

CS 386/586 Winter 2012 Assignment 2
Assigned: Tuesday, January 17, 2012
Due: Tuesday, January 17, 2012 at 4:40PM

(Given the snow delay/cancelling of class on Jan. 17, 2012 – you can turn this assignment in on Thursday, January 26. But note that Assignment 3 will not be delayed. Assignment 3 will be posted on Tuesday, January 24 and will be due on Tuesday, January 31.)

You must have permission from the instructor, based on a health, family, or work-related emergency, to be able to turn in this assignment later than the start of class on Thursday, January 26.

Turn your papers in, on paper, in class. (If you are taking the class remotely, then you can send a single pdf file as an attachment to an e-mail to the TA yuez@cecs.pdx.edu and copy the instructor imd@cs.pdx.edu. If you attend class, please turn it in on paper.)

This assignment can be done individually or in a team of two students. If you work as a team, turn in one paper with both students' names on the paper.

Note: Do NOT ask questions about queries against the Spy or Library database when you post questions on piazza. Rather, write a similar query against the sailors database that allows you to ask your question. Or ask your question in a generic way – that doesn't mention any particular database.

Part I Extensions to the SELECT clause in SQL

1. (4 points) Write an SQL query that lists the agent_id and salary – in Euros based on a conversion rate available for the day that you do your assignment. (You can find a conversion rate online.) Set the name of the attribute in the final query answer (for the salary in Euros) to be 'Salary-in-Euros'. Show the SQL query, (at most) the first five rows of your query answer, and the number of rows in the query answer.
2. (4 points) Write an SQL query that finds the minimum, maximum, and average salaries for agents that are from Germany. Choose appropriate attribute names for all columns of your query answer. Show the SQL query, (at most) the first five rows of your query answer, and the number of rows in the query answer.

Part II Relational Algebra (select, project, join, cross product) and SQL

For questions 3, 4, and 5 write a query in relational algebra, write and execute an equivalent query in SQL against the introdb_spy database. Turn in: (a) your relational

algebra query, (b) your SQL query, and (c) (at most) the first five rows of your query answer along with the count of rows in your query answer.

3. (7 points) List the agent_id, first, middle, and last names of agents with the team_id and team name for the teams that they are assigned to.

(4 points) Write the relational algebra query a second time. (If your first relational algebra query used the join operator, then the second relational algebra query should NOT use the join operator. Conversely, if your first relational algebra query did not use the join operator, then your second relational algebra query SHOULD use the join operator.) Turn in the relational algebra query.
4. (7 points) List all possible pairs of two distinct languages. (That is, do not list a language with itself.)
5. (7 points) List mission names for missions that have an agent who speaks 'German' and an agent who has the skill of 'Sniper'. (It doesn't need to be the same agent who speaks German and who has the skill of sniper; it could be two agents – or it could be one agent. Either situation is fine; a team with either or both of these possibilities should be included in the query answer.)
6. (4 points) Modify the SQL query in question 3 to a skill of 'Sniper' and a language of 'Greek'. List the SQL query and the first 5 rows of the query answer, and the number of rows returned.
7. (4 points) Modify the SQL query in question 3 to list missions where a single agent has the skill of 'Forgery' and also speaks 'Russian'. Show the SQL query, (at most) the first five rows of your query answer, and the number of rows in the query answer.
8. (14 points) List the SQL queries that you used to convince yourself that the answers you got for question 5, 6, and 7 are correct. Explain your reasoning. (You don't need to give the entire query answer for any of these queries but you do need to give the SQL queries and you need to explain your reasoning.)

Part III Outer Joins – Library DB

9. (20 points) Draw a schema diagram for the introdb_library database – in the same style that you used for the diagram that you drew for the introdb_spy database in assignment 1.
10. (4 points) Write an SQL query for the introdb_library DB that lists the PublisherID for all publishers along with the BookDescriptionID for the book that they've published. Be sure to include all publishers – even if they

haven't published a book in this database.

11. (4 points) Write an SQL query for the `introd_b_library` DB that lists the `BookID` for all books along with the `PersonID` of the person who borrowed the book (as recorded in the `Action` table). Be sure to include all books, whether or not they have ever been borrowed. List each appropriate (`BookID`, `PersonID`) pair just once in your query answer.

Show the SQL query, (at most) the first five rows of the query answer, and the total number of rows in the query answer.

(3 points) Also, list the first three publishers that you see in your query answer that haven't published any books. You should identify these three publishers manually – just by looking at the query answer you get from your query.

12. (2 points) Write your SQL query from question 11 again. If you used `LEFT` join in query 11, then write the query using `RIGHT` join. (If you used `RIGHT` join in query 11, then write the query using `LEFT` join.) Show the SQL query, (at most) the first five rows of the query answer, and the total number of rows.

Part IV GROUP BY with HAVING - SQL Queries – `introd_b_spy` DB

13. (4 points) List the `mission_id` with the count of the number of members on the team associated with the mission, for all missions. Show the SQL query, (at most) the first five rows of your query answer, and the number of rows in the query answer.
14. (4 points) List the `agent_id` with the count of the number of languages that the agent speaks for all agents that speak at least five languages. Make sure that your query answer is sorted into ascending order by language id. Show the SQL query and ALL of the rows in your query answer. (There should be fewer than 30 rows.)
15. (4 points) List the language id plus the count of the number of agents who speak that language, for all languages in the database. Show the SQL query, (at most) the first five rows of your query answer, and the number of rows in the query answer.