CS411 Database Systems
Fall 2009

University of Illinois at Urbana-Champaign

Midterm Examination
October 13, 2009
Time Limit: 75 minutes

- Print your name and NetID below. In addition, print your NetID in the upper left corner of every page.

Name: _____ **NetID:** _____

- Closed notes; closed book; no sheet of formulas permitted.
- Please write your answers directly on the exam sheet. The space we left for your answers is often more than what you actually need. Please use the back side of the exam as scratch paper.
- In case you find a question ambiguous, please write down your assumption and answer the question accordingly.
- You may use temporary relations, if you like, for any of the queries below. If you use the same temporary relation for a second exam question, you must redefine the relation in your answer to the second question. In other words, your answer to each question should be self-contained.

Problem	1	2	3	4	5	6	Total
Points	10	16	15	15	20	14	90
Score							
Grader							

Turn over the page when instructed to do so. Good Luck!

Problem 1 Basic Concepts (10 points)

For each of the following statements, indicate whether it is $TRUE$ or $FALSE$ by circling your choice. You will get 1 *point* for each correct answer, -1 *point* for each incorrect answer, and 0 *point* for each answer left blank.

1) TRUE FALSE
A weak entity set does not have any attribute that belongs to its key.

2) TRUE FALSE
Every isa relation in the ER model is one-one.

3) TRUE FALSE
If a functional dependency $A \rightarrow B$ holds in relation R(A, B, C), then $AC \rightarrow B$ also holds.

4) TRUE FALSE
If a relation R is in 3NF, R is also in BCNF.

5) TRUE FALSE
The BCNF decomposition algorithm does not preserve functional dependencies of the initial relation R.

6) TRUE FALSE
If a multi-valued dependency $A_1, \ldots, A_n \twoheadrightarrow B1, \ldots, B_m$ holds, then a functional dependency $A_1, \ldots, A_n \rightarrow B_1, \ldots, B_m$ also holds.

7) TRUE FALSE
Natural join is one the five basic operations in relational algebra.

8) TRUE FALSE
In the WHERE clause of a SQL query, the condition 'Smith' = NULL is evaluated to be false.

9) TRUE FALSE
It is possible to declare multiple different sets of attributes as UNIQUE in the same relation R.

10) TRUE FALSE
If we define a foreign key in relation R, the DBMS checks the foreign-key constraint whenever a tuple in R is deleted.

Problem 2 ER Diagram (16 points)

a) Draw an ER diagram based on the following description: Suppose we have two entity sets, `People` and `Email`. Suppose we also use a relationship `Owns`, which connects these two entities. A person may own multiple email accounts, but an email account can only be owned by a single person. You do not have to draw the attributes for the entities. (2 Points)

b) Draw an ER diagram based on the following description: Suppose we have three entity sets, `Customers`, `Accounts`, and `Branch`. Suppose we also use a relationship `Has` that connects these three entities. For every combination of a customer and a branch, there is a single account. For every combination of a customer and an account, there is a single branch. You do not have to draw the attributes for the entities. (2 Points)

c) Drawn an ER diagram based on the following description: Suppose we have two entity sets, `Account` and `CheckingAccount`. `Account` has two attributes `AccountID` and `Balance`. `AccountID` uniquely identifies Account. `CheckingAccount` has one attribute `Overdraft`. `CheckingAccount` is a subclass of `Account`. (2 Points)

d) Based on the diagram you drew in part c), convert this ER diagram into a set of relations using the following approaches:

1. ER-style approach (2 Points):

2. Object-Oriented approach (2 Points):

e) Which approach, ER-style or Object-Oriented, requires the least amount of storage for the converted relations? Justify your answer. (2 Points)

f) Figure 1 shows an ER diagram.

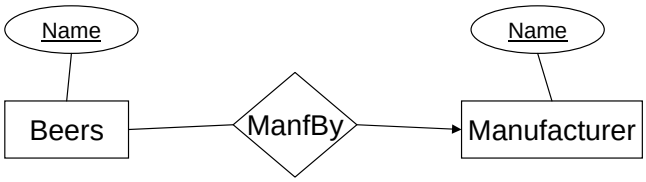

Figure 1: Problem 2f

1) Explain why this is not a good design. (2 Points)

2) Convert this ER diagram into a good design. (2 Points)

Problem 3 Functional Dependencies and Keys (15 points)

a) Consider the relation Treatment and FDs below. Describe, with examples, two different types of anomaly that can arise. (3 points)
Treatment(doctorID, doctorName, patientID, diagnosis)

$doctorID \rightarrow doctorName$
$doctorID, patientID \rightarrow diagnosis$

b) Prove that every two-attribute relation is in BCNF. (4 points)

c) Prove that if relation R is in 3NF and every key is simple (i.e., a single attribute), then R is in BCNF. (4 points)

d) Given a relation R(A,B,C,D,E) and FDs $D \rightarrow B$, $DE \rightarrow A$, $C \rightarrow AD$, decompose R into BCNF. (4 points)

Problem 4 Relational Algebra (15 points)

a) Given two relations R and S with the same set of attributes, select all the equality equations that hold for every instances of R and S. We here assume the set semantics. (3 points)

1) $R \cap S = R - (R - S)$
2) $R \bowtie S = S \bowtie R$
3) $R \cap S = R \bowtie S$
4) $R \cap S = R \times S$

b) Write a relational algebra expression for each of the following questions, using the set of relations provided below. (12 points, 4 points each)

MovieFan(movieFanName, age, isStudent)
Frequents(movieFanName, cinema)
Likes(movieFanName, movieTitle)
Shows(cinema, movieTitle, ticketPrice)

1) Find all `cinemas` that shows at least one movie satisfying the following two conditions:
- The price of the movie is less than $8.00.
- James likes that movie.

2) Find all `cinemas` that are frequented by only students (Assume that isStudent is a Boolean attribute containing either true or false).

3) Find the `movieFanName` of all students who frequent cinemas showing at least one movie they like.

Problem 5 SQL (20 points)

We have an employee database with three tables. The `Employee` table stores information about employees. Every employee is identified by an employeeID. The `Department` table stores information about each department. The `Vacations` table stores information about the total number of days each employee takes for vacation.

Employee(EmployeeID, FirstName, LastName, Office, Email, DepartmentID)
Department(DepartmentID, DepartmentName)
Vacations(EmployeeID, Days)

Answer the following queries using SQL.

a) List the `Office` and `Email` of employee "John Smith." (3 Points)

b) List the number of vacation `Days` taken by each employee with the name "John Smith". (3 Points)

c) List the `FirstName` and `LastName` of all employees who never took a vacation day. (6 Points)

d) List the `DepartmentName` of every department whose total number of vacation days taken by its employees is the largest among those of all the departments. (8 Points)

Problem 6 Data Modification, constraints and Trigger (14 points)

a) Describe a situation where an attribute-based constraint with the keyword CHECK is violated. (4 Points)

b) Explain how the cascade policy maintains foreign-key constraints. (5 Points)

c) Explain the main reason that we sometimes prefer to use triggers rather than assertions. (5 Points)