CS145 Midterm Examination

Spring 2000, Prof. Widom

- Please read all instructions (including these) carefully.
- There are 10 problems on the exam, with a varying number of points for each problem and subproblem for a total of 75 points. *You should look through the entire exam before getting started, in order to plan your strategy.*
- The exam is closed book and closed notes, but you may refer to your three pages of prepared notes.
- Please write your solutions in the spaces provided on the exam. Make sure your solutions are neat and clearly marked. You may use the blank areas and backs of the exam pages for scratch work. Please do not use any additional scratch paper.
- *Simplicity and clarity of solutions will count.* You may get as few as 0 points for a problem if your solution is far more complicated than necessary, or if we cannot understand your solution.

NAME: _____

In accordance with both the letter and spirit of the Honor Code, I have neither given nor received assistance on this examination.

SIGNATURE: _____

Problem	1	2	3	4	5	6	7	8	9	10	TOTAL
Max. points	16	4	5	4	6	8	10	6	4	12	75
Points											

1. **E/R Diagrams and Relations** (16 points)

Consider an entity-relationship diagram containing two entity sets: entity set E_1 with attributes K_1 (a key) and A_1, and entity set E_2 with attributes K_2 (a key) and A_2. The entity sets are connected by a binary relationship set R with attribute A. Perhaps you'd like to sketch the E/R diagram (not graded):

We are interested in translating this E/R diagram into relations. Specifically:

- We want as few relations as possible.
- We want all relations to be in BCNF.
- We do not want `NULL` values to appear in our relations.

For each of the following scenarios, specify an appropriate **set of relations** given the criteria above, and **underline a key** for each relation.

(a) R is a many-to-many relationship set, however each entity in E_1 can be related to a given entity in E_2 at most once. (That is, informally: for an entity $e_1 \in E_1$ and an entity $e_2 \in E_2$, (e_1, e_2) is in the relationship set at most once.)

(b) R is a many-to-many relationship. Each entity in E_1 can be related to a given entity in E_2 more than once, but there is a distinct value of attribute A for each relationship instance. (That is, informally: for an entity $e_1 \in E_1$, an entity $e_2 \in E_2$, and a value a, (e_1, e_2, a) is in the relationship set at most once.)

(c) R is a many-one relationship set from E_1 to E_2.

(d) R is a many-one relationship set from E_1 to E_2, and each entity in E_1 participates in the relationship set. (That is, informally: for each entity $e_1 \in E_1$, some (e_1, e_2) is in the relationship set.)

(e) R is a many-one relationship set from E_1 to E_2, and each entity in E_2 participates in the relationship set.

(f) R is a one-one relationship set.

(g) R is a one-one relationship set, and each entity in E_1 participates in the relationship set.

(h) R is a one-one relationship set, and each entity in E_1 and each entity in E_2 participates in the relationship set.

2. **Multiway Relationship Sets in E/R** (4 points)

Consider an entity-relationship diagram containing three entity sets: *Flights*, representing airline flights, *Customers*, representing people who take flights, and *Agencies*, representing agencies that can reserve flights on behalf of customers. There is a multiway relationship set *Reservations* connecting the three entity sets. Perhaps you'd like to sketch the E/R diagram (not graded):

(a) Suppose there is an arrow pointing to the *Customers* entity set, and no other arrows in the E/R diagram. *Briefly* state in English the real-world requirement that is encoded by the arrow.

(b) Suppose we wish to encode the requirement that a customer may use only one agency for all of his flight reservations. Can we encode this requirement by placing appropriate arrows in the E/R diagram?

Circle one: YES NO

If you answered yes, to which entity set(s) should the arrow(s) point?

3. **Functional Dependencies and Keys** (5 points)

Consider a relation $R(A, B, C, D, E, F, G, H)$ and the following functional dependencies over R:

$$\begin{aligned} A &\rightarrow BCD \\ AD &\rightarrow E \\ EFG &\rightarrow H \\ F &\rightarrow GH \end{aligned}$$

(a) Based on these functional dependencies, there is one key for R.

What is it? [ ]

(b) One of the four functional dependencies can be removed without altering the key.

Which one? [ ]

4. **Functional Dependencies and BCNF** (4 points)

Consider a relation $R(A, B, C, D)$ and the following functional dependencies over R:

$$\begin{array}{lcl} A & \rightarrow & B \\ D & \rightarrow & B \\ C & \rightarrow & A \end{array}$$

For each of the following decompositions of R, state whether all relations in the decomposition are in BCNF.

