

Chapter 2

The relational Model of data

Relational algebra

Contents

- What is a data model?
- Basics of the relational model

- How to define?
- How to query?
- Constraints on relations

An algebraic query language

- What is an “Algebra”?
- Mathematical system consisting of:
 - *Operands* --- variables or values from which new values can be constructed.
 - *Operators* --- symbols denoting procedures that construct new values from given values.
 - Eg. how many students in my classroom? $E = (x + y)$

What is Relational Algebra?

- An algebra whose **operands** are relations or variables that represent relations.
- **Operators** are designed to do with relations in a database.
 - The result is used as a *query language* for relations.
 - how many students in my classroom?
(**x union y**)

Core Relational Algebra

- Union, intersection, and difference.
 - Usual set operations, but *both operands must have the same relation schema*.
- Selection: picking certain rows.
- Projection: picking certain columns.
- Products and joins: compositions of relations.
- Renaming of relations and attributes.

Selection

□ $R1 := \sigma_C(R2)$

- C is a condition (as in “if” statements) that refers to attributes of $R2$.
- $R1$ is all those tuples of $R2$ that satisfy C .

Example: Selection

Relation Sells:

| bar | beer | price |
|-------|--------|-------|
| Joe's | Bud | 2.50 |
| Joe's | Miller | 2.75 |
| Sue's | Bud | 2.50 |
| Sue's | Miller | 3.00 |

JoeMenu := $\sigma_{\text{bar}='Joe's'}(\text{Sells})$:

| bar | beer | price |
|-------|--------|-------|
| Joe's | Bud | 2.50 |
| Joe's | Miller | 2.75 |

Projection

□ $R1 := \pi_L(R2)$

- L is a list of attributes from the schema of $R2$.
- $R1$ is constructed by looking at each tuple of $R2$, extracting the attributes on list L , in the order specified, and creating from those components a tuple for $R1$.
- Eliminate duplicate tuples, if any.

Example: Projection

Relation Sells:

| bar | beer | price |
|-------|--------|-------|
| Joe's | Bud | 2.50 |
| Joe's | Miller | 2.75 |
| Sue's | Bud | 2.50 |
| Sue's | Miller | 3.00 |

Prices := $\pi_{\text{beer,price}}(\text{Sells})$:

| beer | price |
|--------|-------|
| Bud | 2.50 |
| Miller | 2.75 |
| Miller | 3.00 |

Extended Projection

- Using the same Π_L operator, we allow the list L to contain arbitrary expressions involving attributes:
 1. Arithmetic on attributes, e.g., $A+B \rightarrow C$.
 2. Duplicate occurrences of the same attribute.

Example: Extended Projection

$$R = \left(\begin{array}{|c|c|} \hline A & B \\ \hline 1 & 2 \\ \hline 3 & 4 \\ \hline \end{array} \right)$$
$$\pi_{A+B \rightarrow C, A, A}(R) = \begin{array}{|c|c|c|} \hline C & A1 & A2 \\ \hline 3 & 1 & 1 \\ \hline 7 & 3 & 3 \\ \hline \end{array}$$

Product (cross join)

□ $R3 := R1 \times R2$

- Pair each tuple $t1$ of $R1$ with each tuple $t2$ of $R2$.
- Concatenation $t1t2$ is a tuple of $R3$.
- Schema of $R3$ is the attributes of $R1$ and then $R2$, in order.
- If attribute A has the same name in $R1$ and $R2$: use $R1.A$ and $R2.A$.

Example: $R3 := R1 \times R2$

R1(

| A, | B) |
|----|----|
| 1 | 2 |
| 3 | 4 |

R2(

| B, | C) |
|----|----|
| 5 | 6 |
| 7 | 8 |
| 9 | 10 |

R3(

| A, | R1.B, | R2.B, | C) |
|----|-------|-------|----|
| 1 | 2 | 5 | 6 |
| 1 | 2 | 7 | 8 |
| 1 | 2 | 9 | 10 |
| 3 | 4 | 5 | 6 |
| 3 | 4 | 7 | 8 |
| 3 | 4 | 9 | 10 |

Theta-Join

- $R3 := R1 \bowtie_C R2$
 - Take the product $R1 \times R2$.
 - Then apply σ_C to the result.
- σ, C can be any boolean-valued condition.
 - Historic versions of this operator allowed only $A \theta B$, where θ is $=, <, >$, etc.; hence the name “theta-join.”

