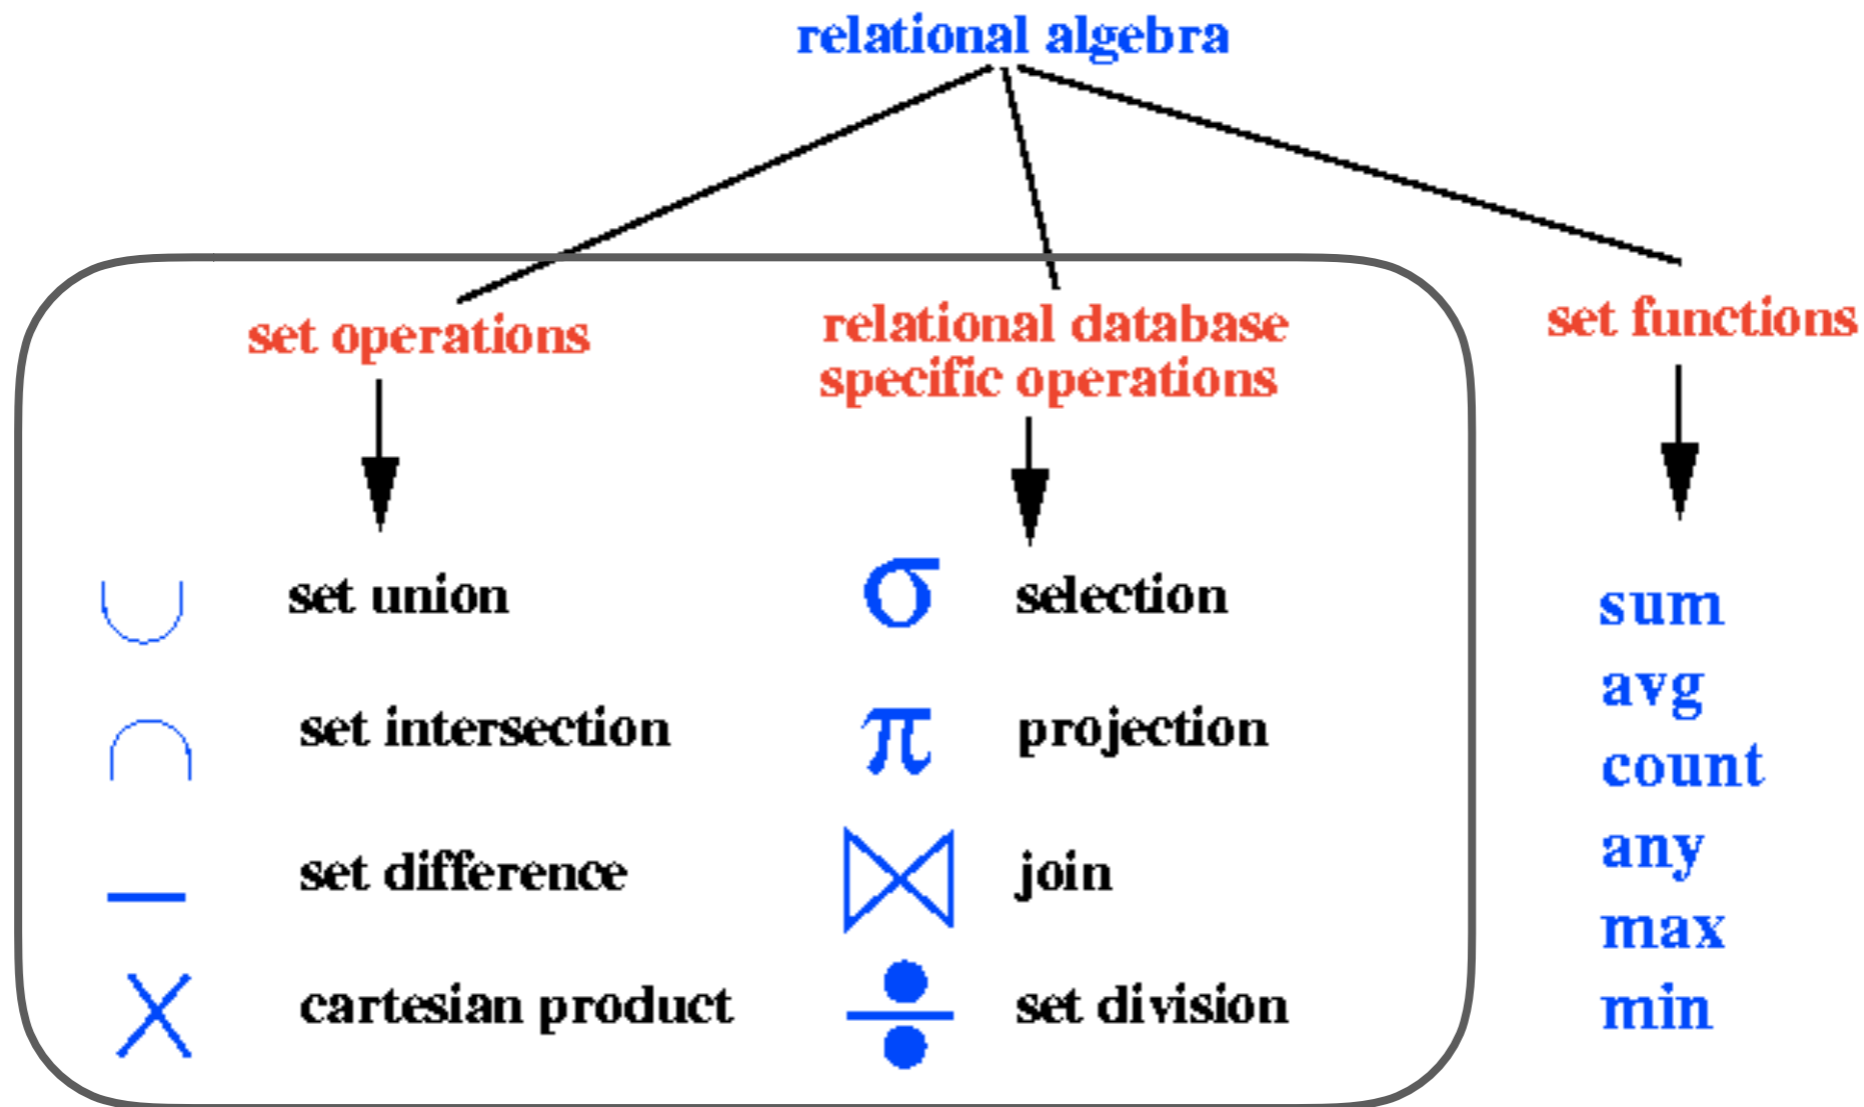


Relational Algebra (II) & Calculus


CS 377: Database Systems

Recap: Relational Algebra Part I



Set (Aggregate) Functions

- Operates on a set of values and produce a single value
- Can also be known as aggregate functions
- Common functions include SUM, AVERAGE, MAXIMUM, MINIMUM, and COUNT



Example: Set Functions

$A = \{1, 4, 5, 10, 15\}$

Function	Description	Value
sum(A)	sum of all values in the (numeric) set	35
avg(A)	average of all values in the (numeric) set	7
max(A)	maximum value of all values in the set	15
min(A)	minimum value of all values in the set	1
any(A)	TRUE if set is not empty, otherwise FALSE	TRUE
count(A)	cardinality (number of elements) of set	5

Additional Operations: Generalized Projection

- Allows functions of attributes to be included in the projection list

$$\pi_{f_1(a_1), f_2(a_2), \dots, f_n(a_n)}(R)$$

- Examples:

$$\pi_{LNAME, FNAME, SALARY * 1.03}(\text{EMPLOYEE})$$

$$\pi_{SSN, FNAME, AGE / 2 + 7, SEX}(\text{EMPLOYEE})$$


Additional Operations: Group By Aggregate

- Groups are formed using one more attribute value(s)
- Aggregate functions applied independently to each group
- Examples:
 - How many people bought an iPad?
 - What is the average age of students in the Database Systems class?
 - What is the average salary of the different departments?

