

Relational Algebra

- By Jyoti Tryambake



Outline

- Relational Algebra the basis of the widely used SQL query language.
- Tuple Relational Calculus
- Domain Relational Calculus

declarative query languages based on mathematical logic



Relational Algebra in DBMS





Relational Algebra

- Relational algebra is the basic set of operations for the relational model.
- These operations enable the user to specify basic retrieval requests (or queries).
- The operators take one or two relations as inputs and produce a new relation as a result.
- Procedural language
- Six basic operators
 - select: σ
 project: Π Unary operations
 rename: ρ
 union: ∪
 set difference: Binary operations
 Cartesian product: x



Select Operation

- Selects tuples that satisfy a given predicate
- **D** Notation: $\sigma_p(r)$
- □ *p* is called the **selection predicate**
- Defined as:

$$\sigma_p(\mathbf{r}) = \{t \mid t \in r \text{ and } p(t)\}$$

Where *p* is a formula in propositional calculus consisting of terms connected by : \land (and), \lor (or), \neg (not) Each term is one of:

<attribute> op <attribute> or <constant>

where *op* is one of: =, \neq , >, \geq . <. \leq



Select Operation

- Examples of selection:
 - Select those tuples of the instructor relation where the instructor is in the "Physics" department

 $\sigma_{dept_name="Physics"}(instructor)$

Select the employee tuples whose salary is greater than \$30,000:

 $\sigma_{SALARY > 30,000}$ (EMPLOYEE)

to find the instructors in Physics with a salary greater than \$90,000

σ_{dept_name ="Physics"∧ salary>90000} (instructor)

comparisons between two attributes, To find all departments whose name is the same as their building name

 $\sigma_{dept_name = building}$ (department)



Project Operation

Notation:

 $\prod_{A_1,A_2,\ldots,A_k}(r)$

where A_1 , A_2 are attribute names and *r* is a relation name.

- This operation keeps certain columns (attributes) from a relation and discards the other columns.
- PROJECT creates a vertical partitioning
 - The list of specified columns (attributes) is kept in each tuple
 - The other attributes in each tuple are discarded
- Example: To eliminate the dept_name attribute of instructor

 $\Pi_{ID, name, salary}$ (instructor)



Combining Select and Project

Find the name of all instructors in the Physics department

 $\Box \prod_{name} (\sigma_{dept name = "Physics"} (instructor))$



(a) $\sigma_{\text{(Dno=4 AND Salary>25000) OR (Dno=5 AND Salary>30000)}}$ (EMPLOYEE).

Results of SELECT and PROJECT operations rightarrow (b) $\pi_{\text{Lname, Fname, Salary}}$ (EMPLOYEE).

(c) $\pi_{\text{Sex, Salary}}$ (EMPLOYEE).

(a)

Fname	Minit	Lname	<u>Ssn</u>	Bdate	Address	Sex	Salary	Super_ssn	Dno
Franklin	Т	Wong	333445555	1955-12-08	638 Voss, Houston, TX	М	40000	888665555	5
Jennifer	S	Wallace	987654321	1941-06-20	291 Berry, Bellaire, TX	F	43000	888665555	4
Ramesh	К	Narayan	666884444	1962-09-15	975 Fire Oak, Humble, TX	М	38000	333445555	5

(b)

Lname	Fname	Salary
Smith	John	30000
Wong	Franklin	40000
Zelaya	Alicia	25000
Wallace	Jennifer	43000
Narayan	Ramesh	38000
English	Joyce	25000
Jabbar	Ahmad	25000
Borg	James	55000

(c)

Sex	Salary
М	30000
М	40000
F	25000
F	43000
М	38000
М	25000
М	55000



- Takes the set of rows in each table and combines them, eliminating duplicates.
- Participating relations must be <u>compatible</u>, i. e. have the same number of columns, and the same column names, domains, and data types.



- **Notation:** $r \cup s$
- Defined as:

 $r \cup s = \{t \mid t \in r \text{ or } t \in s\}$

For $r \cup s$ to be valid.

1. *r*, *s* must have the *same* **arity** (same number of attributes)

2. The attribute domains must be **compatible** (example: 2^{nd} column of *r* deals with the same type of values as does the 2^{nd} column of *s*)



Example 1: Find all courses taught in the Fall 2009 semester, or in the Spring 2010 semester, or in both

 $\Pi_{course_id}(\sigma_{semester="Fall"} \land year=2009 (section))$

$$\Pi_{course_{id}}(\sigma_{semester="Spring" \land year=2010} \\ (section))$$

Example 2: List all cities where there is either a branch office or a property for rent.

