

# Outline

- ▶ Overview of modeling
- ▶ Relational Model (Chapter 2)
  - Basics
  - Keys
  - Relational operations
  - Relational algebra basics
- ▶ SQL
  - Basic Data Definition (3.2)
  - Setting up the PostgreSQL database
  - Basic Queries (3.3-3.5)
  - Null values (3.6)
  - Aggregates (3.7)

# Relational Query Languages

- ▶ Example schema:  $R(A, B)$
- ▶ Practical languages
  - SQL
    - select A from R where B = 5;
  - Datalog (sort of practical)
    - $q(A) :- R(A, 5)$
- ▶ Formal languages
  - Relational algebra  
 $\pi_A(\sigma_{B=5}(R))$
  - Tuple relational calculus  
 $\{ t : \{A\} \mid \exists s : \{A, B\} ( R(A, B) \wedge s.B = 5 ) \}$
  - Domain relational calculus
    - Similar to tuple relational calculus

# Modeling Languages

- ▶ Some of languages are “procedural” and provide a set of operations
  - Each operation takes one or two relations as input, and produces a single relation as output
  - Examples: Relational Algebra
- ▶ The “non-procedural” (also called “declarative”) languages specify the output, but don’t specify the operations
  - SQL, Relational calculus
  - Datalog (used as an intermediate layer in quite a few systems today)

## Select Operation

Choose a subset of the tuples that satisfies some predicate  
Denoted by  $\sigma$  in relational algebra

Relation r

A	B	C	D
$\alpha$	$\alpha$	1	7
$\alpha$	$\beta$	5	7
$\beta$	$\beta$	12	3
$\beta$	$\beta$	23	10

$\sigma_{A=B \wedge D > 5}(r)$

A	B	C	D
$\alpha$	$\alpha$	1	7
$\beta$	$\beta$	23	10

# Project

Choose a subset of the columns (for all rows)  
 Denoted by  $\Pi$  in relational algebra

Relation  $r$

A	B	C	D
$\alpha$	$\alpha$	1	7
$\alpha$	$\beta$	5	7
$\beta$	$\beta$	12	3
$\beta$	$\beta$	23	10

$\Pi_{A,D}(r)$

A	D
$\alpha$	7
$\alpha$	7
$\beta$	3
$\beta$	10

A	D
$\alpha$	7
$\beta$	3
$\beta$	10

Relational algebra following “set” semantics – so no duplicates  
 SQL allows for duplicates – we will cover the formal semantics later

# Set Union, Difference

Relation  $r, s$

A	B
$\alpha$	1
$\alpha$	2
$\beta$	1

$r$

A	B
$\alpha$	2
$\beta$	3

$s$

$r \cup s$ :

A	B
$\alpha$	1
$\alpha$	2
$\beta$	1
$\beta$	3

$r - s$ :

A	B
$\alpha$	1
$\beta$	1

Must be compatible schemas

What about intersection ?

Can be derived

$$r \cap s = r - (r - s);$$

$r \cap s$ :

A	B
$\alpha$	2

# Cartesian Product

Combine tuples from two relations

If one relation contains N tuples and the other contains M tuples, the result would contain N\*M tuples

The result is rarely useful – almost always you want pairs of tuples that satisfy some condition

Relation r, s

A	B	C	D	E
$\alpha$	1	$\alpha$	10	a
$\beta$	2	$\beta$	10	a
		$\beta$	20	b
		$\gamma$	10	b

r

s

$r \times s$ :

A	B	C	D	E
$\alpha$	1	$\alpha$	10	a
$\alpha$	1	$\beta$	10	a
$\alpha$	1	$\beta$	20	b
$\alpha$	1	$\gamma$	10	b
$\beta$	2	$\alpha$	10	a
$\beta$	2	$\beta$	10	a
$\beta$	2	$\beta$	20	b
$\beta$	2	$\gamma$	10	b

# Joins

Combine tuples from two relations if the pair of tuples satisfies some constraint (equivalent to Cartesian Product followed by a Select)

Relation r, s

A	B	C	D	E
$\alpha$	1	$\alpha$	10	a
$\beta$	2	$\beta$	10	a
		$\beta$	20	b
		$\gamma$	10	b

r

s

$r \bowtie_{A=C} s$ :

