
Relational Algebra

- The Relational Model consists of the elements: relations, which are made up of attributes.

Relational Algebra

- A relation is a set of attributes with values for each attribute such that:
 - ❑ Each attribute value must be a single value only (atomic).
 - ❑ All values for a given attribute must be of the same type (or domain).
 - ❑ Each attribute name must be unique.
 - ❑ The order of attributes is insignificant
 - ❑ No two rows (tuples) in a relation can be identical.
 - ❑ The order of the rows (tuples) is insignificant.

Relational Algebra

- Relational Algebra is a *collection of operations* on Relations.
- Relations are *operands* and the result of an operation is another relation.

Relational Algebra

- Two main collections of relational operators:
 - Set theory operations:
 - Union, Intersection, Difference and Cartesian product.
 - Specific Relational Operations:
 - Selection, Projection, Join, Division.

Set Theoretic Operations

Consider the following relations **R** and **S**.

R

First	Last	Age
Bill	Smith	22
Sally	Green	28
Mary	Keen	23
Tony	Jones	32

S

First	Last	Age
Forrest	Gump	36
Sally	Green	28
DonJuan	DeMarco	27

Union: RUS

- Result: Relation with tuples from R and S with duplicates removed.

First	Last	Age
Bill	Smith	22
Sally	Green	28
Mary	Keen	23
Tony	Jones	32
Forrest	Gump	36
DonJuan	DeMarco	27

Difference: R - S

- Result: Relation with tuples from R but not from S

First	Last	Age
Bill	Smith	22
Mary	Keen	23
Tony	Jones	32

Intersection: $R \cap S$

- Result: Relation with tuples that appear in both R and S.

First	Last	Age
Sally	Green	28

Union Compatible Relations

- Attributes of relations need not be identical to perform union, intersection and difference operations.
- However, they must have the same number of attributes or *arity* and the *domains* for corresponding attributes must be identical.

Union Compatible Relations

- **Domain** is the datatype and size of an attribute.
- The **degree** of relation R is the number of attributes it contains.
- **Definition:** Two relations R and S are *union compatible* if and only if they have the same degree and the domains of the corresponding attributes are the same.

Additional properties

- Union, Intersection and difference operators may only be applied to Union Compatible relations.
- Union and Intersection are commutative operations
- Difference operation is NOT commutative.

Cartesian Product: $R \times S$

- Produce all combinations of tuples from two relations.

R

First	Last	Age
Bill	Smith	22
Mary	Keen	23
Tony	Jones	32

S

Dinner	Dessert
Steak	Ice Cream
Lobster	Cheesecake

Cartesian Product: $R \times S$

- $R \times S$:

First	Last	Age	Dinner	Dessert
Bill	Smith	22	Steak	Ice Cream
Bill	Smith	22	Lobster	Cheesecake
Mary	Keen	23	Steak	Ice Cream
Mary	Keen	23	Lobster	Cheesecake
Tony	Jones	32	Steak	Ice Cream
Tony	Jones	32	Lobster	Cheesecake

Relational Algebra

- Two main collections of relational operators:
 - Set theory operations:
 - Union, Intersection, Difference and Cartesian product.
 - Specific Relational Operations:
 - Selection, Projection, Join, Division.

Selection Operator

- Selection and Projection are *unary* operators.
- The selection operator is sigma: σ
- The selection operation acts like a *filter* on a relation by returning only a certain number of tuples.

Selection Operator

- The resulting relation will have the same degree as the original relation.
- The resulting relation may have fewer tuples than the original relation.
- The tuples to be returned are dependent on a *condition* that is part of the selection operator.

