



UNIVERSITY OF COLOMBO, SRI LANKA

UNIVERSITY OF COLOMBO SCHOOL OF COMPUTING

DEGREE OF BACHELOR OF INFORMATION TECHNOLOGY (EXTERNAL)

Academic Year 2012/2013 – 1st Year Examination – Semester 2

IT2304 – Database Systems I
Multiple Choice Question Paper

28th July, 2013
(TWO HOUR)

Important Instructions :

- The duration of the paper is **2 (two) hour**.
- The medium of instruction and questions is English.
- The paper has **40 questions** and **17 pages**.
- All questions are of the MCQ (Multiple Choice Questions) type.
- All questions should be answered.
- Each question will have 5 (five) choices with **one or more** correct answers.
- All questions will carry equal marks.
- There will be a penalty for incorrect responses to discourage guessing.
- The mark given for a question will vary from 0 (*All the incorrect choices are marked & no correct choices are marked*) to +1 (*All the correct choices are marked & no incorrect choices are marked*).
- Answers should be marked on the special answer sheet provided.
- Note that questions appear on both sides of the paper.
If a page is not printed, please inform the supervisor immediately.
- Mark the correct choices on the question paper first and then transfer them to the given answer sheet which will be machine marked. **Please completely read and follow the instructions given on the other side of the answer sheet before you shade your correct choices.**

1) With respect to the ANSI/SPARC architecture, which of the following is/are true?

- (a) Constraints applicable to data are identified in the physical schema.
- (b) File organization is defined in the conceptual schema.
- (c) Indexing mechanism is defined in the conceptual schema.
- (d) File organization is defined in the physical schema.
- (e) Constraints applicable to data are identified in the conceptual schema.

2) Physical data independence is

- (a) making changes at the conceptual level without affecting the physical level.
- (b) making changes at the physical level without affecting the conceptual level.
- (c) making changes at the external level without affecting the conceptual level.
- (d) making changes at the conceptual level without affecting the external level.
- (e) making changes at the physical level without affecting the external level.

3) In a banking system the following scenarios may take place.

- (I) Account balance could go below the minimum allowed.
- (II) Customers are able to find out the bank balance of other customers.
- (III) Incorrect bank balance might be reflected due to simultaneous withdrawals and deposits,
- (IV) If system fails in the middle of a money transfer from one account to another account, it is not possible to establish the consistent state that existed prior to a failure.

Select the correct statement(s) which may provide reasons for the above scenarios.

- (a) (I) occurs due to concurrency control failure.
- (b) (II) occurs due to security enforcement failure.
- (c) (III) occurs due to integrity enforcement failure.
- (d) (IV) occurs due to concurrency control failure.
- (e) (I) occurs due to integrity enforcement failure.

4) Given below are five major steps that one would perform in setting up a database for a particular enterprise.

- (I) Develop a user interface to carry out the task
- (II) Define integrity constraints on the data.
- (III) Define DDL statements for all appropriate types of data and data relationships.
- (IV) Define the conceptual model
- (V) Define the physical level.

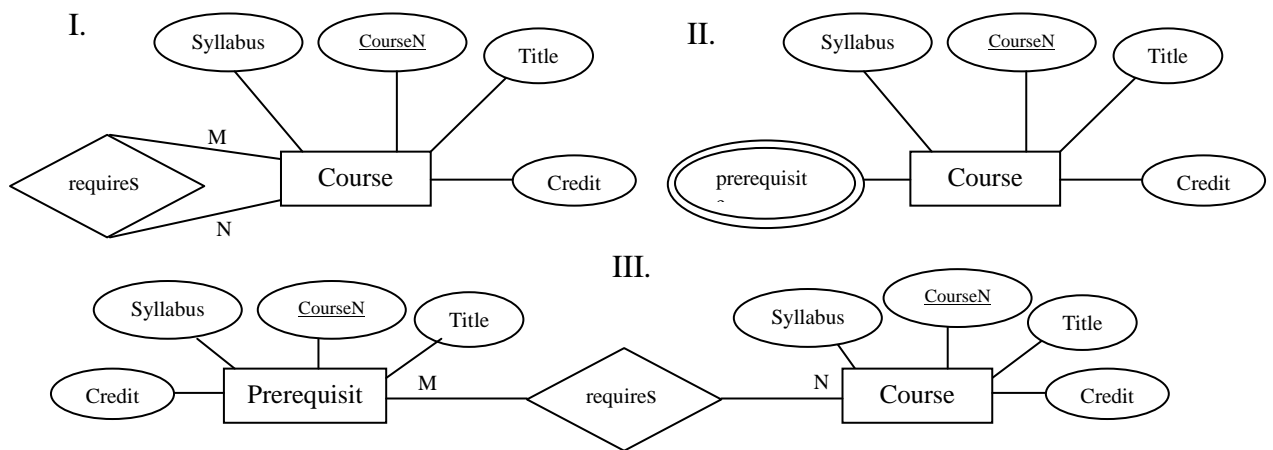
What is the correct order of carrying out the steps given above?

