


## ENTITY TYPES

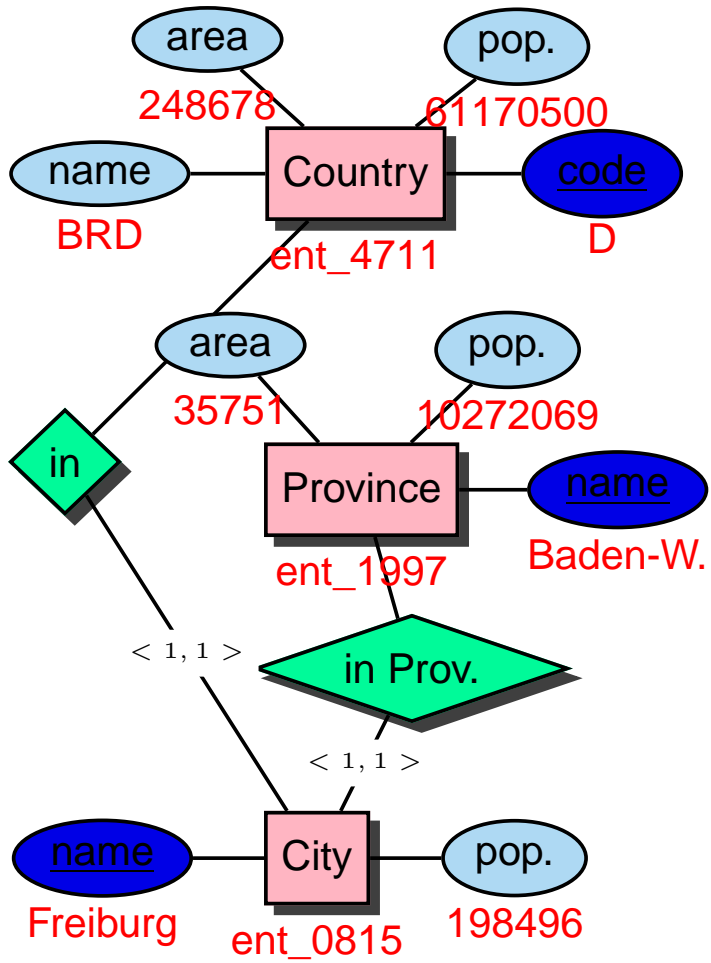
$$(E_{ER}, \{A_1, \dots, A_n\}) \rightarrow E(A_1, \dots, A_n)$$



Continent	
<u>Name</u>	Area
VARCHAR(20)	NUMBER
Europe	9562489.6
Africa	3.02547e+07
Asia	4.50953e+07
America	3.9872e+07
Australia	8503474.56

# WEAK ENTITY TYPES

For weak entity types, the key attributes of the identifying entity type(s) must be added.



City				
<u>Name</u>	<u>Country</u>	<u>Province</u>	Population	...
Freiburg	D	Baden-W.	198496	..
Berlin	D	Berlin	3472009	..
..	..	..	..	..


# RELATIONSHIP TYPES

$(B_{ER}, \{RO_1 : E_1, \dots, RO_k : E_k\}, \{A_1, \dots, A_m\}) \rightarrow$

$B(E_1-K_{11}, \dots, E_1-K_{1p_1}, \dots, E_k-K_{k1}, \dots, E_k-K_{kp_k}, A_1, \dots, A_m)$

where  $\{K_{i1}, \dots, K_{ip_i}\}$  are the primary keys of  $E_i, 1 \leq i \leq k$ .

(it is allowed to rename, e.g., to use *Country* for *Country.Code*)

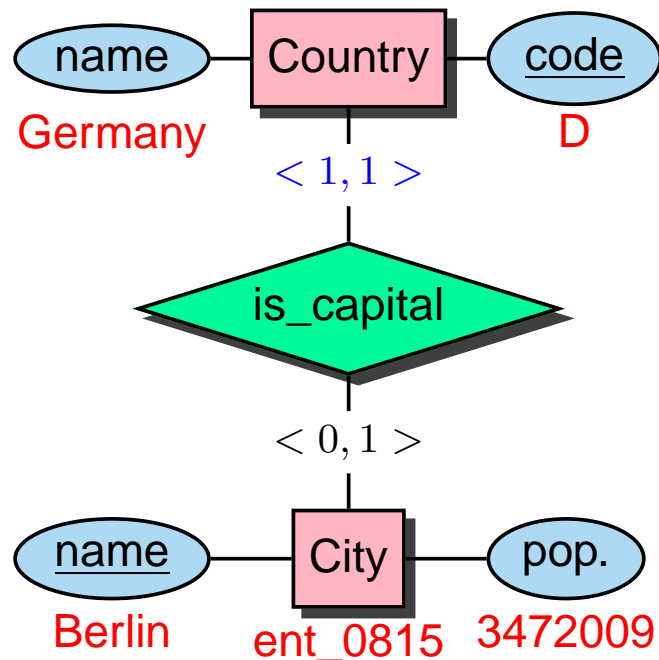


encompasses		
<u>Country</u>	<u>Continent</u>	Percent
VARCHAR(4)	VARCHAR(20)	NUMBER
R	Europe	20
R	Asia	80
D	Europe	100
...	...	...