□ \prod_{city} (Branch) $\cup \prod_{city}$ (PropertyForRent)



Example 1 output

course <u>_</u> id	sec_id	semester	year	building	room_number	time_slot_id
BIO-101	1	Summer	2009	Painter	514	В
BIO-301	1	Summer	2010	Painter	514	А
CS-101	1	Fall	2009	Packard	101	Н
CS-101	1	Spring	2010	Packard	101	F
CS-190	1	Spring	2009	Taylor	3128	E
CS-190	2	Spring	2009	Taylor	3128	Α
CS-315	1	Spring	2010	Watson	120	D
CS-319	1	Spring	2010	Watson	100	В
CS-319	2	Spring	2010	Taylor	3128	С
CS-347	1	Fall	2009	Taylor	3128	А
EE-181	1	Spring	2009	Taylor	3128	С
FIN-201	1	Spring	2010	Packard	101	В
HIS-351	1	Spring	2010	Painter	514	С
MU-199	1	Spring	2010	Packard	101	D
PHY-101	1	Fall	2009	Watson	100	А

Figure 6.4 The section relation.

course <u>_</u> id
CS-101
CS-315
CS-319
CS-347
FIN-201
HIS-351
MU-199
PHY-101



Set Difference Operation

- Takes the set of rows in the first relation but not the second
- Participating relations must be <u>compatible</u>.
- □ Notation r s
- Defined as:

 $r-s = \{t \mid t \in r \text{ and } t \notin s\}$

- □ Set differences must be taken between **compatible** relations.
 - r and s must have the same arity
 - attribute domains of *r* and *s* must be compatible
- Example: to find all courses taught in the Fall 2009 semester, but not in the Spring 2010 semester

$$\Pi_{course_id}(\sigma_{semester="Fall" \land year=2009}(section)) - \Pi_{course_id}(\sigma_{semester="Spring" \land year=2010}(section))$$



Set Difference Operation

Example 1

course <u>_</u> id	sec_id	semester	year	building	room_number	time_slot_id
BIO-101	1	Summer	2009	Painter	514	В
BIO-301	1	Summer	2010	Painter	514	Α
CS-101	1	Fall	2009	Packard	101	Н
CS-101	1	Spring	2010	Packard	101	F
CS-190	1	Spring	2009	Taylor	3128	Е
CS-190	2	Spring	2009	Taylor	3128	А
CS-315	1	Spring	2010	Watson	120	D
CS-319	1	Spring	2010	Watson	100	В
CS-319	2	Spring	2010	Taylor	3128	С
CS-347	1	Fall	2009	Taylor	3128	Α
EE-181	1	Spring	2009	Taylor	3128	С
FIN-201	1	Spring	2010	Packard	101	В
HIS-351	1	Spring	2010	Painter	514	С
MU-199	1	Spring	2010	Packard	101	D
PHY-101	1	Fall	2009	Watson	100	Α



course <u>i</u> d
CS-347
PHY-101

Figure 6.6 Courses offered in the Fall 2009 semester but not in Spring 2010 semester.



Set Difference Operation

Example 2

X1(Name, Age)

Name	Age
Anoop	22
Saurav	22
Rakesh	20
Pritesh	19

X2(Name, Age)

Name	Age
Anoop	22
Anurag	23
Ganesh	21
Saurav	22
Rakesh	20

Query Output(X1 - X2)

Name	Age
Pritesh	19

Query Output(X2 - X1)

_	_	5	-		
		ν	/	>	
		100			

Name	Age
Anurag	23
Ganesh	21



Set-Intersection Operation

- Takes the set of rows that are common to each relation
- Participating relations must be <u>compatible</u>
- **Notation**: $r \cap s$
- Defined as:
- $\Box \quad r \cap s = \{ t \mid t \in r \text{ and } t \in s \}$
- Assume:
 - □ *r*, s have the same arity
 - □ attributes of *r* and *s* are compatible



Set-Intersection Operation

- Example:
 - Find the set of all courses taught in both the Fall 2009 and the Spring 2010 semesters.

