

Name: _____

- This exam is open book, open notes, but there can be no sharing of any material. You can use the Internet, but only as a library. If you are not sure if something is allowed, please ask.

- **COS 480:** Students should answer questions that are not marked with a \star for 40 points in 40 minutes. Questions marked with a \star may also be answered *within the 40 minutes*, for extra credit. (Some students may find a \star question to be easier than a non- \star one, so it is a good idea to quickly read all questions first.)


COS 580: Students should answer all questions, (\star and non- \star) for 60 points in 60 minutes.

- There are 10 questions (including two \star questions) on 7 pages.
- You have *40 minutes* (60 for COS 580) to earn *40 points* (60 for COS 580). You may wish to use this correspondence to plan your time.


1. (1 pt.) Write your name in the space provided above.
2. (4 pts.) Given relations $R_1(A, B, C)$ and $R_2(B, C, D, E)$, provide a relational algebra query *that does not use the natural join operator* and that is equivalent to the following query:

$$R_1 \bowtie \pi_{BDE} R_2$$

3. (5 pts.) What is the primary key of E in the following E-R diagram? Explain your answer briefly.



4. (5 pts.) What is the primary key of F in the following E-R diagram? Explain your answer briefly.



5. (5 pts.) Consider a relation $R(A, B, C, D, E, F)$ with functional dependencies

$$\begin{aligned}A &\rightarrow B \\BC &\rightarrow AE \\E &\rightarrow C \\AD &\rightarrow C\end{aligned}$$

Compute $\{B, E\}^+$, indicating the intermediate steps of the computation.

6. (5 pts.) In the schema of Question 5, do the given dependencies imply the dependency $AC \rightarrow E$? Justify your answer.

7. (10 pts.) In the schema of Question 5, determine all keys of R . Justify your answer, explaining both why the claimed keys are valid and why there are no other keys.

8. (5 pts.) In the schema of Question 5, list all BCNF violations (if any). Explain your answer.

9. (10 pts.) * Given a database with tables $S(A,B,C)$ and $T(C,D)$, with all attributes of integer type, provide an extended bag algebra query that is equivalent to the following SQL query:

```
select S.A, S.C, min(S.B) as X, avg(T.D) as Y
from S, T
where S.C = T.C and S.A > 20
group by S.A, S.C
having max(S.B) < 100
order by X, Y desc, S.A, S.C;
```

10. (10 pts.) ★ Consider an E-R diagram \mathcal{M} containing an entity set E . When this diagram is mapped to a logical schema using the method described in class (and the textbook), E is mapped to a relation R_E whose primary key consists of six attributes. What is the minimum possible number of ovals in \mathcal{M} ? Justify your answer. [Hint: The answer is not six.]