Example: Group By Aggregate

SSN	FName	Other	Sex	DNo	Salary
111-11-1111	John	...	M	4	50,000
242-12-2340	Mary	...	F	5	60,000
222-22-2222	James	...	M	5	80,000
333-33-3333	Jake	...	M	4	60,000

Group by DNO



111-11-1111	John	...	M	4	50,000
333-33-3333	Jake	...	M	4	60,000

avg(salary) =
55,0000

242-12-2340	Mary	...	F	5	60,000
222-22-2222	James	...	M	5	80,000

avg(salary) =
70,0000


Group By Aggregate Operation

- Notation: $a_1, a_2, \dots, a_N \mathcal{F} f_1(a_1), f_2(a_2), \dots, f_M(a_M) (R)$
 - a_1, a_2, \dots, a_N = attributes used to form groups
 - $f_1(a_1), f_2(a_2), \dots, f_M(a_M)$ = set functions applied on each group
- Result is always a relation with the following attributes:
 - Grouping attributes (to differentiate the tuples)
 - Set function values (attributes named after function name)

a1	a2	...	aN	f1	f2	...	fM
----	----	-----	----	----	----	-----	----

Example: Group By Aggregate (2)

SSN	FName	Other	Sex	DNo	Salary
111-11-1111	John	...	M	4	50,000
242-12-2340	Mary	...	F	5	60,000
222-22-2222	James	...	M	5	80,000
333-33-3333	Jake	...	M	4	60,000




$\text{DNO, Sex } \mathcal{F}_{\text{avg}(\text{Salary}), \text{count}(\text{SSN})}(\text{EMPLOYEE})$

DNo	Sex	Avg	Count
4	M	55,000	2
5	M	80,000	1
5	F	60,000	1

No tuple with
DNO=4, Sex='F'
because group
(set) is empty!

Example: Group By Aggregate (3)

SSN	FName	Other	Sex	DNo	Salary
111-11-1111	John	...	M	4	50,000
242-12-2340	Mary	...	F	5	60,000
222-22-2222	James	...	M	5	80,000
333-33-3333	Jake	...	M	4	60,000



$\mathcal{F}_{\text{avg}(\text{Salary}), \text{count}(\text{SSN})}(\text{EMPLOYEE})$


Avg	Count
62,500	4

When no grouping attributes are specified, the set function is applied on ONE group with all the tuples in the relation!

Relational Algebra Operations

Operation	Notation	Purpose
SELECT	$\sigma_{\langle \text{selection condition} \rangle} (R)$	Selects all tuples that satisfy the selection condition from a relation R
PROJECT	$\pi_{\langle \text{attribute list} \rangle} (R)$	New relation with subset of attributes of R and removes duplicate tuples
THETA_JOIN	$R_1 \bowtie_{\langle \text{join condition} \rangle} R_2$	All combinations of tuples from R ₁ and R ₂ that satisfy the join condition
EQUIJOIN	$R_1 \bowtie_{\langle \text{join condition} \rangle} R_2$	Theta join with only equality join comparisons
NATURAL JOIN	$R_1 *_{\langle \text{join condition} \rangle} R_2$	Equijoin except join attributes of R ₂ are not included in the resulting relation
UNION	$R_1 \cup R_2$	Relation that includes all tuples in R ₁ or R ₂
INTERSECTION	$R_1 \cap R_2$	Relation that includes all tuples in both R ₁ and R ₂
DIFFERENCE	$R_1 - R_2$	Relation that includes all tuples in R ₁ that are not in R ₂
CARTESIAN PRODUCT	$R_1 \times R_2$	Relation with attributes of R ₁ and R ₂ and includes tuples with all possible combinations of tuples of R ₁ and R ₂
DIVISION	$R_1(Z) \div R_2(Y)$	Relation that includes all tuples t[X] in R ₁ (Z) that appear in R ₁ in combination with every tuple from R ₂ (Y) where $Z = X \cup Y$

Example: Company Database



Example: RA Queries (1)

Find the name and address of all employees who work in the Research department

Example: RA Queries (2)

Find fname and lname of employees who earn more than 'John Smith'

Example: RA Queries (3)

