

Chapter 3

Relational algebra and calculus

Query languages for relational databases

- Operations on databases:
 - queries: "read" data from the database
 - updates: change the content of the database
- Both can be modeled as functions from databases to databases
- Foundations can be studied with reference to query languages:
 - relational algebra, a "procedural" language
 - relational calculus, a "declarative" language
 - (briefly) Datalog, a more powerful language
- Then, we will study SQL, a practical language (with declarative and procedural features) for queries and updates

Relational algebra

- A collection of operators that
 - are defined on relations
 - produce relations as resultsand therefore can be combined to form complex expressions
- Operators
 - union, intersection, difference
 - renaming
 - selection
 - projection
 - join (natural join, cartesian product, theta join)

Union, intersection, difference

- Relations are sets, so we can apply set operators
- However, we want the results to be relations (that is, homogeneous sets of tuples)
- Therefore:
 - it is meaningful to apply union, intersection, difference only to pairs of relations defined over the same attributes

Union

Graduates

Number	Surname	Age
7274	Robinson	37
7432	O'Malley	39
9824	Darkes	38

Managers

Number	Surname	Age
9297	O'Malley	56
7432	O'Malley	39
9824	Darkes	38

Graduates \cup Managers

Number	Surname	Age
7274	Robinson	37
7432	O'Malley	39
9824	Darkes	38
9297	O'Malley	56

Intersection

Graduates

Number	Surname	Age
7274	Robinson	37
7432	O'Malley	39
9824	Darkes	38

Managers

Number	Surname	Age
9297	O'Malley	56
7432	O'Malley	39
9824	Darkes	38

Graduates \cap Managers

Number	Surname	Age
7432	O'Malley	39
9824	Darkes	38

Difference

Graduates

Number	Surname	Age
7274	Robinson	37
7432	O'Malley	39
9824	Darkes	38

Managers

Number	Surname	Age
9297	O'Malley	56
7432	O'Malley	39
9824	Darkes	38

Graduates - Managers

Number	Surname	Age
7274	Robinson	37

A meaningful but impossible union

Paternity

Father	Child
Adam	Cain
Adam	Abel
Abraham	Isaac
Abraham	Ishmael

Maternity

Mother	Child
Eve	Cain
Eve	Seth
Sarah	Isaac
Hagar	Ishmael

Paternity \cup Maternity ???

- the problem: Father and Mother are different names, but both represent a "Parent"
- the solution: rename attributes

Renaming

- unary operator
- "changes attribute names" without changing values
- removes the limitations associated with set operators
- notation:
 - $\rho_{Y \leftarrow X}(r)$
- example:
 - $\rho_{\text{Parent} \leftarrow \text{Father}}(\text{Paternity})$
- if there are two or more attributes involved then ordering is meaningful:
 - $\rho_{\text{Location, Pay} \leftarrow \text{Branch, Salary}}(\text{Employees})$

Renaming, example

Paternity

Father	Child
Adam	Cain
Adam	Abel
Abraham	Isaac
Abraham	Ishmael

$\rho_{\text{Parent}} \leftarrow \text{Father.}(\text{Paternity})$

Father	Child
Adam	Cain
Adam	Abel
Abraham	Isaac
Abraham	Ishmael

Renaming and union

Paternity

Father	Child
Adam	Cain
Adam	Abel
Abraham	Isaac
Abraham	Ishmael

Maternity

Mother	Child
Eve	Cain
Eve	Seth
Sarah	Isaac
Hagar	Ishmael

$$\rho_{\text{Parent} \leftarrow \text{Father.}}(\text{Paternity}) \cup \rho_{\text{Parent} \leftarrow \text{Mother.}}(\text{Maternity})$$

Parent	Child
Adam	Cain
Adam	Abel
Abraham	Isaac
Abraham	Ishmael
Eve	Cain
Eve	Seth
Sarah	Isaac
Hagar	Ishmael