- | | | |
|-----------------------------|-----------------------------|-----------------------------|
| (a) (I) (V) (IV) (III) (II) | (b) (IV) (III) (II) (V) (I) | (c) (III) (II) (IV) (V) (I) |
| (d) (III) (II) (V) (IV) (I) | (e) (I) (IV) (II) (III) (V) | |

5) Which of the following statement(s) is/are true about candidate key?

- (a) There can be only one Candidate Key in a relation.
- (b) There can be many Candidate Keys in a relation.
- (c) The primary *key* is one of the candidate keys that is chosen by the database designer.
- (d) Candidate keys are considered as foreign keys in a relation.
- (e) Candidate Key is unique and that alone is sufficient to qualify as a Primary Key.


6) Consider a scenario where a course in turn requires some other courses to be taught as prerequisites. One course may be taught as a prerequisite in many other courses. This scenario is to be depicted in the following diagrams given from (I) to (III).



Which diagram(s) is/are able to produce the correct scenario

- (a) (I) Only
- (b) (II) Only
- (c) (III) Only
- (d) (I) and (II) Only
- (e) (I) and (III) Only


7) Consider the following ER diagram.



What would one get when **Section** in the above ER diagram maps to the corresponding relation(s)?

- (a) Section(SecNo, Semester, Cno).
- (b) Section(SecNo, Semester, Cno).
- (c) Section(SecNo, Semester, Cno, Empid).
- (d) Section(SecNo, Semester, Cno, Empid).
- (e) The given ER diagram is incorrect and it cannot be mapped since **Section** as a weak entity cannot participate in a relationship with **Teacher** entity.

- 8) Consider the following ER diagram.



What would one get when **CD** in the above ER diagram maps to the corresponding relation(s)?

- (a) CD(Number, CDTitle, FirstName, LastName, Author) only.
- (b) CD(Number, CDTitle, FirstName, LastName) and CD_Songs(Number, Author).
- (c) CD(Number, CDTitle) and CD_Performers(Number, FirstName, LastName Author).
- (d) CD(Number, FirstName, LastName, CDTitle) and CD_Songs(Number, Author).
- (e) CD(Number, CDTitle), CD_Performers(Number, FirstName, LastName) and CD_Songs(Number, Author).

Consider the following scenario to answer the question from (9) – (11).

A database is required by the university to keep track of the industrial placement process of the 3rd year students. There is a coordinator who is recruited as a permanent staff member with an *Empid* at the university to manage this process. To this end the coordinator collects student details and he forwards these details to companies. Each student has a unique registration no and a company is identified through its name. One student's details would be sent to many companies and a company may receive many students' details depending on the number of vacancies and the specialty. A student may have several specialties. The coordinator blacklists the companies when they do not provide proper training to students. The companies inform the selected list of students to the coordinator. The coordinator always produces status reports (which are identified through the dates) containing information such as who has been placed to what companies and the number of students placed at each company.

- 9) In the corresponding conceptual database model of the above scenario, what entities would be represented?

- (a) Coordinator, Company, Status_Report, Student
- (b) Company, Student, Coordinator
- (c) Student, Company
- (d) Company, Student, Status_Report
- (e) Coordinator, Company, Student, Specialty

10) Which of the following relationship(s) could exist in the corresponding ER/EER diagram?

