Name: _____

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

   Write your name in the space provided above.

2. (8 pts.) Consider a directed graph with colored edges encoded by a relation `Edge(S, D, color)` with $(s, d, c) \in$ `Edge` iff there is a c-colored edge from vertex s to vertex d. Write a Datalog query for the pairs of vertices (x, y) such that there is a directed path from x to y but there is no alternating red-green path from x to y. An alternating red-green path is a path composed of only red- and green-colored edges with no consecutive edges of the same color. The vertices x and y in this definition need not be distinct. Explain your query briefly and prove that it is safe and stratified.

3. (5 pts.) Provide a SQL equivalent of the query of Question 2.

4. (8 pts.) Let E denote an ER diagram composed of only entity sets (single and double boxes), relationships (single and double diamonds), and attributes on only entity sets. If e, r, and a denote the number of entity sets, relationships, and attributes in E then we may define the *weight* of E to be $2e + 2r + a$.

   What is the minimum weight of such an ER diagram that includes at least one entity set whose key has at least 20 attributes? Justify your answer.

5. (8 pts.) Trace the naive evaluation of the query of Question 2 on the following instance.

Edge		
S	D	color
1	2	red
1	5	green
2	3	green
2	4	red
3	1	red
3	2	blue
3	4	green
4	1	red
5	3	red

6. (10 pts.) $\star$ Repeat Question 5 for semi-naive evaluation.