Renaming and union, with more attributes

Employees

Surname	Branch	Salary
Patterson	Rome	45
Trumble	London	53

Staff


Surname	Factory	Wages
Patterson	Rome	45
Trumble	London	53

$$\rho_{\text{Location, Pay} \leftarrow \text{Branch, Salary}}(\text{Employees}) \cup \rho_{\text{Location, Pay} \leftarrow \text{Factory, Wages}}(\text{Staff})$$

Surname	Location	Pay
Patterson	Rome	45
Trumble	London	53
Cooke	Chicago	33
Bush	Monza	32

Selection and projection

- Two unary operators, in a sense orthogonal:
 - selection for "horizontal" decompositions
 - projection for "vertical" decompositions



Selection

- Produce results
 - with the same schema as the operand
 - with a subset of the tuples (those that satisfy a condition)
- Notation:
 - $\sigma_F(r)$
- Semantics:
 - $\sigma_F(r) = \{ t \mid t \in r \text{ and } t \text{ satisfies } F \}$

Selection, example

Employees

Surname	FirstName	Age	Salary
Smith	Mary	25	2000
Black	Lucy	40	3000
Verdi	Nico	36	4500
Smith	Mark	40	3900

$\sigma_{\text{Age}<30 \text{ Salary}>4000}$ (Employees)

Surname	FirstName	Age	Salary
Smith	Mary	25	2000
Verdi	Nico	36	4500

Selection, another example

Citizens

Surname	FirstName	PlaceOfBirth	Residence
Smith	Mary	Rome	Milan
Black	Lucy	Rome	Rome
Verdi	Nico	Florence	Florence
Smith	Mark	Naples	Florence

$\sigma_{\text{PlaceOfBirth}=\text{Residence}}$ (Citizens)

Surname	FirstName	PlaceOfBirth	Residence
Black	Lucy	Rome	Rome
Verdi	Nico	Florence	Florence

Projection

- Produce results
 - over a subset of the attributes of the operand
 - with values from all its tuples
- Notation (given a relation $r(X)$ and a subset Y of X):
 - $\pi_Y(r)$
- Semantics:
 - $\pi_Y(r) = \{ t[Y] \mid t \in r \}$

Projection, example

Employees

Surname	FirstName	Department	Head
Smith	Mary	Sales	De Rossi
Black	Lucy	Sales	De Rossi
Verdi	Mary	Personnel	Fox
Smith	Mark	Personnel	Fox

$\pi_{\text{Surname, FirstName}}(\text{Employees})$

Surname	FirstName
Smith	Mary
Black	Lucy
Verdi	Mary
Smith	Mark

Projection, another example

Employees

Surname	FirstName	Department	Head
Smith	Mary	Sales	De Rossi
Black	Lucy	Sales	De Rossi
Verdi	Mary	Personnel	Fox
Smith	Mark	Personnel	Fox

$\pi_{\text{Department, Head}}(\text{Employees})$

Department	Head
Sales	De Rossi
Personnel	Fox

Cardinality of projections

- The result of a projections contains at most as many tuples as the operand
- It can contain fewer, if several tuples "collapse"
- $\pi_Y(r)$ contains as many tuples as r if and only if Y is a superkey for r ;
 - this holds also if Y is "by chance" (not defined as a superkey in the schema, but superkey for the specific instance), see the example

Tuples that collapse

Students

RegNum	Surname	FirstName	BirthDate	DegreeProg
284328	Smith	Luigi	29/04/59	Computing
296328	Smith	John	29/04/59	Computing
587614	Smith	Lucy	01/05/61	Engineering
934856	Black	Lucy	01/05/61	Fine Art
965536	Black	Lucy	05/03/58	Fine Art

$\pi_{\text{Surname, DegreeProg}}(\text{Students})$

Surname	DegreeProg
Smith	Computing
Smith	Engineering
Black	Fine Art

Tuples that do not collapse, "by chance"

Students	RegNum	Surname	FirstName	BirthDate	DegreeProg
	296328	Smith	John	29/04/59	Computing
	587614	Smith	Lucy	01/05/61	Engineering
	934856	Black	Lucy	01/05/61	Fine Art
	965536	Black	Lucy	05/03/58	Engineering

$\pi_{\text{Surname, DegreeProg}}(\text{Students})$

Surname	DegreeProg
Smith	Computing
Smith	Engineering
Black	Fine Art
Black	Engineering

Join

- The most typical operator in relational algebra
- allows to establish connections among data in different relations, taking into advantage the "value-based" nature of the relational model
- Two main versions of the join:
 - "natural" join: takes attribute names into account
 - "theta" join
- They are all denoted by the symbol \bowtie

A natural join

r_1

Employee	Department
Smith	sales
Black	production
White	production

r_2

Department	Head
production	Mori
sales	Brown

$r_1 \bowtie r_2$

Employee	Department	Head
Smith	sales	Brown
Black	production	Mori
White	production	Mori

Natural join: definition

- $r_1(X_1), r_2(X_2)$
- $r_1 \bowtie r_2$ (natural join of r_1 and r_2) is a relation on X_1X_2 (the union of the two sets):

$$\{ t \text{ on } X_1X_2 \mid t[X_1] \in r_1 \text{ and } t[X_2] \in r_2 \}$$

or, equivalently

$$\{ t \text{ on } X_1X_2 \mid \text{exist } t_1 \in r_1 \text{ and } t_2 \in r_2 \text{ with } t[X_1] = t_1 \text{ and } t[X_2] = t_2 \}$$

Natural join: comments

- The tuples in the result are obtained by combining tuples in the operands with equal values on the common attributes
- The common attributes often form a key of one of the operands (remember: references are realized by means of keys, and we join in order to follow references)

Another natural join

Offences

<u>Code</u>	Date	Officer	Dept	Registartion
143256	25/10/1992	567	75	5694 FR
987554	26/10/1992	456	75	5694 FR
987557	26/10/1992	456	75	6544 XY
630876	15/10/1992	456	47	6544 XY
539856	12/10/1992	567	47	6544 XY

Cars

<u>Registration</u>	<u>Dept</u>	Owner	...
6544 XY	75	Cordon Edouard	...
7122 HT	75	Cordon Edouard	...
5694 FR	75	Latour Hortense	...
6544 XY	47	Mimault Bernard	...

Offences \bowtie Cars

<u>Code</u>	Date	Officer	Dept	Registration	Owner	...
143256	25/10/1992	567	75	5694 FR	Latour Hortense	...
987554	26/10/1992	456	75	5694 FR	Latour Hortense	...
987557	26/10/1992	456	75	6544 XY	Cordon Edouard	...
630876	15/10/1992	456	47	6544 XY	Cordon Edouard	...
539856	12/10/1992	567	47	6544 XY	Cordon Edouard	...

Yet another join

- Compare with the union:
 - the same data can be combined in various ways

Paternity

Father	Child
Adam	Cain
Adam	Abel
Abraham	Isaac
Abraham	Ishmael

Maternity

Mother	Child
Eve	Cain
Eve	Seth
Sarah	Isaac
Hagar	Ishmael

Paternity \bowtie Maternity

Father	Child	Mother
Adam	Cain	Eve
Abraham	Isaac	Sarah
Abraham	Ishmael	Hagar

Joins can be "incomplete"

- If a tuple does not have a "counterpart" in the other relation, then it does not contribute to the join ("dangling" tuple)

r ₁	Employee	Department
	Smith	sales
	Black	production
	White	production

r ₂	Department	Head
	production	Mori
	purchasing	Brown

r₁ ⋈ r₂

Employee	Department	Head
Black	production	Mori
White	production	Mori

Joins can be empty

- As an extreme, we might have that no tuple has a counterpart, and all tuples are dangling

r_1

Employee	Department
Smith	sales
Black	production
White	production

r_2

Department	Head
marketing	Mori
purchasing	Brown

$r_1 \bowtie r_2$

Employee	Department	Head

The other extreme

- If each tuple of each operand can be combined with all the tuples of the other, then the join has a cardinality that is the product of the cardinalities of the operands

r_1

Employee	Project
Smith	A
Black	A
White	A

r_2

Project	Head
A	Mori
A	Brown

$r_1 \bowtie r_2$

Employee	Project	Head
Smith	A	Mori
Black	A	Brown
White	A	Mori
Smith	A	Brown
Black	A	Mori
White	A	Brown

How many tuples in a join?

