Name: _____

1. (1 pt.)
   - **Read all material carefully.**
   - You may refer to your books, papers, and notes during this test.
   - No computer or network access of any kind is allowed (or needed).
   - Write, and draw, carefully. Ambiguous or cryptic answers receive zero credit.
   - Use the conventions used in class and the textbook for all material.
   - COS 480 students should answer non-⋆ questions; ⋆ questions are for extra credit.
   - COS 580 students should answer all questions, including ⋆ questions.

   Write your name in the space provided above.

2. (14 pts.) Recall, from class exercises, the database composed of relations `Students(id, name, year)`, `Courses(id, title, ta)`, and `Enrolls(student, course, credits)`, with the semantics described there. We say student x has TA y if x is enrolled in a class that has y as a TA. Provide a standard SQL (as defined in the textbook) query for the list of all tuples (n, i, m) such that the student with name n and ID i has m distinct TAs. The tuples should be sorted in descending order by m and (secondary order) ascending order by n. You may use views to present your query. *Briefly explain why your query is correct.*

[additional space for answering the earlier question]

3. (15 pts.) Provide an extended bag algebra query that is equivalent to the query of Question 2. You may use all the bag operators defined in the textbook, and may use the linear notation to present your query. *Briefly explain why your query is correct.*

4. (15 pts.) For the database of Question 2, provide a standard SQL query for pairs of students for who are enrolled in identical sets of classes. In more detail, if C_x and C_y denote the classes taken by students x and y, respectively, then the unordered pair (i.e., 2-element set) $\{x, y\}$ qualifies if $C_x = C_y$. The output should list each such unordered pair exactly once, as either (x, y) or (y, x), but not both. *Briefly explain why your query is correct.*

[additional space for answering the earlier question]

5. ⋆ (15 pts.) Either (a) provide, *with explanation*, a relational algebra query that is equivalent to the query of Question 4 using only the operators selection, projection, cross product, union, difference, and renaming, along with a small (constant) number of constant database relations of your choice; or (b) prove that no such query exists. You may use the linear notation for algebra.

[additional space for answering the earlier question]

6. $\star$ (15 pts.) Either (a) provide, *with explanation*, a Datalog query (using any feature described in the textbook and its Web supplement assigned as reading) that is equivalent to the query of Question 4; or (b) prove that no such query exists.