A	B	C	D	E
$\alpha$	1	$\alpha$	10	a
<del><math>\alpha</math></del>	<del>1</del>	<del><math>\beta</math></del>	<del>10</del>	<del>a</del>
<del><math>\alpha</math></del>	<del>1</del>	<del><math>\beta</math></del>	<del>20</del>	<del>b</del>
<del><math>\alpha</math></del>	<del>1</del>	<del><math>\gamma</math></del>	<del>10</del>	<del>b</del>
<del><math>\beta</math></del>	<del>2</del>	<del><math>\alpha</math></del>	<del>10</del>	<del>a</del>
$\beta$	2	$\beta$	10	a
$\beta$	2	$\beta$	20	b
<del><math>\beta</math></del>	<del>2</del>	<del><math>\gamma</math></del>	<del>10</del>	<del>b</del>

# Natural Join

Combine tuples from two relations if the pair of tuples agree on the common columns (with the same name)

dept_name	building	budget
Biology	Watson	90000
Comp. Sci.	Taylor	100000
Elec. Eng.	Taylor	85000
Finance	Painter	120000
History	Painter	50000
Music	Packard	80000
Physics	Watson	70000

Figure 2.5 The *department* relation.

department  $\bowtie$  instructor:

ID	name	salary	dept_name	building	budget
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
12121	Wu	90000	Finance	Painter	120000
15151	Mozart	40000	Music	Packard	80000
22222	Einstein	95000	Physics	Watson	70000
32343	El Said	60000	History	Painter	50000
33456	Gold	87000	Physics	Watson	70000
45565	Katz	75000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
76543	Singh	80000	Finance	Painter	120000
76766	Crick	72000	Biology	Watson	90000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000

Figure 2.12 Result of natural join of the *instructor* and *department* relations.

ID	name	dept_name	salary
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

Figure 2.4 Unsorted display of the *instructor* relation.

# Rename Operation

- ▶ Allows us to name, and therefore to refer to, the results of relational-algebra expressions.
- ▶ Allows us to refer to a relation by more than one name.

Example:

$$\rho_X(E)$$

returns the expression  $E$  under the name  $X$

If a relational-algebra expression  $E$  has arity  $n$ , then

$$\rho_X(A_1, A_2, \dots, A_n)(E)$$

returns the result of expression  $E$  under the name  $X$ , and with the attributes renamed to  $A_1, A_2, \dots, A_n$ .

# Relational Algebra

- ▶ Those are the basic operations
- ▶ What about SQL Joins ?
  - Compose multiple operators together

$$\sigma_{A=C}(r \times s)$$

- ▶ Additional Operations
  - Set intersection
  - Natural join
  - Division
  - Assignment

# Additional Operators

- ▶ Set intersection ( $\cap$ )
  - $r \cap s = r - (r - s)$ ;
  - SQL Equivalent: intersect
- ▶ Assignment ( $\leftarrow$ )
  - A convenient way to right complex RA expressions
  - Essentially for creating “temporary” relations
    - $temp1 \leftarrow \prod_{R-S}(r)$
  - SQL Equivalent: “create table as...”

# Additional Operators: Joins

## ▶ Natural join ( $\bowtie$ )

- A Cartesian product with equality condition on common attributes
- Example:
  - if  $r$  has schema  $R(A, B, C, D)$ , and if  $s$  has schema  $S(E, B, D)$
  - Common attributes:  $B$  and  $D$
  - Then:

$$r \bowtie s = \Pi_{r.A, r.B, r.C, r.D, s.E} (\sigma_{r.B = s.B \wedge r.D = s.D} (r \times s))$$

## ▶ SQL Equivalent:

- `select r.A, r.B, r.C, r.D, s.E from r, s where r.B = s.B and r.D = s.D, OR`
- `select * from r natural join s`

# Additional Operators: Joins

## ▶ *Equi-join*

- A join that only has equality predicates

## ▶ *Theta-join* ( $\bowtie_{\theta}$ )

- $r \bowtie_{\theta} s = \sigma_{\theta}(r \times s)$  (arbitrary predicates)