- (a) 1:N – Collects between Coordinator and Student
- (b) M:N – Sends between Coordinator and Company
- (c) 1:N – Blacklists between Coordinator and Company
- (d) M:N – is_Placed between Student and Company
- (e) 1:N – Produces between Coordinator and Status_Report

11) The coordinator requires to keep track of the date that each student detail was sent to each company.
How would this *sent_date* be represented?

- (a) As an attribute in Student
- (b) As an attribute in Status Report
- (c) As an attribute in Company
- (d) As an attribute in the Sends relationship between Coordinator and Company
- (e) As an attribute in the is_Placed relationship between Student and Company

12) Suppose that two relations $R(A,B)$ and $S(A,B)$ have exactly the same schema and consider the following equalities.

- (I) $R \cap S = R - (R - S)$
- (II) $R \cap S = S - (S - R)$
- (III) $R \cap S = (R \cup S) - ((R - S) \cup (S - R))$
- (IV) $R \cap S = R * S$ // * is natural join

Which of the above equalities hold in relational algebra?

- (a) (I) and (III) only.
- (b) (I) and (II) only.
- (c) (I),(II) and (III) only.
- (d) (II), (III) and (IV) only.
- (e) All.

- 13) Let $R = (A, B, C)$ be a schema with attributes A,B,C and let $r1$ and $r2$ both be relations on schema R . The following relational algebra expressions are given along with SQL statements and each SQL statement is considered as producing the equivalent output as its corresponding relational algebra expression.

| | |
|---|---|
| (I) $\pi_{BC}(r1 \cap r2)$ | SELECT * FROM r1 WHERE(A, B, C) IN (SELECT * FROM r2); |
| (II) $\pi_{BC}(r1 - r2)$ | SELECT DISTINCT B, C FROM r1 WHERE (A, B, C) NOT IN (SELECT * FROM r2); |
| (III) $\pi_{AC}(r1) \bowtie \pi_{BC}(r2)$ | SELECT r1.A, r2.B, r2.C FROM r1, r2 WHERE r1.C = r2.C; |

Which of the following statement(s) is/are true with respect to the output equivalence of SQL statements to its relational algebra expression?

- | | |
|---|---------------------------------------|
| (a) Only (III) is equivalent. | (b) Only (I) and (II) are equivalent. |
| (c) Only (II) and (III) are equivalent. | (d) All are equivalent. |
| (e) Only (II) is equivalent. | |

Answer the question from (14) – (17) considering the following part of the database to keep track of students and their marks for the courses that they follow.

Student(sid, sname, year)
Registered(sid, cid, mark, grade)
Course(cid, cname, no_credit, deptid)

- 14) Which of the following relational algebra expression(s) would list the student id (sid) and name (sname) of each student along with the course id (cid) and name of the courses (cname) for which the student has registered? Students who are not registered for any course must also be listed.

- | |
|---|
| (a) $\pi_{sid, sname, cid, cname}((Student \bowtie_{sid=sid} Registered) * Course)$ |
| (b) $\pi_{sid, sname, cid, cname}((Student \bowtie_{sid=sid} Registered) * Course)$ |
| (c) $\pi_{sid, sname, cid, cname}((Student \bowtie_{sid=sid} Registered) * Course)$ |
| (d) $\pi_{sid, sname, cid, cname}((Student * Registered) * Course)$ |
| (e) $\pi_{sid, sname, cid, cname}((Registered \bowtie_{sid=sid} Student) * Course)$ |

- 15) Select the relational algebra expression(s) which would find the course id and department id of all courses in which no student has ever got "F".

(a) $\text{Res1} \leftarrow \pi_{\text{deptid}, \text{cid}} (\sigma_{\text{grade} = 'F'}(\text{Registered}) * \text{Course})$
 $\text{Result} \leftarrow \text{Course} - \text{Res1}$

(b) $\text{Result} \leftarrow \pi_{\text{deptid}, \text{cid}} (\sigma_{\text{grade} <> 'F'}(\text{Registered}) * \text{Course})$

(c) $\text{Res1} \leftarrow \pi_{\text{deptid}, \text{cid}} (\sigma_{\text{grade} = 'F'}(\text{Registered}) * \text{Course})$
 $\text{Result} \leftarrow \pi_{\text{deptid}, \text{cid}} (\text{Course}) - \text{Res1}$