 $\begin{array}{l} \Pi_{course_id} \; (\sigma_{semester = "Fall" \land year = 2009} \; (section)) \cap \\ \Pi_{course_id} \; (\sigma_{semester = "Spring" \land year = 2010} \; (section)) \end{array}$

course <u>_</u> id	sec_id	semester	year	building	room_number	time_slot_id
BIO-101	1	Summer	2009	Painter	514	В
BIO-301	1	Summer	2010	Painter	514	А
CS-101	1	Fall	2009	Packard	101	Н
CS-101	1	Spring	2010	Packard	101	F
CS-190	1	Spring	2009	Taylor	3128	E
CS-190	2	Spring	2009	Taylor	3128	А
CS-315	1	Spring	2010	Watson	120	D
CS-319	1	Spring	2010	Watson	100	В
CS-319	2	Spring	2010	Taylor	3128	С
CS-347	1	Fall	2009	Taylor	3128	Α
EE-181	1	Spring	2009	Taylor	3128	С
FIN-201	1	Spring	2010	Packard	101	В
HIS-351	1	Spring	2010	Painter	514	С
MU-199	1	Spring	2010	Packard	101	D
PHY-101	1	Fall	2009	Watson	100	А

Figure 6.4 The section relation.





- It defines a relation that is the concatenation of every tuple of relation R with every tuple of relation S.
- □ Notation *r* x s
- Defined as:

 $r \ge s = \{t \ q \mid t \in r \text{ and } q \in s\}$

- □ Assume that attributes of r(R) and s(S) are disjoint. (That is, $R \cap S = \emptyset$).
- If attributes of r(R) and s(S) are not disjoint, then renaming must be used.





- Example find the names of all instructors in the Physics department together with the course id of all courses they taught
 - need the information in both the instructor relation and the teaches relation

inst.ID	name	dept <u>_</u> name	salary	teaches.ID	course_id	sec_id	semester	year
22222	Einstein	Physics	95000	10101	CS-437	1	Fall	2009
22222	Einstein	Physics	95000	10101	CS-315	1	Spring	2010
22222	Einstein	Physics	95000	12121	FIN-201	1	Spring	2010
22222	Einstein	Physics	95000	15151	MU-199	1	Spring	2010
22222	Einstein	Physics	95000	22222	PHY-101	1	Fall	2009
22222	Einstein	Physics	95000	32343	HIS-351	1	Spring	2010
33456	Gold	Physics	87000	10101	CS-437	1	Fall	2009
33456	Gold	Physics	87000	10101	CS-315	1	Spring	2010
33456	Gold	Physics	87000	12121	FIN-201	1	Spring	2010
33456	Gold	Physics	87000	15151	MU-199	1	Spring	2010
33456	Gold	Physics	87000	22222	PHY-101	1	Fall	2009
33456	Gold	Physics	87000	32343	HIS-351	1	Spring	2010

 $\sigma_{dept_name = "Physics"}(instructor \times teaches)$

20



 $\sigma_{instructor.ID = teaches.ID}(\sigma_{dept_name = "Physics"}(instructor \times teaches))$

The names of all instructors in the Physics department together with the course id of all courses they taught

 $\Pi_{name, \ course_id} \ (\sigma_{instructor.ID = teaches.ID}(\sigma_{dept_name = "Physics"}(instructor \ \times \ teaches)))$

Instructor Gold is in the Physics department, he does not teach any course (as recorded in the *teaches relation*), and therefore does not appear in the result.

name	course_id
Einstein	PHY-101



□ Are following expressions are equivalent:

 $\Pi_{name, \ course_id} \ (\sigma_{instructor.ID = teaches.ID}(\sigma_{dept_name = "Physics"}(instructor \ \times \ teaches)))$

or

 $\Pi_{name, \ course_id} (\sigma_{instructor.ID = teaches.ID}((\sigma_{dept_name = "Physics"}(instructor)) \times \ teaches))$



Division Operation

Division operator A+B can be applied if and only if:

- Attributes of B is proper subset of Attributes of A.
- The relation returned by division operator will have
 attributes = (All attributes of A All Attributes of B)
- The relation returned by division operator will return those tuples from relation A which are associated to every B's tuple.