Selection Operator

- $\sigma_C(R)$ Returns only those tuples in R that satisfy condition **C**
- A condition C can be made up of any combination of comparison or logical operators that operate on the attributes of R.
 - Comparison operators:

= < > ≥ ≤ ≠

- Logical operators:

∧ ∨ ¬

∧	T	F
T	T	F
F	F	F

∨	T	F
T	T	T
F	T	F

¬	T	F
	F	T

Selection Examples

- Assume the following relation EMP has the following tuples:

Name	Office	Dept	Rank
Smith	400	CS	Assistant
Jones	220	Econ	Adjunct
Green	160	Econ	Assistant
Brown	420	CS	Associate
Smith	500	Fin	Associate

Selection Examples

Select only those Employees in the CS department:

$\sigma_{\text{Dept} = \text{'CS'}}(\text{EMP})$

Result:

Name	Office	Dept	Rank
Smith	400	CS	Assistant
Brown	420	CS	Associate

Selection Examples

- Select only those Employees with last name Smith who are assistant professors:

$\sigma_{\text{Name} = \text{'Smith'} \wedge \text{Rank} = \text{'Assistant'}} (\text{EMP})$

Result:

Name	Office	Dept	Rank
Smith	400	CS	Assistant

Selection Examples

- Select only those Employees who are either Assistant Professors or in the Economics department:

$\sigma_{\text{Rank} = \text{'Assistant'} \vee \text{Dept} = \text{'Econ'}} (\text{EMP})$

Result:

Name	Office	Dept	Rank
Smith	400	CS	Assistant
Jones	220	Econ	Adjunct
Green	160	Econ	Assistant

Selection Examples

- Select only those Employees who are not in the CS department or Adjuncts:

$\sigma_{\neg (\text{Rank} = \text{'Adjunct'} \vee \text{Dept} = \text{'CS'})} (\text{EMP})$

Result:

Name	Office	Dept	Rank
Green	160	Econ	Assistant
Smith	500	Fin	Associate

Projection Operator

- Projection is also a Unary operator.
- The Projection operator is π
- Projection limits the *attributes* that will be returned from the original relation.
- The general syntax is: $\pi_{\text{attributes}} R$
Where *attributes* is the list of attributes to be displayed and R is the relation.

Projection Operator

- The resulting relation will have the same number of tuples as the original relation (unless there are duplicate tuples produced).
- The degree of the resulting relation may be equal to or less than that of the original relation.

Projection Examples

Project only the names and departments of the employees:

$\pi_{\text{name, dept}}(\text{EMP})$

Results:

Name	Dept
Smith	CS
Jones	Econ
Green	Econ
Brown	CS
Smith	Fin

Combining Selection and Projection

- The selection and projection operators can be combined to perform both operations.
- Show the names of all employees working in the CS department:

$\pi_{\text{name}} \sigma (\text{Dept} = \text{'CS'} (\text{EMP}))$

Results:

Name
Smith
Brown

Combining Selection and Projection

- Show the name and rank of those Employees who are not in the CS department or Adjuncts:

$\pi_{\text{name, rank}} \sigma(\neg (\text{Rank} = \text{'Adjunct'} \vee \text{Dept} = \text{'CS'})) (\text{EMP})$

Results:

Name	Rank
Green	Assistant
Smith	Associate

Aggregate Functions

- We can also apply *Aggregate functions* to attributes and tuples:
 - ❑ SUM
 - ❑ MINIMUM
 - ❑ MAXIMUM
 - ❑ AVERAGE, MEAN, MEDIAN
 - ❑ COUNT

Aggregate Functions

- Assume the relation EMP has the following tuples:

Name	Office	Dept	Salary
Smith	400	CS	45000
Jones	220	Econ	35000
Green	160	Econ	50000
Brown	420	CS	65000
Smith	500	Fin	60000

Aggregate Functions Examples

- Find the minimum Salary: $F_{\text{MIN}}(\text{salary})$ (EMP)
Results:

MIN(salary)
35000

Aggregate Functions Examples

- Find the average Salary: $F_{AVG(salary)}(EMP)$

Results:

AVG(salary)
51000

Aggregate Functions Examples

- Count the number of employees in the CS department: $F_{\text{COUNT}(\text{name})} \sigma (\text{Dept} = \text{'CS'} (\text{EMP}))$
Results:

COUNT(name)
2

Aggregate Functions Examples

- Find the total payroll for the Economics department: $F_{\text{SUM}(\text{salary})} \sigma_{(\text{Dept} = \text{'Econ'})}(\text{EMP})$
Results:

SUM(salary)
85000

Join Operation

- Join operations bring together two relations and combine their attributes and tuples in a specific fashion.
- The generic join operator (called the *Theta Join* is: \bowtie
- It takes as arguments the attributes from the two relations that are to be joined.