- Note: for weak entity types, the global key must be used (exercise: is\_capital)

## RELATIONSHIP TYPES: 1:N-RELATIONSHIPS

In case that  $k = 2$  (binary relationship) and a (0,1)- or (1,1)-relationship complexity (i.e., n:1-relations), the relation schema of the relationship type and that of the entity type can be merged (into the relation schema for the entity type)

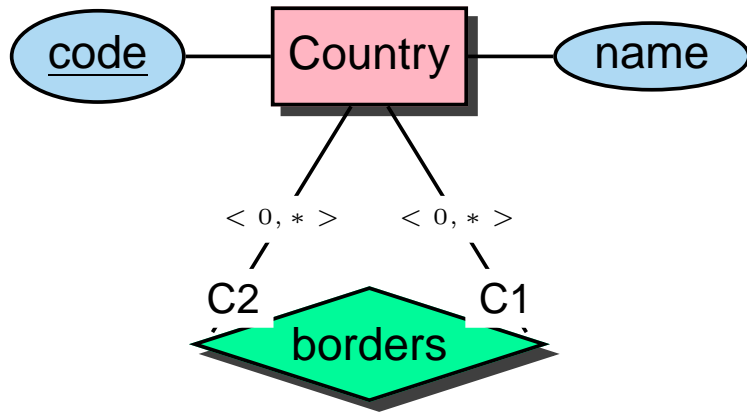


Country					
Name	<u>code</u>	Population	Capital	Province	...
Germany	D	83536115	Berlin	Berlin	..
Austria	A	8023244	Vienna	Vienna	..
Canada	CDN	28820671	Ottawa	Quebec	..
Bolivia	BOL	7165257	La Paz	Bolivia	..
..	..	..	..	..	..

Other examples: flows\_into, headquarters of organizations

# RELATIONSHIP TYPES

In case that for some relationship type, the keys of involved entity types have coinciding names, the role specifications may be used to guarantee the uniqueness of key attributes in the relationship type.



borders	
<u>Country1</u>	<u>Country2</u>
D	F
D	CH
CH	F
..	..

# Relational Algebra

- Recall, the Relational Model consists of the elements: relations, which are made up of attributes.
- A relation is a set of attributes with values for each attribute such that:
  1. Each attribute value must be a single value only (atomic).
  2. All values for a given attribute must be of the same type (or domain).
  3. Each attribute name must be unique.
  4. The order of attributes is insignificant
  5. No two rows (tuples) in a relation can be identical.
  6. The order of the rows (tuples) is insignificant.
- Relational Algebra is a collection of operations on Relations.
- Relations are operands and the result of an operation is another relation.
- Two main collections of relational operators:
  1. Set theory operations:  
Union, Intersection, Difference and Cartesian product.
  2. 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:  $R \cup S$   
Result: Relation with tuples from R and S with duplicates removed.
- Difference:  $R - S$   
Result: Relation with tuples from R but not from S
- Intersection:  $R \cap S$   
Result: Relation with tuples that appear in both R and S.

**R ∪ S**

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

**R - S**

<b>First</b>	<b>Last</b>	<b>Age</b>
Bill	Smith	22
Mary	Keen	23
Tony	Jones	32

**R ∩ S**

<b>First</b>	<b>Last</b>	<b>Age</b>
Sally	Green	28



# Cartesian Product

- Produce all combinations of tuples from two relations.

**R**

<b>First</b>	<b>Last</b>	<b>Age</b>
Bill	Smith	22
Mary	Keen	23
Tony	Jones	32

**S**

<b>Dinner</b>	<b>Dessert</b>
Steak	Ice Cream
Lobster	Cheesecake

**R X S**

<b>First</b>	<b>Last</b>	<b>Age</b>	<b>Dinner</b>	<b>Dessert</b>
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

# Selection Operator

- Selection and Projection are *unary* operators.
- The selection operator is sigma:  $\sigma$
- The selection operation acts like a *filter* on a relation by returning only a certain number of tuples.
- 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.
- $\sigma_C (R)$  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:  $\wedge$   $\vee$   $\neg$
- Use the Truth tables (memorize these) for logical expressions:

$\wedge$	<b>T</b>	<b>F</b>
<b>T</b>	<b>T</b>	<b>F</b>
<b>F</b>	<b>F</b>	<b>F</b>

$\vee$	<b>T</b>	<b>F</b>
<b>T</b>	<b>T</b>	<b>T</b>
<b>F</b>	<b>T</b>	<b>F</b>

$\neg$	<b>T</b>	<b>F</b>
	<b>F</b>	<b>T</b>

## Selection Examples

Assume the following relation EMP has the following tuples:

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

- . Select only those Employees in the CS department:

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

Result:

<b>Name</b>	<b>Office</b>	<b>Dept</b>	<b>Rank</b>
Smith	400	CS	Assistant
Brown	420	CS	Associate

- 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

- 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

- 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  $\pi$ :
- 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.
- 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.