## ▶ *Left outer join* ( $\bowtie_{\leftarrow}$ )

- If have  $r(A, B), s(B, C)$ , then:
  - $r \bowtie_{\leftarrow} s = (r \bowtie s) \cup$  (“all non-matching rows in  $r$  w/ nulls for  $s$ 's attributes”)
- What is  $(r - \Pi_{r.A, r.B}(r \bowtie s))$ ? (rows of  $r$  that match rows of  $s$ )
- $r \bowtie_{\leftarrow} s = (r \bowtie s) \cup \rho_{temp(A, B, C)}((r - \Pi_{r.A, r.B}(r \bowtie s)) \times \{\text{NULL}\})$

## Additional Operators: Join Variations

- ▶ Tables:  $r(A, B), s(B, C)$

name	Symbol	SQL Equivalent	RA expression
cross product	$\times$	select * from r, s;	$r \times s$
natural join	$\bowtie$	natural join	$\pi_{r.A, r.B, s.C} \sigma_{r.B = s.B}(r \times s)$
equi-join	$\bowtie_{\theta}$ ( <i>theta must be equality</i> )		
theta join	$\bowtie_{\theta}$	from .. where $\theta$ ;	$\sigma_{\theta}(r \times s)$
left outer join	$r \bowtie\!\!\!\!\! \leftarrow s$	left outer join (with “on”)	(see previous slide)
full outer join	$r \bowtie\!\!\!\!\! \leftarrow\!\!\!\!\! s$	full outer join (with “on”)	-
(left) semijoin	$r \bowtie\!\!\!\!\! \leftarrow s$	none	$\pi_{r.A, r.B}(r \bowtie\!\!\!\!\! \leftarrow s)$
(left) antijoin	$r \triangleleft s$	none	$r - \pi_{r.A, r.B}(r \bowtie\!\!\!\!\! \leftarrow s)$

## Additional Operators: Division

- ▶ Assume  $r(R), s(S)$ , for queries where  $S \subseteq R$ :
  - $r \div s$
- ▶ Think of it as “opposite of Cartesian product”
  - $r \div s = t$  *iff*  $t \times s \subseteq r$

A	B	C	D	E
$\alpha$	1	$\alpha$	10	a
$\alpha$	1	$\beta$	10	a
$\alpha$	1	$\beta$	20	b
$\alpha$	1	$\gamma$	10	b
$\beta$	2	$\alpha$	10	a
$\beta$	2	$\beta$	10	a
$\beta$	2	$\beta$	20	b
$\beta$	2	$\gamma$	10	b

$\div$ 

A	B
$\alpha$	1
$\beta$	2

 $=$ 

C	D	E
$\alpha$	10	a
$\beta$	10	a
$\beta$	20	b
$\gamma$	10	b



# Relational Algebra Examples

Find all loans of over \$1200:

$$\sigma_{\text{amount} > 1200} (\text{loan})$$

Find the loan number for each loan of an amount greater than \$1200:

$$\pi_{\text{loan-number}} (\sigma_{\text{amount} > 1200} (\text{loan}))$$

Find names of all customers who have a loan, account, or both, from the bank:

$$\pi_{\text{customer-name}} (\text{borrower}) \cup \pi_{\text{customer-name}} (\text{depositor})$$

# Relational Algebra Examples

Find names of customers who have a loan and an account at bank:

$$\pi_{\text{customer-name}} (\text{borrower}) \cap \pi_{\text{customer-name}} (\text{depositor})$$

Find names of customers who have a loan at the UMD branch:

$$\pi_{\text{customer-name}} (\sigma_{\text{branch-name}=\text{"UMD"}} (\sigma_{\text{borrower.loan-number} = \text{loan.loan-number}} (\text{borrower} \times \text{loan})))$$

# Relational Algebra Examples

Find *largest* account balance, assume relation is  $\{(1), (2), (3)\}$ :

1  
2  
3

Rename the account relation to d:

$\Pi_{\text{balance}}(\text{account}) - \Pi_{\text{account.balance}}(\sigma_{\text{account.balance} < \text{d.balance}}(\text{account} \times \rho_d(\text{account})))$