Example: Theta Join

Sells(

| bar, | beer, | price |
|-------|--------|-------|
| Joe's | Bud | 2.50 |
| Joe's | Miller | 2.75 |
| Sue's | Bud | 2.50 |
| Sue's | Coors | 3.00 |

)

Bars(

| name, | addr |
|-------|-----------|
| Joe's | Maple St. |
| Sue's | River Rd. |

)

BarInfo := Sells $\bowtie_{\text{Sells.bar} = \text{Bars.name}}$ Bars

BarInfo(

| bar, | beer, | price, | name, | addr |
|-------|--------|--------|-------|-----------|
| Joe's | Bud | 2.50 | Joe's | Maple St. |
| Joe's | Miller | 2.75 | Joe's | Maple St. |
| Sue's | Bud | 2.50 | Sue's | River Rd. |
| Sue's | Coors | 3.00 | Sue's | River Rd. |

)

Natural Join

- A useful join variant (*natural join*) connects two relations by:
 - Equating attributes of the same name, and
 - Projecting out one copy of each pair of equated attributes.
- Denoted $R3 := R1 \bowtie R2$.

Example: Natural Join

| Sells(| bar, | beer, | price |) | Bars(| bar, | addr |) |
|--------|-------|--------|-------|---|-------|-------|-----------|---|
| | Joe's | Bud | 2.50 | | | Joe's | Maple St. | |
| | Joe's | Miller | 2.75 | | | Sue's | River Rd. | |
| | Sue's | Bud | 2.50 | | | | | |
| | Sue's | Coors | 3.00 | | | | | |

BarInfo := Sells \bowtie Bars

Note: Bars.name has become Bars.bar to make the natural join "work."

| BarInfo(| bar, | beer, | price, | addr |) |
|----------|-------|--------|--------|-----------|---|
| | Joe's | Bud | 2.50 | Maple St. | |
| | Joe's | Miller | 2.75 | Maple St. | |
| | Sue's | Bud | 2.50 | River Rd. | |
| | Sue's | Coors | 3.00 | River Rd. | |

Example:

two or more common attributes

□ $R(A,B,C,D)$ natural join $S(A,B,F)$

| A | B | C | D | | A | B | F | = | A | B | C | D | F |
|----|----|----|----|---|----|----|----|---|----|----|----|----|----|
| a1 | b1 | c1 | d1 | ⋈ | a1 | b1 | f1 | = | a1 | b1 | c1 | d1 | f1 |
| a1 | b1 | c2 | d2 | | a1 | b2 | f2 | | a1 | b1 | c2 | d2 | f1 |

If theta join on $R.A=S.B$ and $R.B=S.B$? or R product S ?

| A | B | C | D | A | B | F |
|----|----|----|----|----|----|----|
| a1 | b1 | c1 | d1 | a1 | b1 | f1 |
| a1 | b1 | c2 | d2 | a1 | b2 | f2 |

Renaming

- The ρ operator gives a new schema to a relation.
- $R1 := \rho_{R1(A1, \dots, An)}(R2)$ makes R1 be a relation with attributes $A1, \dots, An$ and the same tuples as R2.
- Simplified notation: $R1(A1, \dots, An) := R2$.

Example: Renaming

Bars(

| name, | addr |
|-------|-----------|
| Joe's | Maple St. |
| Sue's | River Rd. |

)

$R(\text{bar}, \text{addr}) := \text{Bars}$

R(

| bar, | addr |
|-------|-----------|
| Joe's | Maple St. |
| Sue's | River Rd. |

)

Building Complex Expressions

- Combine operators with *parentheses* and *precedence rules*.
- Precedence of relational operators:
 1. $[\sigma, \pi, \rho]$ (highest).
 2. $[x, \bowtie]$.
 3. \cap .
 4. $[\cup, -]$

Three notations:

1. Sequences of assignment statements. $:=$
2. Expressions with several operators.
3. Expression trees

Example: a Query

- Bars(name, addr)
- Sells(bar, beer, price)

Query: find the names of all the bars that are **either** on Maple St. **or** sell Bud for less than \$3.

Three notations to represent the query.

1. Sequences of Assignments

- Create temporary relation names.
- Renaming can be implied by giving relations a list of attributes.

$B1 := \pi_{\text{name}} \sigma_{\text{Address}='Maper str.'} (\text{Bars})$

$B2 := \pi_{\text{bar}} \sigma_{\text{beer}='Bud' \text{ and } \text{price} < 3} (\text{Sells})$

$B3 := \rho_{\text{name}} (B2)$

$B4 = B1 \cup B3$

2. Expressions in a Single Assignment

Bar: =

$\pi_{\text{name}} (\sigma_{\text{address}='Maper str.'} (\text{Bars})) \cup$


$\rho_{\text{name}} (\pi_{\text{bar}} (\sigma_{\text{beer}='Bud' \text{ and } \text{price} < 3} (\text{Sells})))$

3. Expression Trees

- Leaves are **operands** --- either variables standing for relations or particular, constant relations.
- Interior nodes are **operators**, applied to their child or children.