Join Operation

For example assume we have the EMP relation as above and a separate DEPART relation with (Dept, MainOffice, Phone) :

EMP \bowtie EMP.Dept = DEPART.Dept DEPART

- The join condition can be = < > ≥ ≤ ≠
- When the join condition operator is = then we call this an *Equijoin*
- Note that the attributes in common are repeated.

Join Examples

- Assume we have the EMP relation from above and the following DEPART relation:

Name	Office	Dept	Salary
Smith	400	CS	45000
Jones	220	Econ	35000
Green	160	Econ	50000
Brown	420	CS	65000
Smith	500	Fin	60000

Dept	MainOffice	Phone
CS	404	555-1212
Econ	200	555-1234
Fin	501	555-4321
Hist	100	555-9876

Join Examples

- Find all information on every employee including their department info:

EMP ⋈_{emp.Dept = depart.Dept} DEPART

Join Examples

- **EMP** ⋈ **DEPART**
emp.Dept = depart.Dept

Name	Office	EMP.Dept	Salary	DEPART.Dept	MainOffice	Phone
Smith	400	CS	45000	CS	404	555-1212
Jones	220	Econ	35000	Econ	200	555-1234
Green	160	Econ	50000	Econ	200	555-1234
Brown	420	CS	65000	CS	404	555-1212
Smith	500	Fin	60000	Fin	501	555-4321

Join Examples

- Find all information on every employee including their department info, where the employee works in an office numbered less than the department main office:

$EMP \bowtie_{(emp.office < depart.mainoffice) \wedge (emp.dept = depart.dept)} DEPART$

Join Examples

- **EMP** ⋈_{(emp.office < depart.mainoffice) ∧ (emp.dept = depart.dept)} **DEPART**

Name	Office	EMP.Dept	Salary	DEPART.Dept	MainOffice	Phone
Smith	400	CS	45000	CS	404	555-1212
Green	160	Econ	50000	Econ	200	555-1234
Smith	500	Fin	60000	Fin	501	555-4321

Natural Join

- Notice in the generic (Theta) join operation, any attributes in common (such as dept above) are repeated.
- The *Natural Join* operation removes these duplicate attributes.
- The natural join operator is: *
- We can also assume using * that the join condition will be = on the two attributes in common.

Natural Join Example

- Example: EMP * DEPART

Results:

Name	Office	Dept	Salary	MainOffice	Phone
Smith	400	CS	45000	404	555-1212
Jones	220	Econ	35000	200	555-1234
Green	160	Econ	50000	200	555-1234
Brown	420	CS	65000	404	555-1212
Smith	500	Fin	60000	501	555-4321

Outer Join

- Often in joining two relations, a tuple in one relation does not have a matching tuple in the other relation: there is no matching value in the join attributes.
- To display rows in the result that do not have matching values in the join column, use Outer join.
- Types of outer joins:
 - Left Outer Join
 - Right Outer Join
 - Full Outer Join

Left Outer Join

$R \bowtie S$

- (Left) outer join is a join in which tuples from R that do not have matching values in common columns of S are also included in result relation.
- Missing values in the second relation are set to null.
- The advantage of an Outer join is that information is preserved, that is, the Outer join preserves tuples that would have been lost by other types of tuples.

Outer Join Examples

- Assume we have two relations:

PEOPLE:

Name	Age	Food
Alice	21	Hamburger
Bill	24	Pizza
Carl	23	Beer
Dina	19	Shrimp

MENU:

Food	Day
Pizza	Monday
Hamburger	Tuesday
Chicken	Wednesday
Pasta	Thursday
Tacos	Friday

Left Outer Join

- **PEOPLE** ⋈_{people.food = menu.food} **MENU**

Name	Age	people.Food	menu.Food	Day
Alice	21	Hamburger	Hamburger	Tuesday
Bill	24	Pizza	Pizza	Monday
Carl	23	Beer	<i>NULL</i>	<i>NULL</i>
Dina	19	Shrimp	<i>NULL</i>	<i>NULL</i>

