WOMEN'S UNIVERSITY IN AFRICA

Addressing gender disparity and fostering equity in University Education

FACULTY OF MANAGEMENT AND ENTREPRENEURIAL SCIENCES

BSc HONOURS DEGREE IN INFORMATION SYSTEMS

BSc HONOURS DEGREE IN COMPUTER SCIENCE

MAIN PAPER

IS214: ADVANCED DATABASE SYSTEMS

INTAKE 22: THIRD YEAR SECOND SEMESTER

INTAKE 1: THIRD YEAR SECOND SEMESTER

TIME: 2 HOURS MORNING

INSTRUCTIONS TO CANDIDATES

Answer any **four** questions.

Question 1

A university registrar's office maintains data about the following entities: (a)courses, including number, title, credits, syllabus, and prerequisites; (b) course offerings, including course number, year, semester, section number, instructor(s), timings, and classroom; (c) students, including student-id, name, and program; and (d) instructors, including identification number, name, department, and title. Further, the enrolment of students in courses and grades awarded to students in each course they are enrolled for must be appropriately modelled.

a) Construct an E-R diagram for the registrar's office. Document all assumptions that you make about the mapping constraints. [12]

b) Convert E-R diagram created in (a) into relational schema. [13]

Question 2

Suppose relation R(A,B,C) has the following tuples:

A	B	C
1	2	3
4	2	3
4	5	6
2	5	3
1	2	6

and relation S(A,B,C) has the following tuples:

A	B	C
2	5	3
2	5	4
4	5	6
1	2	3

Compute

a) The natural join of R and S. [4]

b) The theta-join of R and S with the condition R.B = S.B AND R.A <S.C. [4]

c) The projection $\pi_{A,B}(R)$ [2]

d) The Cartesian product $R \times S$ [3]

e) The selection $\sigma_{B=5}(R)$ [2]

f) The union of R and S. [3]

g) The intersection of the relations R and S. [3]

h) (R - S) union (S - R), often called the "symmetric difference" of R and S. [4]

Question 3

Show how you can specify the following relational algebra operations in both tuple and domain relational calculus.

a) σ A=C(R(A, B, C)) [4]

b) π <A, B>(R(A, B, C)) [4]

c) R(A, B, C) × S(C, D, E) [5]

d) R(A, B, C) ∪ S(A, B, C) [6]

e) R(A, B, C) ∩ S(A, B, C) [6]

Question 4

You've started a new movie-rating website, and you've been collecting data on reviewers' ratings of various movies. There's not much data yet, but you can still try out some interesting queries. The schema is as follows:

Movie (mID, title, year, director)

English: There is a movie with ID number *mID*, a *title*, a release *year*, and a *director*.

Reviewer (rID, name)

English: The reviewer with ID number *rID* has a certain *name*.

Rating (rID, mID, stars, ratingDate)

English: The reviewer *rID* gave the movie *mID* a number of *stars* rating (1-5) on a certain *ratingDate*

Write the following quires in SQL:

a) Find the titles of all movies directed by Steven Spielberg. [3]

b) Find all years that have a movie that received a rating of 4 or 5, and sort them in increasing order. [3]

c) Find the titles of all movies that have no ratings. [3]

d) Some reviewers didn't provide a date with their rating. Find the names of all reviewers who have ratings with a NULL value for the date. [3]

e) Write a query to return the ratings data in a more readable format: reviewer name, movie title, stars, and ratingDate. Also, sort the data, first by reviewer name, then by movie title, and lastly by number of stars. [4]

f) For all cases where the same reviewer rated the same movie twice and gave it a higher rating the second time, return the reviewer's name and the title of the movie. [3]

g) For each movie that has at least one rating, find the highest number of stars that movie received. Return the movie title and number of stars. Sort by movie title. [3]

h) List movie titles and average ratings, from highest-rated to lowest-rated. If two or more movies have the same average rating, list them in alphabetical order. [3]

Question 5

Consider schedules H1 and H2 given below:

H1 = r1(x); r2(z); r1(z); r3(x); r3(y);w1(x);w3(y); r2(y);w2(z);w2(y)

H2 = r1(x); r2(z); r3(x); r1(z); r2(y); r3(y);w1(x);w2(z);w3(y);w2(y)

These schedules are generated by the following transactions:

T1 = r1(x); r1(z);w1(x)

T2 = r2(z); r2(y);w2(z);w2(y)

T3 = r3(x); r3(y);w3(y)

a) Draw the serialization graph for H1, H2 and state whether or not it is serializable. If it is serializable, give the equivalent serial schedule. [12]

b) Discuss the 'ACIDity' property of transactions. [13]

Question 6

a) Suppose you are given a relation R=(A,B,C,D,E) with the following functional dependencies:

{BC $\rightarrow$ ADE, D $\rightarrow$ B}.

i. Find all candidate keys. [5]

ii. Identify the best normal form that R satisfies (1NF, 2NF, 3NF, or BCNF). [5]

iii. Decompose R until it becomes BCNF. At each step, identify a new relation, decompose and re-compute the keys and the normal forms they satisfy. For each decomposition show if your decomposition is good or bad. [10]

b) Suppose you are given a relation R=(A,B,C,D,E) with the following functional dependencies:

BD $\rightarrow$ E, A $\rightarrow$C.

Show that the decomposition into R1=(A,B,C) and R2=(D,E) is lossy. You can show using any method. [5]

END