Division Operation

- Suited to queries that include the phrase "for all".
- Let r and s be relations on schemas R and S respectively

 $r \div s$

•
$$R = (A_1, ..., A_m, B_1, ..., B_n)$$

•
$$S = (B_1, ..., B_n)$$

The result of r ÷ s is a relation on schema

$$R - S = (A_1, ..., A_m)$$

$$r \div s = \{ t \mid t \in \prod_{R - S} (r) \land \forall u \in s (tu \in r) \}$$

* **u** representing any tuple in s

Where *tu* means the concatenation of a tuple *t* and *u* to produce a single tuple

* for every tuple in R-S (called t), there are a set of tuples in R, such that for all tuples (such as u) in s, the tu is a tuple in R.

Division Operation – Example

n Relations r, s:



Α

α



e.g.

A is customer name

B is branch-name

1and 2 here show two specific branchnames

(Find customers who have an account in all branches of the bank)

n *r÷s*:



Examples of Division A/B pno sno pno B1 p2 s1p1 p2 s1sno p3 s1s1p4 s2s1s3 s2p1 s4s2p2 s3p2 A/B1 Which have p2 s4p2 in A p4 s4



- Allows us to name, and therefore to refer to, the results of relationalalgebra 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

$$\rho_{x(A_1,A_2,...,A_n)}(E)$$

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



- Example Find the highest salary in the university
- □ Strategy is to
 - (1) compute first a temporary relation consisting of those salaries that are *not the largest*
 - (2) take the set difference between the relation salary
 (instructor) and the temporary relation just computed, to
 obtain the result.



- Example Find the highest salary in the university
- Strategy is to
 - (1) compute first a temporary relation consisting of those salaries that are *not the largest*
 - computing the Cartesian product instructor × instructor and forming a selection to compare the value of any two salaries appearing in one tuple
 - the rename operation to rename one reference to the instructor relation; to reference the relation twice without ambiguity
 - write the temporary relation that consists of the salaries that are not the largest:

 $\Pi_{instructor.salary}$ ($\sigma_{instructor.salary < d.salary}$ (instructor × ρ_d (instructor)))



- Example Find the highest salary in the university
- Strategy is to
 - (1) compute first a temporary relation consisting of those salaries that are *not the largest*

ID	name	salary
10101	Srinivasan	65000
12121	Wu	90000
15151	Mozart	40000
22222	Einstein	95000
32343	El Said	60000
33456	Gold	87000
45565	Katz	75000
58583	Califieri	62000
76543	Singh	80000
76766	Crick	72000
83821	Brandt	92000
98345	Kim	80000

salary
65000
90000
40000
60000
87000
75000
62000
72000
80000
92000

Database System Concepts - 6th Edition



- Example Find the highest salary in the university
- Strategy is to
 - (2) take the set difference between the relation salary (instructor) and the temporary relation just computed, to obtain the result.
 - The query to find the largest salary in the university can be written as:

 $\Pi_{salary} (instructor) - \\ \Pi_{instructor.salary} (\sigma_{instructor.salary < d.salary} (instructor \times \rho_d (instructor)))$

salary
95000

Figure 6.12 Highest salary in the university.



Formal Definition of the Relational Algebra

- A basic expression in the relational algebra consists of either one of the following:
 - □ A relation in the database
 - A constant relation
- □ Let E_1 and E_2 be relational-algebra expressions; the following are all relational-algebra expressions:
 - $\Box E_1 \cup E_2$
 - $\Box E_1 E_2$
 - $\Box E_1 \times E_2$
 - $\Box \sigma_p(E_1), P \text{ is a predicate on attributes in } E_1$
 - \square $\prod_{s}(E_{1})$, S is a list consisting of some of the attributes in E_{1}
 - $\rho_{x}(E_{1})$, x is the new name for the result of E_{1}



Natural Join

- The natural join is a binary operation that allows us to combine certain selections and a Cartesian product into one operation.
- \square It is denoted by the **join symbol** \bowtie
- Example: Find the names of all instructors together with the course id of all courses they taught

 $\Pi_{name, course_id}$ (instructor \bowtie teaches)

the schemas for instructor and teaches have the attribute ID in common, the natural-join operation considers only pairs of tuples that have the same value on ID. *name*

name	course_id
Srinivasan	CS-101
Srinivasan	CS-315
Srinivasan	CS-347
Wu	FIN-201
Mozart	MU-199
Einstein	PHY-101
El Said	HIS-351
Katz	CS-101
Katz	CS-319
Crick	BIO-101
Crick	BIO-301
Brandt	CS-190
Brandt	CS-319
Kim	EE-181



Natural Join – Formal Definition

□ The natural join of r and s, denoted by r s, is a relation on schema R U S formally defined as follows:

 $r \bowtie s = \prod_{R \cup S} (\sigma_{r,A_1=s,A_1 \land r,A_2=s,A_2 \land \dots \land r,A_n=s,A_n} (r \times s))$

- □ if r (R) and s(S) are relations without any attributes in common, that is, $R \cap S = \emptyset$, then r \bowtie s = r × s.
- Example Find the names of all instructors in the Comp. Sci. department together with the course titles of all the courses that the instructors teach

 $\Pi_{name,title} (\sigma_{dept_name} = "Comp. Sci." (instructor \bowtie teaches \bowtie course))$

name	title	
Brandt	Game Design	1
Brandt	Image Processing	I
Katz	Image Processing	l
Katz	Intro. to Computer Science	
Srinivasan	Intro. to Computer Science	
Srinivasan	Robotics	
Srinivasan	Database System Concepts	



Natural Join

□ the natural join is **associative**

(instructor \bowtie teaches) \bowtie course instructor \bowtie (teaches \bowtie course)

- The theta join operation is a variant of the natural-join operation that allows us to combine a selection and a Cartesian product into a single operation
- Consider relations r(R) and s(S), and let be a predicate on attributes in the schema $R \cup S$.
- □ The theta join operation is defined as follows:

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



Theta join example

List the names and comments of all clients who have viewed a property for rent.

(Π_{clientNo, fName, IName}(Client)) (Π_{clientNo, propertyNo, comment}(Viewing))

client.clientNo	fName	IName	Viewing.clientNo	propertyNo	comment
CR76	John	Kay	CR76	PG4	too remote
CR56	Aline	Stewart	CR56	PA14	too small
CR56	Aline	Stewart	CR56	PG4	
CR56	Aline	Stewart	CR56	PG36	
CR62	Mary	Tregear	CR62	PA14	no dining room



The Assignment Operation

□ The assignment operation, denoted by ←, works like assignment in a programming language.

$$\Pi_{R \cup S} \left(\sigma_{r.A_1 = s.A_1 \wedge r.A_2 = s.A_2 \wedge \ldots \wedge r.A_n = s.A_n} \left(r \times s \right) \right)$$

Written as;

$$\begin{array}{l} temp1 \ \leftarrow \ R \times \ S \\ temp2 \ \leftarrow \ \sigma_{r.A_1 = s.A_1 \wedge r.A_2 = s.A_2 \wedge \ldots \wedge r.A_n = s.A_n} \ (temp1) \\ result \ = \ \Pi_{R \ \cup \ S} \ (temp2) \end{array}$$

□ Just a convenient way to express complex queries.



Outer Join

- The outer-join operation is an extension of the join operation to deal with missing information
- It preserves those tuples that would be lost in an join by creating tuples in the result containing null values
- □ Types-
 - Left outer
 - Right outer
 - Full outer



Left Outer Join

All rows from R are retained and unmatched rows of S are padded with NULL

ID	name	dept_name	salary	course_id	sec_id	semester	year
10101	Srinivasan	Comp. Sci.	65000	CS-101	1	Fall	2009
10101	Srinivasan	Comp. Sci.	65000	CS-315	1	Spring	2010
10101	Srinivasan	Comp. Sci.	65000	CS-347	1	Fall	2009
12121	Wu	Finance	90000	FIN-201	1	Spring	2010
15151	Mozart	Music	40000	MU-199	1	Spring	2010
22222	Einstein	Physics	95000	PHY-101	1	Fall	2009
32343	El Said	History	60000	HIS-351	1	Spring	2010
33456	Gold	Physics	87000	null	null	null	null
45565	Katz	Comp. Sci.	75000	CS-101	1	Spring	2010
45565	Katz	Comp. Sci.	75000	CS-319	1	Spring	2010
58583	Califieri	History	62000	null	null	null	null
76543	Singh	Finance	80000	null	null	null	null
76766	Crick	Biology	72000	BIO-101	1	Summer	2009
76766	Crick	Biology	72000	BIO-301	1	Summer	2010
83821	Brandt	Comp. Sci.	92000	CS-190	1	Spring	2009
83821	Brandt	Comp. Sci.	92000	CS-190	2	Spring	2009
83821	Brandt	Comp. Sci.	92000	CS-319	2	Spring	2010
98345	Kim	Elec. Eng.	80000	EE-181	1	Spring	2009

Figure 6.17 Result of *instructor* $\exists \forall$ *teaches*.

