UNIVERSITY OF RUHUNA

Faculty of Engineering

End-Semester 5 Examination in Engineering: August 2018

Module Number: EE5203 **Module Name: Data Management Systems**

[Three Hours]

[Answer **all questions**, each question carries 10 marks]

Q1 a) Describe what is a weak entity.

[1 Mark]

b) What is the difference between the primary key and the foreign key?

[1 Mark]

c) Draw the ER diagram and the table structure with key constraints based on the following case study.

Assume Sri Lanka has 10 wind power plants which are managed by a private company. Each plant contains more than 5 wind turbines and the generated power is connected to the national grid. Usually a power station has one technical officer and 5 support staff to manage the operation at the ground level. Company has 3 electrical engineers to manage the entire operation and accountants to handle office work. Assume you are hired to develop a database management system to track the operation and optimize the usage.

[4 Marks for the ER]
[4 Marks for the Tables]

Q2 a) Describe the following relational operations.

i) Join
ii) Outer join
iii) Intersection
iv) Left outer join
v) Union

[5 Marks]

b) Answer the following questions based on the case study and the table structure given.

Telecom service provider owns 125 base stations in Sri Lanka. All the base stations are distributed so that they provide the full coverage to the country. In this operation, they have employed technical officer and support staff member for each base station. There is one engineer for 25 base stations.

Assume you are the chief engineer who is heading the operation for all 125 base stations and you have designed the following tables of the database management system to manage the staff details, basic inventory and the power utilization.

In the table structure the primary keys are underlined and the foreign keys are bold.

Employee(emp_id, first_name, last_name, salary, employee_type)

Station(station_id, name, location, power_utilization, **technical_officer_id**, **support_staff_id**, **engineer_id**)

Inventory(inv_id, item_name, bought_date, price, **assigned_ engineer_id**)

i) List first name of all support staff who works in station 5.
ii) List the salary of the technical officers of all the power plants.
iii) List the names of the items in the inventory which are assigned to the engineer in station 3.
iv) List the station names with its power utilization and its technical officers first names.
v) Find the total power utilization of all the stations.

[5 Marks]

Q3 a) What are the advantages of having users and roles in a database management system?

[2 Marks]

b) Nuwan, Shehan and Eranga are working on a database management system. Nuwan is the database administrator and has full privileges on the database. All three of them simultaneously log on to the database and issue following commands.

#	User	Command
1	Nuwan	create table marks (sid number(2), marks number(4));
2	Nuwan	insert into marks values (1,87);
3	Nuwan	savepoint a;
4	Nuwan	select * from marks;
5	Nuwan	grant select on marks to Shehan with grant option;
6	Shehan	grant select on Nuwan.marks to Eranga;
7	Eranga	select * from Nuwan.marks;
8	Nuwan	grant insert, delete on marks to Eranga;
9	Eranga	insert into Nuwan.marks values(3,66);
10	Eranga	savepoint a;
11	Eranga	select * from Nuwan.marks;
12	Eranga	rollback to a;
13	Nuwan	insert into marks values(2,74);
14	Eranga	commit;
15	Shehan	select * from Nuwan.marks;
16	Nuwan	rollback to a;
17	Nuwan	commit;
18	Nuwan	select * from marks;
19	Eranga	insert into Nuwan.marks values(5,57);
20	Shehan	select * from Nuwan.marks;
21	Eranga	commit;
22	Nuwan	Revoke select on marks from Shehan;
23	Eranga	select * from Nuwan.marks;
24	Nuwan	revoke select on marks from Shehan;
25	Eranga	select * from Nuwan.marks;

26	Eranga	insert into Nuwan.marks values(10,84);
27	Nuwan	select * from marks;
28	Eranga	commit;
29	Nuwan	commit;
30	Nuwan	select * from marks;
31	Shehan	select * from Nuwan.marks;

Write down the output of all the "select" statements according to the command number.