Given $r_1 (X_1)$, $r_2 (X_2)$

- the join has a cardinality between zero and the products of the cardinalities of the operands:

$$0 \leq | r_1 \bowtie r_2 | \leq | r_1 | \times | r_2 |$$

($| r |$ is the cardinality of relation r)

- moreover:
 - if the join is complete, then its cardinality is at least the maximum of $| r_1 |$ and $| r_2 |$
 - if $X_1 \cap X_2$ contains a key for r_2 , then $| r_1 \bowtie r_2 | \leq | r_1 |$
 - if $X_1 \cap X_2$ is the primary key for r_2 , and there is a referential constraint between $X_1 \cap X_2$ in r_1 and such a key, then $| r_1 \bowtie r_2 | = | r_1 |$

Outer joins

- A variant of the join, to keep all pieces of information from the operands
- It "pads with nulls" the tuples that have no counterpart
- Three variants:
 - "left": only tuples of the first operand are padded
 - "right": only tuples of the second operand are padded
 - "full": tuples of both operands are padded

Outer joins

r_1

Employee	Department
Smith	sales
Black	production
White	production

r_2

Department	Head
production	Mori
purchasing	Brown

$r_1 \bowtie_{\text{LEFT}} r_2$

Employee	Department	Head
Smith	Sales	NULL
Black	production	Mori
White	production	Mori

$r_1 \bowtie_{\text{RIGHT}} r_2$

Employee	Department	Head
Black	production	Mori
White	production	Mori
NULL	purchasing	Brown

$r_1 \bowtie_{\text{FULL}} r_2$

Employee	Department	Head
Smith	Sales	NULL
Black	production	Mori
White	production	Mori
NULL	purchasing	Brown

N-ary join

- The natural join is
 - commutative: $r_1 \bowtie r_2 = r_2 \bowtie r_1$
 - associative: $(r_1 \bowtie r_2) \bowtie r_3 = r_1 \bowtie (r_2 \bowtie r_3)$
- Therefore, we can write n-ary joins without ambiguity:

$$r_1 \bowtie r_2 \bowtie \dots \bowtie r_n$$

N-ary join

r_1

Employee	Department
Smith	sales
Black	production
Brown	marketing
White	production

r_2

Department	Division
production	A
marketing	B
purchasing	B

r_2

Division	Head
A	Mori
B	Brown

$r_1 \bowtie r_2 \bowtie r_3$

Employee	Department	Division	Head
Black	production	A	Mori
Brown	marketing	B	Brown
White	production	A	Mori

Cartesian product

- The natural join is defined also when the operands have no attributes in common
- in this case no condition is imposed on tuples, and therefore the result contains tuples obtained by combining the tuples of the operands in all possible ways

Cartesian product: example

Employees

Employee	Project
Smith	A
Black	A
Black	B

Projects

Code	Name
A	Venus
B	Mars

Employees \bowtie Projects

Employee	Project	Code	Name
Smith	A	A	Venus
Black	A	A	Venus
Black	B	A	Venus
Smith	A	B	Mars
Black	A	B	Mars
Black	B	B	Mars

Theta-join

- In most cases, a cartesian product is meaningful only if followed by a selection:
 - **theta-join**: a derived operator

$$r_1 \bowtie_F r_2 = \sigma_F(r_1 \bowtie r_2)$$

- if F is a conjunction of equalities, then we have an **equi-join**

Equi-join: example

Employees

Employee	Project
Smith	A
Black	A
Black	B

Projects

Code	Name
A	Venus
B	Mars

Employees $\bowtie_{\text{Project=Code}}$ Projects

Employee	Project	Code	Name
Smith	A	A	Venus
Black	A	A	Venus
Black	B	B	Mars

Queries

- A query is a function from database instances to relations
- Queries are formulated in relational algebra by means of expressions over relations

