

# CS3DB3 / SE4DB3 / SE6DB3 TUTORIAL

Xiao Jiao Wang

Feb 25, 2015

# Relational Algebra

- **IMPORTANT:** relational engines work on **bags**, no set !!!

# Union, intersection, and difference

- Union:  $\cup$     Intersection:  $\cap$     Difference:  $-$
- **Note:** Both operands must have the **same relation schema**.
- Example 1

**R:**

Product Name	Unit price
Melon	800G
Strawberry	150G
Melon	800G

**S:**

Product Name	Unit price
Melon	800G
Strawberry	150G
Apple	120G

**R-S**

Product Name	Unit price
Melon	800G

# Selection and projection

- Selection:  $\sigma_c(R)$ 
  - Picking all tuples of R that satisfy C.
  - C is a condition that refers to attributes of R.
- Projection:  $\pi_L(R)$ 
  - L is a list of attributes from the schema of R.
  - Constructed by looking at each tuple of R.
  - Schema of result contains exactly the fields in the projection list, with the same names that they had in relation R.
- Example 2

**R:**

Sname	Rating
Yuppy	9
Lubber	8
Guppy	5
Rusty	10

$\pi_{\text{name, rating}}(\sigma_{\text{rating} > 8}(R))$

Sname	Rating
Yuppy	9
Rusty	10

Don't forget schema

# Renaming and Product

- Products and joins: compositions of relations.
- Renaming :  $\rho_{R1(A1, \dots, An)}(R2)$ 
  - Gives a new schema to a relation.
  - Makes R1 be a relation with attributes A1, ... , An and the same tuples as R2.
- Product:  $R3 := R1 \times R2$ 
  - Also called cross-product or Cartesian product.
  - Pair each tuple t1 of R1 with each tuple t2 of R2 and concatenation t1 t2 is a tuple of R3.
    - # of tuples in R3 = (# of tuples in R1) × (# of tuples in R2)
  - Schema of R3 is the attributes of R1 and then R2, in order.
  - **Beware:** R1 and R2 have the common attribute A
    - In relational algebra, use renaming to distinguish.

# Renaming and Product (Cont.)

## □ Example 3

**R1**  
(2 tuples)

Name	Price
Melon	800G
Apple	120G

**R2**  
(3 tuples)

Fruit	Place
Melon	Canada
Lemon	Spain
Apple	France

**R1 × R2**  
(2\*3=6 tuples)

Name	Price	Fruit	Place
Melon	800G	Melon	Canada
Melon	800G	Lemon	Spain
Melon	800G	Apple	France
Apple	120G	Melon	Canada
Apple	120G	Lemon	Spain
Apple	120G	Apple	France

Schema:  
the attributes of R1 and  
then R2, in order.

R1 and R2 have no  
common attributes.

# Renaming and Product (Cont.)

## □ Example 4

**R1**  
(2 tuples)

Name	Price
Melon	800G
Apple	120G

**R2**  
(3 tuples)

Name	Place
Melon	Canada
Lemon	Spain
Apple	France

**R1 × R2**  
(2\*3=6 tuples)

R1.Name	Price	R2.Name	Place
Melon	800G	Melon	Canada
Melon	800G	Lemon	Spain
Melon	800G	Apple	France
Apple	120G	Melon	Canada
Apple	120G	Lemon	Spain
Apple	120G	Apple	France

R1 and R2 have a  
common attribute.

# Theta-Join

- Theta-Join:  $R3 := R1 \bowtie_c R2$ 
  - Take the product  $R1 \times R2$ .
  - Then apply  $\sigma_c$  to the result.
- Example 5

**R1**

Name	Price
Melon	800G
Apple	120G

**R2**

Fruit	Place
Melon	Canada
Lemon	Spain
Apple	France

$R1 \bowtie_{\text{Name=Fruit}} R2$

Name	Price	Fruit	Place
Melon	800G	Melon	Canada
Apple	120G	Apple	France



# Theta-Join (Cont.)

## □ Example 5

**R1**

Name	Price
Melon	800G
Apple	120G

**R2**

Name	Place
Melon	Canada
Lemon	Spain
Apple	France

**R1**  $\bowtie_{R1.Name=R2.Name}$  **R2**

R1.Name	Price	R2.Name	Place
Melon	800G	Melon	Canada
Apple	120G	Apple	France

There are **4** columns.

# Natural Join

- Natural Join:  $R3 := R1 \bowtie R2$

- Connects two relations by:

- Equating attributes of the same name, and
    - Projecting out one copy of each pair of equated attributes.

- Example 6

**R1**

Name	Price
Melon	800G
Apple	120G

**R2**

Name	Place
Melon	Canada
Lemon	Spain
Apple	France

$R1 \bowtie R2$

Name	Price	Place
Melon	800G	Canada
Apple	120G	France

There are **3** columns.