Right Outer Join

- **PEOPLE** ⋈_{people.food = menu.food} **MENU**

Name	Age	people.Food	menu.Food	Day
Bill	24	Pizza	Pizza	Monday
Alice	21	Hamburger	Hamburger	Tuesday
<i>NULL</i>	<i>NULL</i>	<i>NULL</i>	Chicken	Wednesday
<i>NULL</i>	<i>NULL</i>	<i>NULL</i>	Pasta	Thursday
<i>NULL</i>	<i>NULL</i>	<i>NULL</i>	Tacos	Friday

Full Outer Join

- **PEOPLE** ⋈_{people.food = menu.food} **MENU**

Name	Age	people.Food	menu.Food	Day
Alice	21	Hamburger	Hamburger	Tuesday
Bill	24	Pizza	Pizza	Monday
Carl	23	Beer	<i>NULL</i>	<i>NULL</i>
Dina	19	Shrimp	<i>NULL</i>	<i>NULL</i>
<i>NULL</i>	<i>NULL</i>	<i>NULL</i>	Chicken	Wednesday
<i>NULL</i>	<i>NULL</i>	<i>NULL</i>	Pasta	Thursday
<i>NULL</i>	<i>NULL</i>	<i>NULL</i>	Tacos	Friday

Relational algebra and SQL

- SELECT statement

SELECT [DISTINCT | ALL]

{* | [columnExpression [AS newName]] [,...]}

FROM TableName [alias] [, ...]

[WHERE condition]

[GROUP BY columnList] [HAVING condition]

[ORDER BY columnList]

SELECT Statement

FROM	Specifies table(s) to be used.
WHERE	Filters rows.
GROUP BY	Forms groups of rows with same column value.
HAVING	Filters groups subject to some condition.
SELECT	Specifies which columns are to appear in output.
ORDER BY	Specifies the order of the output.

Relational algebra and SQL

- Projection
- Example: The table **E** (for **EMPLOYEE**)

nr	name	salary
1	John	100
5	Sarah	300
7	Tom	100

Relational algebra and SQL - Projection

SQL	Result	Relational algebra								
<code>select distinct salary from E</code>	<table border="1"><thead><tr><th>salary</th></tr></thead><tbody><tr><td>100</td></tr><tr><td>300</td></tr></tbody></table>	salary	100	300	$\pi_{\text{salary}}(E)$					
salary										
100										
300										
<code>select nr, salary from E</code>	<table border="1"><thead><tr><th>nr</th><th>salary</th></tr></thead><tbody><tr><td>1</td><td>100</td></tr><tr><td>5</td><td>300</td></tr><tr><td>7</td><td>100</td></tr></tbody></table>	nr	salary	1	100	5	300	7	100	$\pi_{\text{nr, salary}}(E)$
nr	salary									
1	100									
5	300									
7	100									

Relational algebra and SQL - Selection

SQL	Result	Relational algebra									
<pre>select * from E where salary < 200</pre>	<table border="1"><thead><tr><th>nr</th><th>name</th><th>salary</th></tr></thead><tbody><tr><td>1</td><td>John</td><td>100</td></tr><tr><td>7</td><td>Tom</td><td>100</td></tr></tbody></table>	nr	name	salary	1	John	100	7	Tom	100	$\sigma_{\text{salary} < 200}(\mathbf{E})$
nr	name	salary									
1	John	100									
7	Tom	100									
<pre>select * from E where salary < 200 and nr >= 7</pre>	<table border="1"><thead><tr><th>nr</th><th>name</th><th>salary</th></tr></thead><tbody><tr><td>7</td><td>Tom</td><td>100</td></tr></tbody></table>	nr	name	salary	7	Tom	100	$\sigma_{\text{salary} < 200 \text{ and } nr \geq 7}(\mathbf{E})$			
nr	name	salary									
7	Tom	100									

Combination of projection and selection

SQL	Result	Relational algebra						
<pre>select name, salary from E where salary < 200</pre>	<table border="1"><thead><tr><th data-bbox="875 691 1048 774">name</th><th data-bbox="1048 691 1240 774">salary</th></tr></thead><tbody><tr><td data-bbox="875 774 1048 857">John</td><td data-bbox="1048 774 1240 857">100</td></tr><tr><td data-bbox="875 857 1048 940">Tom</td><td data-bbox="1048 857 1240 940">100</td></tr></tbody></table>	name	salary	John	100	Tom	100	$\pi_{\text{name, salary}} (\sigma_{\text{salary} < 200}(E))$
name	salary							
John	100							
Tom	100							