[1 Mark for each correct output]
[Maximum 8 Marks]

Q4 a) What are the advantages and disadvantages of using indexing in databases?

[2 Marks]

b) A movie store has 8000 movies which are categorized into 20 genres. When a customer comes to the store they can search for movies by genre or by main actor by using the provided database management system. Given below is the structure of the table movie. Hint: Average number of movies per actor is 40.

Movie(id number(5), movie_name varchar(100), main_actor varchar(100), genre varchar(50, rating number(2))

The storage system has a block size of 256KB where as a single record to the table movie requires 64KB. If an index is created on the table **Movie**, it will occupy 4 blocks. (4 blocks per each index created)

On a typical day Query 1, Query 2 and Query3 are executed on the system at an average frequency of 30, 50 and 20 times per hour, respectively

Query 1: SELECT movie_name from Movie
WHERE genre = 'Action';

Query 2: SELECT movie_name from Movie
WHERE genre = 'Action' AND main_actor = 'Tom Cruise';

Query 3: UPDATE Movie SET rating =8 WHERE id=47;

What indexes will you create to improve the performance of the system. Clearly state your assumptions. (Answers without supporting calculations will carry no marks).

[8 Marks]

Q5 a) Explain following terms used in database architecture.

i) Block
ii) Tuple
iii) Dirty block
iv) Indexing

[2 Marks]

b) An airline maintains a database to capture information on their operations. This database includes two tables named Delay and Pilot. The structure of the tables are given below.

Delay (delay_id, pilot_id, flight, date, origin, destination, delay_category, delay_mins, aircraft_type)

Pilot (pilot_id, last_name, first_name, birth_date, gender, hired_date, salary, flying_hours, skill_level)

The airline serves 87 cities. Number of flights to each city can be assumed to be uniform. They have a total of 800 pilots out of which 37 are females. Salary of a pilot ranges from 100,000 to 450,000. 20% of pilots earn less than 150,000. 60% of all pilots are older than 30 years. 10% of all pilots are older than 50 years. The administration observes that the majority of delays occurring falls under 'weather' category. Airline owns a fleet of 127 aircrafts out of which 30 are Boeings and others are Airbuses. Skill level of a pilot is ranged from 1-10.

Developers are running following queries on the database. All of them are functionally correct and yield the correct results. Rewrite these queries to improve the performance. If you have any assumptions clearly state them with reasons.

Query 1:
```
SELECT last_name, first_name FROM Pilot p, Delay d
WHERE p.pilot_id = d.pilot_id
AND birth_date - 233 < '16-Aug-1968'
AND gender <> 'Female'
AND aircraft_type = 'Boeing'
AND delay_mins> 45
AND destination = 'Mumbai'
AND flight_date + 150 > '01-Jan-2017'
AND skill_level -2> 6;
```

Query 2:
```
SELECT origin, destination, delay_mins FROM Delay
WHERE date - 150 < '01-Jan-2016'
AND date + 100 > '01-Jan-2016'
AND origin = 'Mumbai
AND delay_mins> (SELECT AVG(delay_mins) FROM Delay WHERE
aircraft_type<> 'Airbus')
AND delay_catgory = 'weather'
AND origin <> 'Colombo';
```

Query 3:
```
SELECT last_name, first_name FROM Pilot p
WHERE hired_date +50 < '16-Aug-2006'
OR gender = 'Male'
OR salary > 175000
OR flying_hours> 450
OR last_name LIKE 'HE%'
OR skill_level + 2> 9;
```

Query 4: SELECT last_name, first_name FROM Pilot p, Delay d
WHERE p.pilot_id = d.pilot_id
AND (birth_date - 233 < '16-Aug-1968'
OR gender <> 'Male'
OR aircraft_type = 'Boeing'
OR delay_mins> 30
OR flight_date + 150 > '01-Jan-2017'
OR delay_mins +5 < 60);

[2 x 4 = 8 Marks]