39



Right Outer Join

All rows from S are retained and unmatched rows of R are padded with NULL

ID	course_id	sec_id	semester	year	name	dept_name	salary
10101	CS-101	1	Fall	2009	Srinivasan	Comp. Sci.	65000
10101	CS-315	1	Spring	2010	Srinivasan	Comp. Sci.	65000
10101	CS-347	1	Fall	2009	Srinivasan	Comp. Sci.	65000
12121	FIN-201	1	Spring	2010	Wu	Finance	90000
15151	MU-199	1	Spring	2010	Mozart	Music	40000
22222	PHY-101	1	Fall	2009	Einstein	Physics	95000
32343	HIS-351	1	Spring	2010	El Said	History	60000
33456	null	null	null	null	Gold	Physics	87000
45565	CS-101	1	Spring	2010	Katz	Comp. Sci.	75000
45565	CS-319	1	Spring	2010	Katz	Comp. Sci.	75000
58583	null	null	null	null	Califieri	History	62000
76543	null	null	null	null	Singh	Finance	80000
76766	BIO-101	1	Summer	2009	Crick	Biology	72000
76766	BIO-301	1	Summer	2010	Crick	Biology	72000
83821	CS-190	1	Spring	2009	Brandt	Comp. Sci.	92000
83821	CS-190	2	Spring	2009	Brandt	Comp. Sci.	92000
83821	CS-319	2	Spring	2010	Brandt	Comp. Sci.	92000
98345	EE-181	1	Spring	2009	Kim	Elec. Eng.	80000

Figure 6.18 Result of teaches \science instructor.

40



Full Outer Join

full outer join(⊐∕⊂)

- Does both the left and right outer join operations, padding tuples from the left relation that did not match any from the right relation, as well as tuples from the right relation that did not match any from the left relation, and adding them to the result of the join.
- Note: teaches tuples always have matching instructor tuples, and
- □ Will this teaches → instructor give the same result as teaches → instructor ?



Extended Relational-Algebra-Operations

Generalized ProjectionAggregate Functions



Generalized Projection

Extends the projection operation by allowing arithmetic functions to be used in the projection list.

$$\prod_{\mathsf{F1, F2, ..., Fn}} (E)$$

□ *E* is any relational-algebra expression

□ Each of F_1 , F_2 , ..., F_n are arithmetic expressions such as +,-, *, and ÷ on numeric valued attributes, numeric constants, and on expressions that generate a numeric result in the schema of *E*.



Generalized Projection

Examples:

- Following example gives the ID, name, dept name, and the monthly salary of each instructor
 - □ Π_{ID,name,dept name,salary÷12}(instructor)
- Given relation credit-info(customer-name, limit, credit-balance), find how much more each person can spend:

 $\Pi_{\text{customer-name, limit - credit-balance}}$ (credit-info)



Aggregate Functions and Operations

Aggregation function takes a collection of values and returns a single value as a result.

avg: average value
min: minimum value
max: maximum value
sum: sum of values
count: number of values

□ Aggregate operation in relational algebra

G1, G2, ..., Gn $g_{F1(A1), F2(A2),..., Fn(An)}(E)$

- □ *E* is any relational-algebra expression
- $G_1, G_2, ..., G_n$ is a list of attributes on which to group (can be empty)
- □ Each *F*_{*i*} is an aggregate function
- **Each** A_i is an attribute name



□ Relation *r*.

A	В	С
α	α	7
α	β	7
β	β	3
β	β	10

 $g_{sum(c)}(r)$





Database System Concepts



□ Relation *account* grouped by *branch-name*:

branch-name	account-number	balance
Perryridge	A-102	400
Perryridge	A-201	900
Brighton	A-217	750
Brighton	A-215	750
Redwood	A-222	700

branch-name **g** sum(balance) (account)

branch-name	balance
Perrvridae	1300
Drighton	1500
Биушон	1500
Redwood	700
Reamona	100





Aggregate Functions (Cont.)

Result of aggregation does not have a name

- Can use rename operation to give it a name
- For convenience, we permit renaming as part of aggregate operation

branch-name **g** sum(balance) as sum-balance (account)



Modification of the Database

- The content of the database may be modified using the following operations:
 - Deletion
 - Insertion
 - Updating
- All these operations are expressed using the assignment operator.