(a) $R_1(A, B)$, $R_2(D, B)$, $R_3(C, A)$
All in BCNF? Circle one: YES NO

(b) $R_1(A, B)$, $R_2(B, C)$, $R_3(C, D)$
All in BCNF? Circle one: YES NO

(c) $R_1(A, C, D)$, $R_2(A, B)$
All in BCNF? Circle one: YES NO

(d) $R_1(B, C, D)$, $R_2(A, C)$
All in BCNF? Circle one: YES NO

5. **Functional Dependencies and Multivalued Dependencies** (6 points)

Consider a relation $R(A, B)$ with two tuples: $R = \{(1, 2), (1, 3)\}$

(a) Does $A \rightarrow B$ hold for this instance of R?
Circle one: YES NO

(b) Does $A \twoheadrightarrow B$ hold for this instance of R?
Circle one: YES NO

Now consider a relation $R(A, B, C)$ with two tuples: $R = \{(1, 2, 3), (1, 3, 2)\}$

(c) Does $A \rightarrow B$ hold for this instance of R?
Circle one: YES NO

(d) Does $A \rightarrow C$ hold for this instance of R?
Circle one: YES NO

(e) Does $A \twoheadrightarrow B$ hold for this instance of R?
Circle one: YES NO

(f) Does $A \twoheadrightarrow C$ hold for this instance of R?
Circle one: YES NO

6. **Relational Algebra** (8 points)

Consider a relation $R(A, B)$ that contains r tuples, and a relation $S(B, C)$ that contains s tuples. You may assume $r > 0$ and $s > 0$. For each of the following relational algebra expressions, state in terms of r and s the minimum and maximum number of tuples that could be in the result of the expression.

Expression	minimum #tuples	maximum #tuples
$R \cup \rho_{S(A,B)} S$		
$R \bowtie S$		
$\Pi_B(R) - (\Pi_B(R) - \Pi_B(S))$		
$R \rtimes S$		

Note: If you have forgotten what the $\rtimes$ operator does, you may ask one of the staff members for a definition. However, that hint will cost you an automatic loss of 1 point on this problem.

7. **More Relational Algebra** (10 points)

Consider the following relational schema:

```
Name(ID, name)  // ID is a key
GPA(ID, gpa)    // ID is a key
```

Write a relational algebra expression to find the names of all students with the highest GPA in the database.

8. **SQL Queries** (6 points)

Consider relations $R(A)$ and $S(B)$ and the following four SQL queries:

(1) `SELECT * FROM R WHERE A >= ALL (SELECT B FROM S)`

(2) `SELECT * FROM R WHERE A >= ANY (SELECT B FROM S)`

(3) `SELECT * FROM R WHERE A >= (SELECT MAX(B) FROM S)`

(4) `SELECT * FROM R WHERE A >= (SELECT MIN(B) FROM S)`

For each pair of queries in the chart below, choose from among the following options. *You may assume that neither relation is empty.*

(a) The first query always returns a superset (not necessarily strict) of the tuples returned by the second.

(b) The first query always returns a subset (not necessarily strict) of the tuples returned by the second.

(c) Options (a) and (b) are both true, so the two queries always return exactly the same set of tuples.

(d) None of the above.

Fill in the following chart:

First query	Second query	Your selection: (a), (b), (c), or (d)
(1)	(2)	
(1)	(3)	
(1)	(4)	
(2)	(3)	
(2)	(4)	
(3)	(4)	

9. **Programming with SQL** (4 points)

Consider a relation $R(A, B)$ that may contain duplicate tuples, and suppose we are interested in writing an operation that deletes exactly one copy of each tuple appearing in R. For example, if initially $R = \{(1,2),(1,2),(3,4),(5,6),(5,6),(5,6)\}$, then after our deletion operation we want $R = \{(1,2),(5,6),(5,6)\}$.

(a) Is it possible to perform this deletion operation using a single SQL statement?

Circle one: YES NO

Justify your answer *in one sentence or less*:

(b) Is it possible to perform this deletion operation using SQL in conjunction with a programming language (i.e., using embedded SQL, dynamic SQL, or a DBMS programming language such as PL/SQL)?

Circle one: YES NO

Justify your answer *in one sentence or less*:

10. **SQL and Outerjoin** (12 points)

You are to implement the *natural full outerjoin* operation in SQL, *without using SQL's* `outerjoin` *operator*. Specifically, write a SQL query that computes the natural full outerjoin of relations $R(A, B)$ and $S(B, C)$. You may `select` the constant value `NULL` if you find it useful; for example, the following query returns all tuples in $Q(A, B, C)$ with all B values replaced by `NULL`:

```
select A,NULL,C from Q
```

Note: The outerjoin operator was covered in your textbook readings. If you do not recall what it does, you may ask one of the staff members for a definition. However, that hint will cost you an automatic loss of 3 points on this problem.

Write the SQL query here: