

CS 386/586 Winter 2013 Assignment 4

Assigned: Friday, February 1, 2013

Due: Wednesday, February 13, 2013 at midnight

General Information:

You are strongly encouraged to complete substantial work on Assignment 4 by Wednesday, February 6, 2013 (one week before it is due) so that you can ask questions during class.

Submission: you **must** post a **note** to the “**instructors**” with the **attachment** (listed below) in Piazza. You **must** put your note with the attachment in the **turn-in-assign-here4** folder.

Assignment:

1. Explain what the following queries compute in English.

The queries address a database with this schema:

Student(snum, sname, major, level, age)

Class(name, time, room, fid)

Enrolled(snum, cname)

Faculty(fid, fname, deptid)

In this DB, students are represented by student number, name, major, level (either ‘F’, ‘S’, ‘JR’, ‘SR’), and age. Classes are represented simply by the name (e.g., ‘Calculus I’), the time (an integer representing the hour that the class meets, the room (a character string), and the id of the (single) faculty member who teaches the class. The Enrolled table indicates which students are taking which classes, in the current term. Faculty are represented by their id, name, and the id of their department.

Note: You may find it useful to define these tables in postgresql and insert some data; this would allow you to try out these queries (or pieces of these queries) in order to understand what they do. You are not required to do so; you need only turn in the English description of the queries.

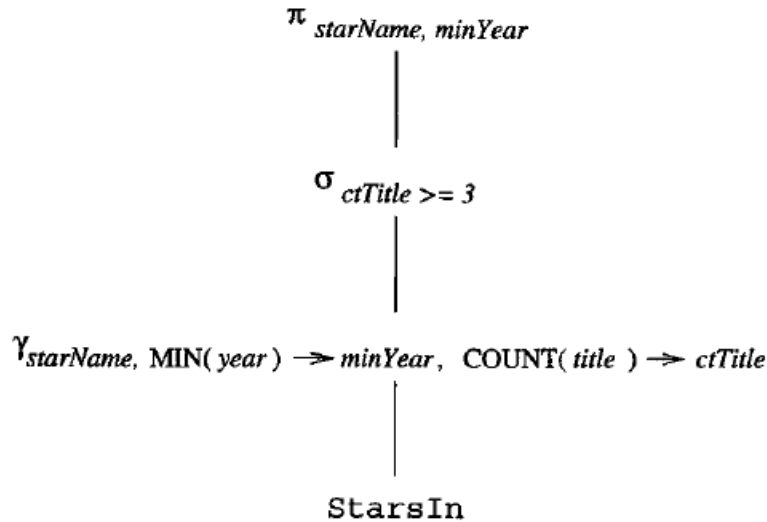
- a.

```
SELECT DISTINCT F.fname
FROM Faculty F
WHERE NOT EXISTS (( SELECT room FROM Class C )
EXCEPT
(SELECT C1.room FROM Class C1
WHERE C1.fid = F.fid ));
```

- b. SELECT DISTINCT S.sname
FROM Student S
WHERE S.snum IN (SELECT E.snum
FROM Enrolled E
GROUP BY E.snum
HAVING COUNT (*) >= ALL (SELECT COUNT (*)
FROM Enrolled E2
GROUP BY E2.snum))
- c. SELECT C.name
FROM Class C
WHERE C.room = 'R128'
OR C.name IN (SELECT E.cname
FROM Enrolled E
GROUP BY E.cname
HAVING COUNT (*) >= 5)
- d. SELECT DISTINCT S.sname
FROM Student S
WHERE S.snum NOT IN (SELECT E.snum
FROM Enrolled E)
- e. SELECT S.level, AVG(S.age)
FROM Student S
WHERE S.level <> 'JR'
GROUP BY S.level

2. Write an SQL query equivalent to the following Algebraic Expression tree.
 StarsIn (title, year, starname)

This table indicates that which movies stars starred in which movies. The title attribute is the title of the movie; year is the year that the movie was produced, and starname is the name of the movie star.



3. For each of the following descriptions, write two different SQL queries. Write each query in two significantly different ways. You should use at least one subquery in your answer (e.g., using different sets of the operators EXISTS, IN, ALL, and ANY). These queries address the Spy database.

Turn in your query; the number of rows in the query answer; and the first ten rows of the query answer. Note that your two queries should be equivalent and thus have identical answers (expect, perhaps, for the order of the rows). So, you only need to turn in the number of rows and the first ten rows once – for each question.

- a. Find the countries whose agents have the maximum salary.
- b. Find the countries where at least one agent speaks English.

4. Write SQL queries that address the following database schema:

Employee(first, last, ssn, gender, salary, super_ssn, dnum)

In this table, an employee is represented by their first and last names, their social security number, their gender, and their salary. The table also indicates the social security number of their supervisor and the number of the department that they are in.

Note: we have provided a csv file of data for this database. Define the above table and load it with the data we provide.

Turn in your SQL query, your query answer (based on the above data), and the number of rows in your query answer.

1. List the names of all employees who work in the department that has the employee with the highest salary among all employees.
2. List the names of all employees whose supervisor's supervisor has ssn = '123456789'
3. List the names of the employees who make at least \$10,000 more than the employee who is paid the least in the company.
4. List the department id of all departments where the number of females is greater than the number of males.