Deletion

- A delete request is expressed similarly to a query, except instead of displaying tuples to the user, the selected tuples are removed from the database.
- □ A deletion is expressed in relational algebra by:

$r \leftarrow r - E$

where *r* is a relation and *E* is a relational algebra query.







Deletion Examples

Delete all account records in the Perryridge branch.

 $account \leftarrow account - \sigma_{branch-name} = "Perryridge" (account)$

Delete all loan records with amount in the range of 0 to 50

loan \leftarrow loan – σ amount \geq 0 and amount \leq 50 (loan)

Delete all accounts at branches located in Needham (branch-city = Needham).







©Silberschatz, Korth and Sudarshan

Insertion Examples

Insert information in the database specifying that Smith has \$1200 in account A-973 at the Perryridge branch.

account \leftarrow account \cup {("Perryridge", A-973, 1200)} depositor \leftarrow depositor \cup {("Smith", A-973)}

Provide as a gift for all loan customers in the Perryridge branch, a \$200 savings account. Let the loan number serve as the account number for the new savings account.

 $r_1 \leftarrow (\sigma_{branch-name = "Perryridge"} (borrower \bowtie loan))$ account \leftarrow account $\cup \prod_{branch-name, loan-number, 200} (r_1)$

depositor \leftarrow depositor $\cup \prod_{customer-name, loan-number}(r_1)$

borrower

loan-number

Modification of the Database – Insertion

Gift for all loan customers of the Perryridge branch: a \$200 savings account. Let the loan number serve as the account number for the new savings account

insert into account

select loan_number, branch_name, 200
from loan
where branch_name = 'Perryridge'
insert into depositor
select customer_name, loan_number
from loan, borrower
where branch_name = 'Perryridge'
and loan.account number = borrower.account number

©Silberschatz, Korth and Sudars

Updating

- A mechanism to change a value in a tuple without charging all values in the tuple
- Use the generalized projection operator to do this task

$$r \leftarrow \prod_{F1, F2, \dots, Fl,} (r)$$

- **Each** F_i is either
 - \Box the *i*th attribute of *r*, if the *i*th attribute is not updated, or,
 - □ if the attribute is to be updated F_i is an expression, involving only constants and the attributes of r, which gives the new value for the attribute

Update Examples

□ Make interest payments by increasing all balances by 5 percent.

 $account \leftarrow \prod_{AN, BN, BAL * 1.05} (account)$

where AN, BN and BAL stand for account-number, branch-name and balance, respectively.

Pay all accounts with balances over \$10,000 6 percent interest and pay all others 5 percent

- In some cases, it is not desirable for all users to see the entire logical model (i.e., all the actual relations stored in the database.)
- Consider a person who needs to know a customer's loan number but has no need to see the loan amount. This person should see a relation described, in the relational algebra, by

 $\Pi_{customer-name, loan-number}$ (borrower loan)

Any relation that is not of the conceptual model but is made visible to a user as a "virtual relation" is called a view.

Consider the view (named *all-customer*) consisting of branches and their customers.

create view all-customer as

 $\Pi_{branch-name, customer-name} (depositor \bowtie account) \\ \cup \Pi_{branch-name, customer-name} (borrower \bowtie loan)$

Assume the following relations:
 BOOKS(DocId, Title, Publisher, Year)
 STUDENTS(StId, StName, Major, Age)
 AUTHORS(AName, Address)
 borrows(DocId, StId, Date)
 has-written(DocId, AName)
 describes(DocId, Keyword)

- 1. List the year and title of each book.
- 2. List all information about students whose major is CS.
- 3. List all books published by McGraw-Hill before 1990.
- 4. List the name of those authors who are living in Davis.
- 5. List the names of all students who have borrowed a book and who are CS majors.

Assume the following relations:
 BOOKS(Docld, Title, Publisher, Year)
 STUDENTS(StId, StName, Major, Age)
 AUTHORS(AName, Address)
 borrows(Docld, StId, Date)
 has-written(Docld, AName)
 describes(Docld, Keyword)

6. List the title of books written by the author 'Silberschatz'.

7. List the title of books written by the author 'Silberschatz' but not books that have the keyword 'database'.

- 8. Find the name of the youngest student.
- 9. List the title of books written by the author 'Ullman'.
- 10. List the authors of the books the student 'Smith' has borrowed.
- 11. Which books have both keywords 'database' and 'programming'?

Korth