Database Systems
Comprehensive Exam
Fall 2015

Name:

1) (30 points) Consider a medical clinic for treating patients. In this clinic all patients are required to make appointments with doctors in advance (no walk-ins allowed). Clinic requirements dictate that the following information be recorded for each visit: the date and time the patient arrives at the clinic, the time they leave, and the total amount billed for the visit (cost). Note that there are no overnight stays, but there can be multiple visits in the same day by the same patient. Also of great concern to the clinic is the policy number for each patient's medical insurance. It is assumed that patients will have only one medical insurance policy, members of the same family may be listed on the same policy, and patients without insurance are not allowed to visit the clinic. Each patient is assigned to exactly one doctor the first time they visit the clinic, and this doctor sees the patient on all future visits. Finally, basic patient and information must also be recorded including name and social security number.

Give an entity-relationship diagram for representing the above information.

2) (40 points) Consider the following relational schemes. Note that attributes forming the primary key for each relation have been underlined.

student(<u>student-name</u>, street, city)	-- Basic student information.
offering(<u>department</u>, <u>number</u>, population)	-- Courses currently offered; for CSE5260 department is "CSE" and number is 5260. Population is the number of students.
titles(<u>department</u>, <u>number</u>, title)	-- Course titles; "CSE5260" is "Database Systems"
enrollment(<u>student-name</u>, <u>department</u>, <u>number</u>)	-- Indicates which students are enrolled in which courses.

(a) Give DDL for each of the above tables.

/* note that the tables from the previous page are repeated below for your convenience */

student(student-name, street, city)
offering(department, number, population)
titles(department, number, title)
enrollment(student-name, department, number)

Give an SQL query for part (b).

(b) For each student, list the course they are enrolled in that has the largest population. Include the students name, department, course number, title, and population in the result. Sort the result by population.

student-name	department	number	title	population
Bob Jones	MTH	1001	Calculus	15
Mary Carson	CSE	4020	Database Systems	25
Joe Smith	PSY	4260	Child Development	62
Carl Brown	CHM	5264	Introduction to Chemistry	83

Give a tuple calculus expression for part (c).

(c) A list of the names of students who are currently enrolled in every class that "Bob Jones" is currently enrolled in, and in every class that "Will Smith" is currently enrolled in. In other words, if student "Mary Brown" is listed in the result of the query, then she is enrolled in all the classes that Bob Jones is enrolled in, and all the classes that Will Smith is enrolled in.

Give a relational algebra expression for part (d).

(d) A list of the names of all students who are enrolled in CSE 5260 or MTH 5100, but not both.

3) Consider the following set F of functional dependencies for the relational scheme R=(A,B,C,D,E,F,G).

A=>BD
C=>ABE
F=>G

(a) (5 points) Circle each of the following that is a candidate-key for the above relational scheme.

A C CF AF ABCDEFG

(b) (10 points) Assume that the above set of functional dependencies is a canonical cover. Give a decomposition of the above relational scheme into a collection of relational schemes that are in 3NF, have a lossless join and preserves dependencies.

4) Consider the following set of functional dependencies for the relational scheme R=(A,B,C,D,E,F,G).

A=>B
A=>D
C=>EA
F=>G

(a) (5 points) Give one non-trivial functional dependency that is in the closure of the above set of functional dependencies. Use Armstrong's axioms to prove that the functional dependency is in the closure.

(b) (5 points) Consider the decomposition of R into R1 = (A,B,F,G) and R2 = (A,C,D,E). Does this decomposition have a lossless join (yes or no)? Be sure to explain your answer.

(c) (5 points) Now consider the following set of functional dependencies for the same relational scheme:

A=>B
A=>D
C=>EA
F=>G
C=>B

Does the decomposition of R into R1 = (A,B,C,F,G) and R2 = (A,C,D,E) have a lossless join (yes or no)? Be sure to explain your answer.