
SELECT DISTINCT  Company.city
FROM     Company
WHERE 100 > ALL  (SELECT price
                  FROM Product
                  WHERE company = cname)
```

Question for Database Fans and their Friends

- Can we unnest the *universal quantifier* query ?

Monotone Queries

- A query Q is monotone if:
  - Whenever we add tuples to one or more of the tables…
  - … the answer to the query cannot contain fewer tuples

- **Fact**: all unnested queries are monotone
  - Proof: using the "nested for loops" semantics

- **Fact**: A query a universal quantifier is not monotone

- **Consequence**: we cannot unnest a query with a universal quantifier

Queries that must be nested

- Queries with universal quantifiers or with negation
- The drinkers-bars-beers example next
- This is a famous example from textbook on databases by Ullman

Rule of Thumb:

Non-monotone queries cannot be unnested. In particular, queries with a universal quantifier cannot be unnested

The drinkers-bars-beers example

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

Challenge: write these in SQL

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

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

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

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

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

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

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

x: $\forall y.$ Frequents(x, y)$\Rightarrow$ $\forall z.$(Serves(y,z) $\Rightarrow$ Likes(x,z))

Aggregation

```
SELECT  avg(price)
FROM     Product
WHERE   maker='Toyota'
```

```
SELECT  count(*)
FROM     Product
WHERE   year > 1995
```

SQL supports several aggregation operations:

sum, count, min, max, avg

Except count, all aggregations apply to a single attribute

Aggregation: Count

COUNT applies to duplicates, unless otherwise stated:

```
SELECT  Count(category)
FROM    Product
WHERE   year > 1995
```

same as Count(*)

We probably want:

```
SELECT  Count(DISTINCT category)
FROM    Product
WHERE   year > 1995
```

More Examples

Purchase(product, date, price, quantity)

```
SELECT  Sum(price * quantity)
FROM    Purchase
```

What do they mean ?

```
SELECT  Sum(price * quantity)
FROM    Purchase
WHERE   product = '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  Sum(price * quantity)
FROM    Purchase
WHERE   product = ‘Bagel’
```

90 (= 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, including aggregates.

1&2. FROM-WHERE-GROUPBY

Product	Price	Quantity
Bagel	3	20
Bagel	1.50	20
Banana	~~0.5~~	~~50~~
Banana	2	10
Banana	4	10

3. SELECT

Product	Price	Quantity
Bagel	3	20
Bagel	1.50	20
Banana	~~0.5~~	~~50~~
Banana	2	10
Banana	4	10

Product	TotalSales
Bagel	40
Banana	20

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

GROUP BY v.s. Nested Quereis

```
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 price > 1)
                            AS TotalSales
FROM        Purchase x
WHERE       price > 1
```

Why twice ?

Another Example

```
SELECT     product,
            sum(quantity) AS SumSales
            max(price) AS MaxQuantity
FROM       Purchase
GROUP BY product
```

What does it mean ?

Rule of thumb:

Every group in a GROUP BY is non-empty !

If we want to include empty groups in the output, then we need either a subquery, or a *left outer join* (see later)

HAVING Clause

Same query, except that we consider only products that had at least 100 buyers.

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

HAVING clause contains conditions on aggregates.

General form of Grouping and Aggregation

SELECT S

FROM $R_1,\ldots,R_n$

WHERE C1

GROUP BY $a_1,\ldots,a_k$

HAVING C2

Why ?

S = may contain attributes $a_1,\ldots,a_k$ and/or any aggregates but NO OTHER ATTRIBUTES

C1 = is any condition on the attributes in $R_1,\ldots,R_n$

C2 = is any condition on aggregate expressions

General form of Grouping and Aggregation

```
SELECT   S
FROM     R1,…,Rn
WHERE    C1
GROUP BY a1,…,ak
HAVING   C2
```

Evaluation steps:

1. Evaluate FROM-WHERE, apply condition C1
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

Advanced SQLizing

1. Unnesting Aggregates

2. Finding witnesses

Unnesting Aggregates

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

Find the number of companies in each city

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

Equivalent queries

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

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

Unnesting Aggregates

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

Find the number of products made in each city

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

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

They are NOT equivalent ! (WHY?)

More Unnesting

Author(<u>login</u>,name)

Wrote(login,url)

- Find authors who wrote ≥ 10 documents:
- Attempt 1: with nested queries

This is SQL by a novice

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

More Unnesting

- Find all authors who wrote at least 10 documents:
- Attempt 2: SQL style (with GROUP BY)

```
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

Finding Witnesses

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

For each store,
find its most expensive products

Finding Witnesses

Finding the maximum price is easy…

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

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

Finding Witnesses

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

```
SELECT Store.sname, Product.pname
FROM Store, Product,
        (SELECT Store.sid AS sid, max(Product.price) AS p
         FROM    Store, Product
         WHERE  Store.sid = Product.sid
          GROUP BY  Store.sid, Store.sname) X
WHERE  Store.sid = Product.sid
      and Store.sid = X.sid and Product.price = X.p
```

Finding Witnesses

There is a more concise solution here:

```
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)
```

NULLS in SQL

- Whenever we don’t have a value, we can put a NULL
- Can mean many things:
  - Value does not exists
  - 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 ?

Null Values

- If x= NULL then 4*(3-x)/7 is still NULL

- If x= NULL then x='Joe'   is UNKNOWN
- In SQL there are three boolean values:

FALSE = 0

UNKNOWN = 0.5

TRUE = 1

Null Values

- C1 AND C2 = min(C1, C2)
- C1 OR C2 = max(C1, C2)
- NOT C1 = 1 – C1

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

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

Rule in SQL: include only tuples that yield TRUE

Null Values

Unexpected behavior:

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

Some Persons are not included !

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

Outerjoins

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

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 Products that never sold will be lost !

Outerjoins

Product(<u>name</u>, category)
Purchase(prodName, store)

If we want the never-sold products, 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

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 ?

Application

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

Product(name, category)

Purchase(prodName, month, store)

```
SELECT Product.name, count(store)
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

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