As a Tree:

find the names of all the bars that are either on Maple St. or sell Bud for less than \$3.



Self Join

□ Sometimes, **conditions** and query **results** are in the same table.

□ **Recursion** situation:

Parents (Parents, child) in DB

We want to know **grandparents** information.

Example: Self-Join

- Using `Sells(bar, beer, price)`, find the bars that sell *two different beers* at the *same price*.

| | | |
|---------|--------|-----|
| Joe's | Bud | 2.5 |
| Joe's | Coors | 3.0 |
| Joe's | Miller | 2.5 |
| Sue's | Bud | 2.5 |
| Sue's | Coors | 3.5 |
| Marry's | Miller | 2.5 |


Example: Self-Join (cont.)

- **Strategy**: by renaming, define a copy of Sells, called $S(\text{bar}, \text{beer1}, \text{price})$.

The natural join of Sells and S consists of quadruples (bar, beer, beer1, price) such that the **bar sells both beers at this price.**

The Tree

Query: find the bars that sell two different beers at the same price.



Query expression: $\Pi_{\text{bar}}(\sigma_{\text{beer} <> \text{beer1}}(\rho_{S(\text{bar}, \text{beer1}, \text{price})}(\text{sell}) \bowtie \text{sell}))$

<> Change to <

| Bar | Beer1 | price | | Bar | Beer | price |
|---------|--------|-------|---|---------|--------|-------|
| Joe's | Bud | 2.5 | ⋈ | Joe's | Bud | 2.5 |
| Joe's | Coors | 3.0 | | Joe's | Coors | 3.0 |
| Joe's | Miller | 2.5 | | Joe's | Miller | 2.5 |
| Sue's | Bud | 2.5 | | Sue's | Bud | 2.5 |
| Sue's | Coors | 3.5 | | Sue's | Coors | 3.5 |
| Marry's | Miller | 2.5 | | Marry's | Miller | 2.5 |

| bar | beer | beer1 | price | | bar | beer | beer1 | price |
|------|--------|--------|-------|----|------|--------|-------|-------|
| joes | bud | bud | 2.5 | => | joes | bud | 2.5 | |
| joes | bud | miller | 2.5 | | Joes | miller | 2.5 | |
| joes | coors | coors | 3.0 | | | | | |
| joes | miller | miller | 2.5 | | | | | |
| Joes | miller | bud | 2.5 | | | | | |
| ... | .. | .. | .. | | | | | |

Beer<>beer1

If we do not want pairs appear twice? What should we do?

Schemas for Results

- **Union, intersection, and difference:** the schemas of the two operands must be the same, so use that schema for the result.
- **Selection:** schema of the result is the same as the schema of the operand.
- **Projection:** list of attributes tells us the schema.

Schemas for Results (cont.)

- **Product**: schema is the attributes of both relations.
 - Use $R.A$, etc., to distinguish two attributes named A .
- **Theta-join**: same as product.
- **Natural join**: union of the attributes of the two relations.
- **Renaming**: the operator tells the schema.

Relational Algebra on Bags

- A *bag* (or *multiset*) is like a set, but an element may appear more than once.
- **Example:** $\{1,2,1,3\}$ is a bag.
- **Example:** $\{1,2,3\}$ is also a bag that happens to be a set.

Why Bags?

- ❑ SQL, the most important query language for relational databases, is actually **a bag language**.
- ❑ Some operations, like projection, are more efficient on bags than sets.

Operations on Bags

- ❑ **Selection** applies to each tuple, so its effect on bags is like its effect on sets.
- ❑ **Projection** also applies to each tuple, but as a bag operator, we **do not eliminate duplicates**.
- ❑ **Products** and **joins** are done on each pair of tuples, so duplicates in bags have no effect on how we operate.

Example: Bag Selection

R(

| A | B |
|---|---|
| 1 | 2 |
| 5 | 6 |
| 1 | 2 |

)

$\sigma_{A+B < 5}(R) =$

| A | B |
|---|---|
| 1 | 2 |
| 1 | 2 |

Example: Bag Projection

R(

| A | B |
|---|---|
| 1 | 2 |
| 5 | 6 |
| 1 | 2 |

)

$\pi_A(R) =$

| |
|---|
| A |
| 1 |
| 5 |
| 1 |

Example: Bag Product

R(

| A | B |
|---|---|
| 1 | 2 |
| 5 | 6 |
| 1 | 2 |

)

S(

| B | C |
|---|---|
| 3 | 4 |
| 7 | 8 |

)

