

# Database Tutorial 3: Relational Algebra and SQL

November 24, 2017

## 1 Repetition

### 1.1 Relational Algebra

$\sigma_C(R)$  Select matching condition C which is a boolean expression

$\pi_A(R)$  Project with attribute list A, can be used for renaming and some computations

$\rho_S(R)$  Rename according to schema  $S = \{S_1 \rightarrow T_1 \dots S_n \rightarrow T_n\}$

$\gamma_{X,G}(R)$  Group according to X and aggregate with G, usually involving renaming

$\tau_A(R)$  Sort according to attributes A

**Set operations** Union, Intersection, Difference, Cross Product

**Joins** Natural Join ( $\bowtie$ )<sup>1</sup>, Conditional Join ( $\bowtie_C$ ), Division ( $\div$ )<sup>2</sup>

### 1.2 SQL

**CREATE** TABLE name (attributes);

**INSERT** INTO table (attributes) VALUES (values);

**UPDATE** table SET name=value WHERE condition;

**DELETE** FROM table WHERE condition;

**SELECT** attributes FROM table;

**Set Operations** UNION, INTERSECT, EXCEPT

### 1.3 Translation

```
SELECT  X, G
FROM    T
WHERE   C       $\tau_Z(\pi_{X,G}(\sigma_D(\gamma_{Y,G}(\sigma_C(T))))))$ 
GROUP BY Y
HAVING D
ORDER BY Z;
```

---

<sup>1</sup> $R \bowtie S = \{r \cup s \mid r \in R \wedge s \in S \wedge \text{Fun}(r \cup s)\}$  with predicate  $\text{Fun}(t)$  true if t is a function

<sup>2</sup> $R \div S = \{t[a_1 \dots a_n] \mid t \in R \wedge \forall s \in S(t[a_i \dots a_n] \cup s \in R)\}$  with  $t[a_1 \dots a_n]$  subset of all attributes in R

## 2 Exercises

- (8 points) (a) (4 points) Given the relation  
`Planets(star, name, position, distance, mass, atmosphere, oxygen, water)`  
write a relational algebra query that returns, for each star with more than 5 planets, the total combined mass of all planets with an atmosphere. The query should return tuples of the form  
`(star, totalMass)`  
(b) (4 points) Given the relations  
`P(star, position, distance, mass, atmosphere, oxygen, water)` and  
`G(star, position, gravity)`,  
translate the following relational algebra query to SQL:  
$$\tau_{maxg}(\pi_{position, maxg}(\gamma_{position, AVG(mass) \rightarrow avgm, MAX(gravity) \rightarrow maxg}(P \bowtie G)))$$

- (8 points) A multi-national company uses a relational database to manage information about its offices in different cities, and its employees. This database has the following relations:

```
Offices(city, supplement)
Departments(city, dname, departmentHead)
Employees(empId, name, salary, dept, city)
```

The company has one office in each city, and several departments can be located at each office. Attribute `supplement` is the monthly salary supplement that each employee working at that office receives (e.g. employees at the London office might receive a supplement of 1000 SEK per month to cover higher living costs in London). The default city supplement is 0 SEK. Attribute `dname` describes the departments function (e.g. sales or personnel). Attribute `departmentHead` is the employee identifier of the head of the department. Employee identifiers (`empId`) are unique. Attribute `salary` is an employees basic monthly salary. The total monthly salary for an employee can be calculated by adding the city supplement to the employees basic monthly salary.

- (4 points) Write a relational algebra expression that finds the employee identifier, name and total monthly salary of all employees (recall that the total monthly salary for an employee can be calculated by adding the city supplement to the employees basic monthly salary). The results should be sorted by employee name.
- (4 points) Write a relational algebra expression that finds the names of cities where there is a sales department (named “sales”) and, for each of these departments, the average basic salary of the employees in that department. You can assume that every department has at least one employee.

Consider the relation `Planets(star, name, distance, mass, atmosphere, oxygen, water)`.

- (12 points) (a) (4 points) Write an SQL table definition with reasonable types and constraints. Store distance in millions of km (For Earth, you would store the value 149.6)  
(b) (4 points) Write an SQL query to determine how many planets are in orbits larger than the orbit of the fictional planet “Duna” of the fictional star “Kerbol”?  
(c) (4 points) We define a planet “habitable” if it satisfies all these conditions:
  - orbits at a distance (in millions of km) between 100 and 200 (inclusive) from its star,
  - has an atmosphere and it has an oxygen percentage between 15% and 25% (inclusive),
  - has water on its surface.

Write an SQL query which returns the star and name of a planet, as well as a column `status` with value “habitable” if the planet is habitable, otherwise “uninhabitable”. (This means, return 3 values per row)

- (10 points) A database system used by a company’s personnel department has the following relations:

Employees(empId, name, year, salary, entitlement, branch)  
ParentalLeave(employee, startDay, startYear, endDay, endYear)

Employee identifiers (empId) are unique. Attribute year is the employee's year of birth. Attribute entitlement is the number of annual vacation days to which the employee is entitled. Employees have 30 days of annual vacation entitlement by default. Branch is the name of the city where the branch is located (assume that there is only one branch in each city). The personnel department keeps a record of all periods of parental leave taken by employees. The attributes startDay and endDay are integers in the range 1-366 that represent the day within the year. For each period of parental leave, the start date must be before the end date.

- (a) (4 points) Suggest keys and references for these relations. Write SQL statements that create these relations with reasonable constraints in PostgreSQL.
- (b) (3 points) Show an SQL query to get the amount of employees per branch that have/had parental leave spanning a period in two different calendar years (this means startYear and endYear are not the same). Example output row if the Stockholm branch has had 3 such employees: ('Stockholm', 3).
- (c) (3 points) Show an SQL query to get the branch name and average of salaries in that branch, for those branches that only have employees born after 1987. Sort the output by average salary. Do not show information about branches that have employees born in or before 1987.