Find fname and lname of employees who have 2 or more dependents

Example: RA Queries (4)

Find fname and lname of employees who have the most number of dependents

Example: RA Queries (5)

Retrieve the names of employees who have no dependents

Example: RA Queries (6)

List the names of managers who have at least one dependent

Example: RA Queries (7)

Find fname and lname of employees who work on more projects than 'John Smith'

Example: RA Queries (8)

For each department, show the department name, number of employees, minimum employee salary and maximum employee salary

Example: RA Queries (9)

Find fname and lname of all employees who work on 2 or more projects controlled by the Research department

Example: RA Queries (10)

Find fname and lname of all employees who work on all projects controlled by the Research department

Example: RA Queries (1 1)

Find fname and lname of all employees who do not work on any projects controlled by the Research department

Example: RA Queries (12)

Find fname and lname of all employees that only work on projects controlled by the Research department

Relational Calculus

- Declarative query language that describes what is to be retrieved rather than how to retrieve it (nonprocedural)
- Two flavors of relational calculus: Tuple relational calculus (TRC) and Domain relational calculus (DRC)
- Relational calculus and relational algebra are logically equivalent (same logical content)

Relational Calculus

- Calculus has variables, constants, comparison operations, logical connectives, and quantifiers
 - TRC: Variables range over (i.e., get bound to) tuples. Similar to SQL
 - DRC: Variables range over domain elements (field values)
- Both are simple subsets of first-order Logic
- Expression in calculus are called formulas

Tuple Relational Calculus (TRC)

- Tuple variable: a variable name that represents data tuples in the database
 - Typically denoted using a lower case letter
- Range relation: the relation that is the range for a tuple variable
 - Expression $R(t)$ is evaluated as follows:
 - $R(t)$ = true if tuple t is a tuple from the relation R
 - $R(t)$ = false if tuple t is not a tuple from the relation R

TRC

- A query in TRC has the form: $\{t \mid \text{CONDITION}(t)\}$



- Returns all tuples for which the condition or formula evaluates to true
- Formula is recursively defined, starting with simple atomic formulas and building more complex operators using the logical operators

TRC Formula

- An atomic formula is one of the following:
 - $t \in R$
 - $R.a \text{ op } S.b$ $\langle, \rangle =, \geq, \leq, \neq$
 - $R.a \text{ op } \text{constant}$
- A formula can be:
 - An atomic formula
 - NOT p , p AND q , p OR q , where p and q are formulas
 - Special quantifiers

TRC Simple Examples

- $\{t \mid \text{Employee}(t) \text{ AND } t.\text{salary} > 50000\}$
 - Retrieve all tuples t such that t is a tuple of the relation EMPLOYEE and their salary amount is greater than 50000
- $\{t.\text{fname}, t.\text{lname} \mid \text{Employee}(t) \text{ AND } t.\text{salary} > 50000\}$
 - Retrieve the first and last name of employees whose salary is greater than 50000
- $\{t.\text{salary} \mid \text{Employee}(t) \text{ AND } t.\text{fname} = \text{'John'} \text{ AND } t.\text{lname} = \text{'Smith'}\}$
 - Retrieve the salary of the employee “John Smith”

Special Formula Quantifiers

Two special quantifiers can appear in formulas

- Universal quantifier $(\forall t) (\text{Condition}(t))$
evaluates to true if all tuples t satisfies $\text{Condition}(t)$
otherwise false
- Existential quantifier $(\exists t) (\text{Condition}(t))$
evaluates to true if there is some (at least one) tuple t
that satisfies $\text{Condition}(t)$

Free and Bound Variables

- The use of special quantifiers in a formula binds the variable t
 - A variable that is not bound is free
- The variable t that appears to the left of $|$ must be the only free variable in the formula

TRC Example (2)

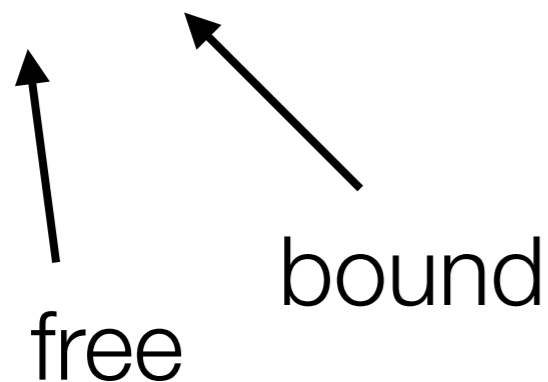
SAILORS (sid, sname, rating, age)