Cartesian Product

table **E** (for **EMPLOYEE**)

enr	ename	edept
1	Bill	A
2	Sarah	C
3	John	A

table **D** (for **DEPARTMENT**)

dnr	dname
A	Marketing
B	Sales
C	Legal

Cartesian Product

SQL	Result					Relational algebra																																																		
<pre>select *from E, D</pre>	<table border="1"> <thead> <tr> <th>enr</th> <th>ename</th> <th>edept</th> <th>dnr</th> <th>dname</th> </tr> </thead> <tbody> <tr><td>1</td><td>Bill</td><td>A</td><td>A</td><td>Marketing</td></tr> <tr><td>1</td><td>Bill</td><td>A</td><td>B</td><td>Sales</td></tr> <tr><td>1</td><td>Bill</td><td>A</td><td>C</td><td>Legal</td></tr> <tr><td>2</td><td>Sarah</td><td>C</td><td>A</td><td>Marketing</td></tr> <tr><td>2</td><td>Sarah</td><td>C</td><td>B</td><td>Sales</td></tr> <tr><td>2</td><td>Sarah</td><td>C</td><td>C</td><td>Legal</td></tr> <tr><td>3</td><td>John</td><td>A</td><td>A</td><td>Marketing</td></tr> <tr><td>3</td><td>John</td><td>A</td><td>B</td><td>Sales</td></tr> <tr><td>3</td><td>John</td><td>A</td><td>C</td><td>Legal</td></tr> </tbody> </table>					enr	ename	edept	dnr	dname	1	Bill	A	A	Marketing	1	Bill	A	B	Sales	1	Bill	A	C	Legal	2	Sarah	C	A	Marketing	2	Sarah	C	B	Sales	2	Sarah	C	C	Legal	3	John	A	A	Marketing	3	John	A	B	Sales	3	John	A	C	Legal	$E \times D$
enr	ename	edept	dnr	dname																																																				
1	Bill	A	A	Marketing																																																				
1	Bill	A	B	Sales																																																				
1	Bill	A	C	Legal																																																				
2	Sarah	C	A	Marketing																																																				
2	Sarah	C	B	Sales																																																				
2	Sarah	C	C	Legal																																																				
3	John	A	A	Marketing																																																				
3	John	A	B	Sales																																																				
3	John	A	C	Legal																																																				

Join ("inner join")

SQL	Result					Relational algebra
<pre>select * from E, D where dept = dnr</pre>						$\sigma_{\text{dept} = \text{dnr}} (E \times D)$ <p><i>or, using the equivalent join operation</i></p> $E \bowtie_{\text{dept} = \text{dnr}} D$
	enr	ename	dept	dnr	dname	
	1	Bill	A	A	Marketing	
	2	Sarah	C	C	Legal	
	3	John	A	A	Marketing	

Aggregate functions

- Table **E** (for **EMPLOYEE**)

nr	name	salary	dept
1	John	100	A
5	Sarah	300	C
7	Tom	100	A
12	Anne	null	C

Sum

SQL	Result	Relational algebra		
<pre>select sum(salary) from E</pre>	<table border="1"><tr><td>sum</td></tr><tr><td>500</td></tr></table>	sum	500	$F_{\text{sum}(\text{salary})}(\mathbf{E})$
sum				
500				

- **Count:**
 - ❑ Duplicates are not eliminated.
 - ❑ **Null** values are ignored.

Count

SQL	Result	Relational algebra		
<pre>select count(salary) from E</pre>	<table border="1"><tr><td>count</td></tr><tr><td>3</td></tr></table>	count	3	$F_{\text{count(salary)}}(E)$
count				
3				
<pre>select count(distinct salary)from E</pre>	<table border="1"><tr><td>count</td></tr><tr><td>2</td></tr></table>	count	2	$F_{\text{count(salary)}}(\pi_{\text{salary}}(E))$
count				
2				

Aggregate Functions

- We can calculate aggregates "grouped by" something:

