Name: ____

- COS 480 students should answer non-⋆ questions; optional ⋆ questions are for extra credit.
- COS 580 students should answer all questions, including ⋆ questions.

1. (1 pt.) Write your name in the space provided above.

2. (19 pts.) Consider a database with the following familiar schema, as discussed in class. (Ask for clarifications if the semantics are unclear.)

```
Students(sid, name, year, major)
Courses(cid, title, ta)
Enrolls(sid, cid, credits)
```

Consider the following constraints.

(a) Values for `credits` must be non-negative integers.

(b) Every student ID appearing in the Enrolls table must also appear in the Students table.

(c) Values for the `year` must be two-character codes.

(d) Students cannot TA courses in which they are enrolled.

(e) Students with major '`COS`' cannot be enrolled in any class for fewer than 3 credits.

(f) Students with major '`COS`' cannot have '`NN`' for the `year` attribute.

Provide standard SQL statements that create the tables by selecting the most appropriate types, keeping the constraints in mind. Justify your selections and *explain briefly* how your selections assist in maintaining consistency. Provide standard SQL statements to declare the constraints listed above. *Use the simplest (least powerful) possible SQL constraint type for each.*

[additional space for answering the earlier question]

[additional space for answering the earlier question]

3. (15 pts.) The *division* operator of relational algebra takes operands $R(A, B, C, D)$ and $S(A, B)$, and produces the *quotient*:

$$R \div S \equiv \{(c, d) \mid \forall (a, b) \in S : (a, b, c, d) \in R\}$$

Express $R{\div}S$ using a relational algebra expression composed of *no operators other than* selection, projection, cross product, union, and difference. *Prove* that your expression is equivalent to the above definition.

[additional space for answering the earlier question]

4. (10 pts.) Given a SQL database with tables `R(a,b,c,d)` and `S(a,b)`, provide a standard SQL statement to declare the constraint that $\texttt{R} \div \texttt{S}$ must be empty. (The tables are assumed to exist and to be populated.)

5. (10 pts.) $\star$ Define an extended bag algebra operator η that, intuitively, is similar to the division operator of Question 3 but that instead of "for all tuples in S" uses "for at least half the tuples in S." In more detail, $R \eta S$ contains tuples (c, d) such that there are at least $\lceil |S|/2 \rceil$ tuples $(a, b) \in S$ such that $(a, b, c, d) \in R$. Duplicates should be preserved from R to $R \eta S$ and counted separately in S.

   Prove or disprove: $R \eta S$ may be expressed using standard extended bag algebra (as defined in the textbook and in class).

6. (5 pts.) $\star$ Repeat Question 4 replacing the operator $\div$ by the operator η of Question 5.