R X S =

| A | R.B | S.B | C |
|---|-----|-----|---|
| 1 | 2 | 3 | 4 |
| 1 | 2 | 7 | 8 |
| 5 | 6 | 3 | 4 |
| 5 | 6 | 7 | 8 |
| 1 | 2 | 3 | 4 |
| 1 | 2 | 7 | 8 |

Example: Bag Theta-Join

R(

| A | B |
|---|---|
| 1 | 2 |
| 5 | 6 |
| 1 | 2 |

)

S(

| B | C |
|---|---|
| 3 | 4 |
| 7 | 8 |

)

R $\bowtie_{R.B < S.B}$ S =

| A | R.B | S.B | C |
|---|-----|-----|---|
| 1 | 2 | 3 | 4 |
| 1 | 2 | 7 | 8 |
| 5 | 6 | 7 | 8 |
| 1 | 2 | 3 | 4 |
| 1 | 2 | 7 | 8 |

Bag Union

- An element appears in the union of two bags the sum of the number of times it appears in each bag.
- **Example:** $\{1,2,1\} \cup \{1,1,2,3,1\} = \{1,1,1,1,1,2,2,3\}$

Bag Intersection

- An element appears in the intersection of two bags the minimum of the number of times it appears in either.
- **Example:** $\{1,2,1,1\} \cap \{1,2,1,3\} = \{1,1,2\}$.

Bag Difference

- An element appears in the difference $A - B$ of bags as many times as it appears in A , minus the number of times it appears in B .
 - But never less than 0 times.
- **Example:** $\{1,2,1,1\} - \{1,2,3\} = \{1,1\}$.

Beware: Bag Laws \neq Set Laws

- *not all* algebraic laws that hold for sets also hold for bags.
- **Example:** the commutative law for union ($R \cup S = S \cup R$) *does* hold for bags.
 - Since addition is commutative, adding the number of times x appears in R and S doesn't depend on the order of R and S .

Example: A Law That Fails

- Set union is *idempotent*, meaning that $S \cup S = S$.
- However, for bags, if x appears n times in S , then it appears $2n$ times in $S \cup S$.
- Thus $S \cup S \neq S$ in general.
 - e.g., $\{1\} \cup \{1\} = \{1,1\} \neq \{1\}$.

Comparison

| Operation ^o | Set model ^o | Bag model ^o |
|--------------------------------|--|---|
| R union S ^o | Elements that are in R or S or both ^o | Sum the times an element appears in the two bags ^o |
| R intersection S ^o | Elements that are in both R and S ^o | Take the minimum of the number of occurrences in each bag ^o |
| R difference S ^o | Elements that are in R but not in S ^o | Proper-subtract the number of occurrences in the two bags ^o |
| Selection $c(R)$ ^o | <u>Tuples</u> which satisfied the condition C ^o | <u>Tuples</u> which satisfied the condition C ^o |
| Projection $L(R)$ ^o | Duplicate <u>tuples</u> eliminated when projecting on L ^o | Duplicate <u>tuples</u> not eliminated when projecting on L ^o |
| R Join S ^o | Eliminate duplicate <u>tuples</u> ^o | Do not eliminate duplicate <u>tuples</u> ^o |
| ... ^o | ... ^o | ^o |

Constraints on Relations

- The ability to restrict the data that may be stored in a database.
- Relational algebra: used as **a constraint language** abstractly.

Two ways to express constraints

R, S : **expressions** of relational algebra

1. $R=0$: “there are no tuples in the result of R ” or the value of R must be empty.
2. $R \subseteq S$: “every tuple in the result of R must be in the result of S ”

For example

Beers (name, manf)

Bars (name, addr, license)

Sells (bar, beer, price)

□ Legal value constrain:

$\sigma_{\text{Price} < 0} (\text{Sells}) = 0$ (empty)

Means **all the price is not allowed to lower than 0**

Referential Integrity Constraint

Beers (name, manf)

Bars (name, addr, license)

Sells (bar, beer, price)

□ Referential Integrity constraint:

$\pi_{\text{bar}}(\text{Sells}) \subseteq \pi_{\text{name}}(\text{Bars})$

$\pi_{\text{beer}}(\text{Sells}) \subseteq \pi_{\text{name}}(\text{Beers})$

Summary of Chapter 2

- Relational Data models
 1. **Structure**: schemas, relations, keys, how to define structure.
 2. **Operations**: relational algebra as a query language (set and bag), three notations
 3. **Constraints**: Relational algebra as a constraint language (two ways to express)

Homework for chapter 2

- ❑ Exercise 2.3.1 (DDL)
- ❑ Exercise 2.4.1 a), f), h) (DML)
- ❑ Exercise 2.5.1 b)

Submit: *electronically*

<ftp://public.sjtu.edu.cn> to <public-files/upload/chapter2>

User name: fli Password: public

Name of your homework is your **studentID**