SQL	Result	Relational algebra						
<pre>select sum(salary) from E group by dept</pre>	<table border="1"><thead><tr><th>dept</th><th>sum</th></tr></thead><tbody><tr><td>A</td><td>200</td></tr><tr><td>C</td><td>300</td></tr></tbody></table>	dept	sum	A	200	C	300	$\text{dept}^F_{\text{sum}(\text{salary})}(\text{E})$
dept	sum							
A	200							
C	300							

Aggregate Functions

- Several aggregates simultaneously:

SQL	Result	Relational algebra									
<pre>select sum(salary), count(*) from E group by dept</pre>	<table border="1"><thead><tr><th>dept</th><th>sum</th><th>count</th></tr></thead><tbody><tr><td>A</td><td>200</td><td>2</td></tr><tr><td>C</td><td>300</td><td>1</td></tr></tbody></table>	dept	sum	count	A	200	2	C	300	1	$\text{dept} \overset{F}{\text{sum(salary), count(*)}}(\text{E})$
dept	sum	count									
A	200	2									
C	300	1									

Outer join

- Example: Table **E** (for **EMPLOYEE**); table **D** (for **DEPARTMENT**)

enr	ename	dept
1	Bill	A
2	Sarah	B
3	John	A

dnr	dname
A	Marketing
B	Sales
C	Legal

- List each employee together with the department he or she works at.

Outer join

- What if we want to know the number of employees at each department?

Outer join

SQL	Result	Relational algebra																									
<pre>select * from (E right outer join D on dept = dnr)</pre>	<table border="1"> <thead> <tr> <th>enr</th> <th>ename</th> <th>dept</th> <th>dnr</th> <th>dname</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>Bill</td> <td>A</td> <td>A</td> <td>Marketing</td> </tr> <tr> <td>2</td> <td>Sarah</td> <td>B</td> <td>B</td> <td>Sales</td> </tr> <tr> <td>3</td> <td>John</td> <td>A</td> <td>A</td> <td>Marketing</td> </tr> <tr> <td>null</td> <td>null</td> <td>null</td> <td>C</td> <td>Legal</td> </tr> </tbody> </table>	enr	ename	dept	dnr	dname	1	Bill	A	A	Marketing	2	Sarah	B	B	Sales	3	John	A	A	Marketing	null	null	null	C	Legal	$E \bowtie_{\text{dept} = \text{dnr}} D$
enr	ename	dept	dnr	dname																							
1	Bill	A	A	Marketing																							
2	Sarah	B	B	Sales																							
3	John	A	A	Marketing																							
null	null	null	C	Legal																							

Outer Join

SQL	Result	Relational algebra												
<pre>select dnr, dname, count(*) from (E right outer join D on dept = dnr)group by dnr, dname</pre>	<table border="1"> <thead> <tr> <th>dnr</th> <th>dname</th> <th>count</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>Marketing</td> <td>2</td> </tr> <tr> <td>B</td> <td>Sales</td> <td>1</td> </tr> <tr> <td>C</td> <td>Legal</td> <td>1</td> </tr> </tbody> </table>	dnr	dname	count	A	Marketing	2	B	Sales	1	C	Legal	1	$\text{dnr, dname } \mathbf{F}_{\text{count}(*)}(\mathbf{E} \bowtie_{\text{dept} = \text{dnr}} \mathbf{D})$
dnr	dname	count												
A	Marketing	2												
B	Sales	1												
C	Legal	1												

Outer Join

SQL	Result	Relational algebra												
<pre>select dnr, dname, count(enr) from (E right outer join D on edept = dnr) group by dnr, dname</pre>	<table border="1"> <thead> <tr> <th>dnr</th> <th>dname</th> <th>count</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>Marketing</td> <td>2</td> </tr> <tr> <td>B</td> <td>Sales</td> <td>1</td> </tr> <tr> <td>C</td> <td>Legal</td> <td>0</td> </tr> </tbody> </table>	dnr	dname	count	A	Marketing	2	B	Sales	1	C	Legal	0	$\text{dnr, dname} \mathbf{F}_{\text{count(enr)}}(\text{E} \bowtie \text{D})$ <p style="text-align: center;">dept = dnr</p>
dnr	dname	count												
A	Marketing	2												
B	Sales	1												
C	Legal	0												