A database for the examples

Employees

Number	Name	Age	Salary
101	Mary Smith	34	40
103	Mary Bianchi	23	35
104	Luigi Neri	38	61
105	Nico Bini	44	38
210	Marco Celli	49	60
231	Siro Bisi	50	60
252	Nico Bini	44	70
301	Steve Smith	34	70
375	Mary Smith	50	65

Supervision

Head	Employee
210	101
210	103
210	104
231	105
301	210
301	231
375	252

Example 1

- Find the numbers, names and ages of employees earning more than 40 thousand.

$$\pi_{\text{Number,Name,Age}}(\sigma_{\text{Salary}>40}(\text{EMPLOYEES})) \quad (3.1)$$

Number	Name	Age
104	Luigi Neri	38
210	Marco Celli	49
231	Siro Bisi	50
252	Nico Bini	44
301	Steve Smith	34
375	Mary Smith	50

Example 2

- Find the registration numbers of the supervisors of the employees earning more than 40 thousand

$$\pi_{\text{Head}}(\text{SUPERVISION} \bowtie_{\text{Employee=Number}}(\sigma_{\text{Salary}>40}(\text{EMPLOYEES}))) \quad (3.2)$$

Head
210
301
375

Example 3

- Find the names and salaries of the supervisors of the employees earning more than 40 thousand

$$\pi_{\text{NameH, SalaryH}} (\rho_{\text{NumberH, NameH, SalaryH, AgeH} \leftarrow \text{Number, Name, Salary, Age}} (\text{EMPLOYEES})$$

$$\bowtie_{\text{NumberH=Head}}$$

$$(\text{SUPERVISION} \bowtie_{\text{Employee=Number}} (\sigma_{\text{Salary}>40}(\text{EMPLOYEES})))) \quad (3.3)$$

NameH	SalaryH
Marco Celli	60
Steve Smith	70
Mary Smith	65

Example 4

- Find the employees earning more than their respective supervisors, showing registration numbers, names and salaries of the employees and supervisors

$$\pi_{\text{Number, Name, Salary, NumberH, NameH, SalaryH}} \left(\sigma_{\text{Salary} > \text{SalaryH}} \left(\rho_{\text{NumberH, NameH, SalaryH, AgeH} \leftarrow \text{Number, Name, Salary, Age}} (\text{EMPLOYEES}) \right) \right. \\ \left. \bowtie_{\text{NumberH=Head}} (\text{SUPERVISION} \bowtie_{\text{Employee=Number}} (\text{EMPLOYEES})) \right) \quad (3.4)$$

Number	Name	Salary	NumberH	NameH	SalaryH
104	Luigi Neri	61	210	Marco Celli	60
252	Nico Bini	70	375	Mary Smith	65

Example 5

- Find the registration numbers and names of the supervisors whose employees all earn more than 40 thousand

$$\pi_{\text{Number, Name}}(\text{EMPLOYEES} \bowtie_{\text{Number=Head}} (\pi_{\text{Head}}(\text{SUPERVISION}) - \pi_{\text{Head}}(\text{SUPERVISION} \bowtie_{\text{Employee=Number}} (\sigma_{\text{Salary} \leq 40}(\text{EMPLOYEES})))))) \quad (3.5)$$

Number	Name
301	Steve Smith
375	Mary Smith

Algebra with null values

People

Name	Age	Salary
Aldo	35	15
Andrea	27	21
Maria	NULL	42

- $\sigma_{\text{Age}>30}(\text{People})$
- which tuples belong to the result?
- The first yes, the second no, but the third?

Material on relational calculus and Datalog will be provided in the near future.

Please contact Paolo Atzeni (atzeni@dia.uniroma3.it)
for more information