- 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

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

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

Result:

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

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

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

<b>MIN(salary)</b>
35000

- Find the average Salary:  $\mathcal{F}_{\text{AVG}(\text{salary})}(\text{EMP})$   
Results:

<b>AVG(salary)</b>
51000

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

<b>COUNT(name)</b>
2



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

<b>SUM(salary)</b>
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.
- For example assume we have the EMP relation as above and a separate DEPART relation with (Dept, MainOffice, Phone) :  
 $\text{EMP} \bowtie_{\text{EMP.Dept} = \text{DEPART.Dept}} \text{DEPART}$

- 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:

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

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

EMP  $\bowtie$ <sub>emp.Dept = depart.Dept</sub> DEPART

Results:

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

# 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.
- 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

## Underlying Data Model

SQL uses the relational model:

- SQL relations are **multisets (bags)** of tuples (i.e., they can contain duplicates)
- Notions: Relation  $\rightsquigarrow$  Table  
Tuple  $\rightsquigarrow$  Row  
Attribute  $\rightsquigarrow$  Column

The relational algebra serves as theoretical base for SQL as a query language.

## The Language

### Basic structure of SQL queries:

SELECT  $A_1, \dots, A_n$       (... corresponds to  $\pi$  in the algebra)  
FROM  $R_1, \dots, R_m$       (... specifies the contributing relations)  
WHERE  $F$       (... corresponds to  $\sigma$  in the algebra)

corresponds to the algebra expression  $\pi[A_1, \dots, A_n](\sigma[F](r_1 \times \dots \times r_m))$

- Note: cartesian product  $\rightarrow$  aliasing

### General Structure:

SELECT  $A_1, \dots, A_n$       list of attributes  
FROM  $R_1, \dots, R_m$       list of relations  
WHERE  $F$       condition(s)  
GROUP BY  $B_1, \dots, B_k$       list of grouping attributes  
HAVING  $G$       condition on groups  
ORDER BY  $H$       sort order

Subqueries correspond to subtrees in the relational algebra.

## Constructing Queries

For each problem there are multiple possible equivalent queries in SQL (cf. Example 4.15). The choice is a matter of personal taste.

- analyze the problem “systematically”:
  - collect all relations (in the FROM clause) that are needed
  - generate a suitable conjunctive WHERE clause

⇒ leads to a single “broad” SFW query  
(cf. conjunctive queries, relational calculus)
- analyze the problem “top-down”:
  - take the relations that directly contribute to the result in the (outer) FROM clause
  - do all further work in correlated subquery/-queries in the WHERE clause

⇒ leads to a “main” part and nested subproblems
- decomposition of the problem into subproblems:
  - subproblems are solved by nested SFW queries that are combined in the FROM clause of a surrounding query

# Projection

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

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

SQL	Result	Relational algebra								
<pre>select salary from E</pre>	<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	$\pi_{\text{salary}}(\mathbf{E})$					
salary										
100										
300										
<pre>select nr, salary from E</pre>	<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	$\pi_{\text{nr, salary}}(\mathbf{E})$
nr	salary									
1	100									
5	300									
7	100									



# Selection

The same table **E** (for **EMPLOYEE**) as above.

SQL	Result	Relational algebra									
<pre>select * from E where salary &lt; 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	$\sigma_{\text{salary} < 200}(\text{E})$
nr	name	salary									
1	John	100									
7	Tom	100									
<pre>select * from E where salary &lt; 200 and nr &gt;= 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	$\sigma_{\text{salary} < 200 \text{ and nr } \geq 7}(\text{E})$			
nr	name	salary									
7	Tom	100									

## Combination of projection and selection

SQL	Result	Relational algebra						
<pre>select name, salary from E where salary &lt; 200</pre>	<table border="1"> <thead> <tr> <th>name</th> <th>salary</th> </tr> </thead> <tbody> <tr> <td>John</td> <td>100</td> </tr> <tr> <td>Tom</td> <td>100</td> </tr> </tbody> </table>	name	salary	John	100	Tom	100	$\pi_{\text{name, salary}} (\sigma_{\text{salary} < 200}(\text{E}))$
name	salary							
John	100							
Tom	100							

# Cartesian product

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

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

Example: The table **D** (for **DEPARTMENT**)

dnr	dname
A	Marketing
B	Sales
C	Legal

SQL	Result	Relational algebra																																																		
<pre>select * from E, D</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>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	dept	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	<p><b>E X D</b></p>
enr	ename	dept	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																																																