RESERVES (sid, bid, day)

BOATS (bid, bname, color)

- $\pi_{\text{sname}}(\sigma_{\text{rating} > 1}(\text{SAILORS}))$

$\{t \mid (\exists s) (\text{SAILORS}(s) \text{ AND } s.\text{rating} > 1 \text{ AND } t.\text{sname} = s.\text{sname})\}$



CONVENTION: the attributes of the free variable t are exactly the ones mentioned in the formula!

TRC Example (3)

Find the department number of the Research department

- $\{d.dno \mid \text{Department}(d) \text{ AND } d.dname = \text{'Research'}\}$

$$\{d.dno \mid \text{Department}(d) \\ \text{AND } (\\ (\exists t) \\ (\text{Department}(t) \\ \text{AND } t.dname = \text{'Research'} \\ \text{AND } t.dno = d.dno) \\)\}$$

TRC Example (4)

- List the name and address of all employees who work for the 'Research' department

$\{t.Fname, t.Lname, t.Address \mid EMPLOYEE(t) \text{ AND } (\exists d)(DEPARTMENT(d) \text{ AND } d.Dname = \text{'Research'} \text{ AND } d.Dnumber = t.Dno))\}$

- List the names of employees who work on some projects controlled by department number 5

$\{e.fname, e.lname \mid Employee(e) \text{ AND } ((\exists p) (\exists w) (Project(p) \text{ AND } Works_on(w) \text{ AND } p.dnum = 5 \text{ AND } p.pnumber = w.pnum \text{ AND } w.essn = e.ssn)))\}$

TRC Example (4)

Employee(e)

e1, John
e2, Kate
e3, Ann

Works_on(w)

e1, p1
e2, p3
e3, p2

Project(p)

p1, 5
p2, 5
p3, 4

- Run through the employee tuples and make the second condition true, we must find tuples such that p is a Project tuple, w is a Works_on tuple, and it matches the 3 conditions with employee number matching.
 - e1 is good since you can find it in all 3 tables and meets the conditions
 - e2 is problematic because $p3 = 4$, which doesn't match our condition
 - e3 is also output because the combination exists that can make the second condition true

TRC Example (5)

- List the names of employees who work on all the projects controlled by department number 5
- Solution 1: Projects that are either not controlled by department 5 of e is working on

$$\{e.fname, e.lname \mid \text{Employee}(e) \\ \text{AND } ((\forall x) (\text{NOT}(\text{Project}(x)) \\ \text{OR NOT } (x.dnum = 5) \\ \text{OR } ((\exists w) (\text{Works_on}(w) \\ \text{AND } w.essn = e.ssn \\ \text{AND } x.pnumber = w.pno))))\}$$

TRC Example (5)

- List the names of employees who work on all the projects controlled by department number 5
- Solution 2: There is no project controlled by department 5 that e is not working on

$$\{e.fname, e.lname \mid \text{Employee}(e) \\ \text{AND } (\text{NOT}(\exists x)(\text{Project}(x) \\ \text{AND } (x.dnum = 5) \\ \text{AND } (\text{NOT}(\exists w)(\text{Works_On}(w) \\ \text{AND } w.essn = e.ssn \text{ AND } x.pnumber = w.pno))))))\}$$

Relational Algebra & Relational Calculus

- (Definition) Expressive power of a query language is the set of all queries that can be written using that query language
- Query language A is more expressive than query language B if the set of all queries written in A is a superset of all queries that can be written in B
- Codd's Theorem: Every relational algebra query can be expressed as a "safe" query in TRC/DRC; the converse is true
 - Relational Algebra and Relational Calculus are equally expressive

Relational Algebra & Calculus: Recap

- Relational Algebra
 - Set Functions
 - Group By Aggregate
- Relational Calculus