# Precedence of relational operators


- $[\sigma, \pi, \rho]$  (highest)
- $[\times, \bowtie]$
- $\cap$
- $[\cup, -]$

# Duplicate Elimination and Sorting

- Duplicate elimination:  $\delta(R)$ 
  - **Recall**: relational engines work on bags.
  - Consists of one copy of each tuple that appears in R2 one or more times.
  - SQL: `SELECT DISTINCT ...`
- Sorting:  $\tau_L(R)$ 
  - L is a list of some of the attributes of R2.
  - Sorted first on the value of the first attribute on L, then on the second attributes of L, and so on.

# Grouping and Aggregation

- Grouping and Aggregation :  $\gamma_L(R)$ 
  - L is a list of elements that are either
    - Grouping attributes
    - $AGG(A)$ , where  $AGG$  is one of the aggregation operators such as **SUM**, **AVG**, **COUNT**, **MIN**, **MAX** and **A** is an attribute.
      - An arrow and a new attribute name renames the component
  - Example:  $\gamma_{A,B,AVG(C) \rightarrow X}(R)$



# SQL and relational algebra


□ **SELECT**  $A_1, A_2, \dots, A_n$   
**FROM**  $R_1, R_2, \dots, R_m$   
**WHERE**  $P$

is equivalent to the **multiset** relational algebra expression

Don't forget the parenthesis since  $\sigma$  has a higher Precedence than  $[\times, \bowtie]$

$$\prod_{A_1, A_2, \dots, A_n} (\sigma_P (R_1 \times R_2 \times \dots \times R_m))$$

# SQL and relational algebra (Cont.)

- Example 1     **Takes** (id, course\_id, semester, year, grade)  
                  **Teaches**(name, course\_id, semester, year)
  - Find the IDs of all courses who were taught by an instructor named Jones.
  - SQL
    - **SELECT** Teaches.course\_id
    - FROM** Takes, Teaches
    - WHERE** name = 'Jones' AND Takes.course\_id = Teaches.course\_id;
  - Relation algebra
    - WAY 1:  $\Pi_{\text{course\_id}}(\sigma_{\text{name}='Jones'}(\text{Takes} \bowtie \text{Teaches}))$  Because of  
common attribute
    - WAY 2:  $\Pi_{\text{Teaches.course\_id}}(\sigma_{\text{name}='Jones'}(\text{Takes} \bowtie_{\text{Takes.course\_id} = \text{Teaches.course\_id}} \text{Teaches}))$
    - WAY 3:  $\Pi_{\text{Teaches.course\_id}}(\sigma_{\text{name}='Jones'} \wedge \text{Takes.course\_id} = \text{Teaches.course\_id}(\text{Takes} \times \text{Teaches}))$

# SQL and relational algebra (Cont.)

- Example 2
  - **Works** (pname, cname, salary)
  - Find the names of all employees who earn more than every employee of “First Bank”.
  - SQL

```
SELECT pname
FROM Works
WHERE salary > ALL (SELECT salary
                    FROM Works
                    WHERE cname= 'First Bank');
```

- Relational algebra

$R1 := \Pi_{w1.pname} (\rho_{w1}(\text{Works}) \bowtie_{w1.salary \leq w2.salary \wedge w2.cname = \text{'First Band'}} \rho_{w2}(\text{Works}))$

$\text{Result} := \Pi_{pname}(\text{Works}) - R1$

Assignment:  
create temporary relation  
names



# SQL and relational algebra (Cont.)

▣ **SELECT** A1, A2, **AGG**(A3) AS AGG3

**FROM** R1, R2, ..., Rm

**WHERE** P

**GROUP BY** A1, A2

■ Is equivalent to the multiset relational algebra expression

$\gamma_{A1, A2, AGG(A3) \rightarrow AGG3}(\sigma_P(R1 \times R2 \times \dots \times Rm))$

▣ If only display attribute A1 and AGG3, then

$\Pi_{A1, AGG3}(\gamma_{A1, A2, AGG(A3) \rightarrow AGG3}(\sigma_P(R1 \times R2 \times \dots \times Rm)))$

# SQL and relational algebra (Cont.)

## □ Example 3

□ **Takes** (student id, course id, semester, year, grade)

□ Find the enrollment of each course that was offered in Fall 2009.

□ SQL

```
SELECT course_id, count(*) as enrollment
FROM Takes
WHERE year=2009 AND semester='Fall'
GROUP BY course_id;
```

□ Relational Algebra

$\gamma_{\text{course\_id, count(*)} \rightarrow \text{enrollment}} (\sigma_{\text{year}=2009 \wedge \text{semester}=\text{"Fall"}}(\text{Takes}))$

# SQL and relational algebra (Cont.)

- Example 4
  - **Takes** (student id, course id, semester, year, grade)
  - Find the maximum enrollment in Fall 2009.

- SQL

```
SELECT MAX(enrollment)
FROM (SELECT course_id, count(*) as enrollment
      FROM Takes
      WHERE year=2009 AND semester='Fall'
      GROUP BY course_id);
```

- Relational Algebra

$R := \gamma_{\text{course\_id, count}(*)} \rightarrow \text{enrollment} (\sigma_{\text{year}=2009 \wedge \text{semester}=\text{“Fall”}}(\text{Takes}))$   
Result :=  $\gamma_{\text{max(enrollment)}}(R)$