(d) $\text{Res1} \leftarrow \sigma_{\text{grade} = 'F'} (\pi_{\text{cid}} (\text{Registered}) * \text{Course})$
 $\text{Result} \leftarrow \pi_{\text{deptid}, \text{cid}} (\text{Course}) - \text{Res1}$

(e) $\text{Res1} \leftarrow \pi_{\text{deptid}, \text{cid}} (\sigma_{\text{grade} = 'F'}(\text{Registered}) * \text{Course})$
 $\text{Result} \leftarrow \text{Res1} - \pi_{\text{deptid}, \text{cid}} (\text{Course})$

- 16) Which of the following sequence of operations would list the names of students who have **not** registered for any course?

(a) $\text{Reg_S} \leftarrow \pi_{\text{Sid}}(\text{Registered})$
 $\text{RESULT1} \leftarrow \pi_{\text{Sid}}(\text{Student})$
 $\text{RESULT2} \leftarrow \text{RESULT1} \cap \text{Reg_S}$
 $\text{RESULT} \leftarrow \pi_{\text{sname}}(\text{RESULT2} \bowtie_{\text{Sid}=\text{Sid}} \text{Student})$

(b) $\text{Reg_S} \leftarrow \pi_{\text{Sid}}(\text{Registered})$
 $\text{RESULT1} \leftarrow \pi_{\text{Sid}}(\text{Student})$
 $\text{RESULT2} \leftarrow \text{RESULT1} \cup \text{Reg_S}$
 $\text{RESULT} \leftarrow \pi_{\text{sname}}(\text{RESULT2} \bowtie_{\text{Sid}=\text{Sid}} \text{Student})$

(c) $\text{Reg_S} \leftarrow \pi_{\text{Sname}}(\text{Registered} \bowtie_{\text{sid}=\text{sid}} \text{Student})$

(d) $\text{Reg_S} \leftarrow \pi_{\text{Sid}}(\text{Registered})$
 $\text{RESULT1} \leftarrow \pi_{\text{Sid}}(\text{Student})$
 $\text{RESULT2} \leftarrow \text{RESULT1} - \text{Reg_S}$
 $\text{RESULT} \leftarrow \pi_{\text{sname}}(\text{RESULT2} \bowtie_{\text{Sid}=\text{Sid}} \text{Student})$

(e) $\text{Reg_S} \leftarrow \pi_{\text{Sid}}(\text{Registered})$
 $\text{RESULT1} \leftarrow \pi_{\text{Sid}}(\text{Student})$
 $\text{RESULT2} \leftarrow \text{RESULT1} \bowtie_{\text{sid}=\text{sid}} \text{Reg_S}$
 $\text{RESULT} \leftarrow \pi_{\text{sname}}(\text{RESULT2} \bowtie_{\text{Sid}=\text{Sid}} \text{Student})$

17) Consider the sequence of operations given below on the above relations.

- (I) $RESULT1 \leftarrow \pi_{cid} (\sigma_{deptid=3} (Course))$
- (II) $RESULT2 \leftarrow Registered \div RESULT1$
- (III) $RESULT \leftarrow \pi_{sname} (RESULT2 \bowtie_{sid=sid} Student)$

What will the above sequence of operations performed on the given relations produce?

- (a) Find the name of Students who are attached to department No.3 and follow some of the courses delivered by the same department.
- (b) Find the name of Students who have registered for any course offered by department No.3.
- (c) Find the name of students who have not registered for any course offered by department No.3.
- (d) Find the name of students who have registered for all the courses offered by department No. 3.
- (e) Find the name of students who have registered for only the courses offered by department No.3.

Answer the question from (18) – (21) considering part of the database which keeps track of courses and their enrolment details. For one course there can be many Training sessions and each Training session is identified by a unique Sno. Enrolment for each Training session is performed by several lecturers and a lecturer can be an enroller for several Training sessions.

Course(Cid, Title, Duration)
Training_session(Sno, Cid, Date, Lno, Scancel)
Lecturer(Lno, Lname, Salary)
Enrolment(Sno, Lno, Ecancel)

Scancel - indicates if the Training session is cancelled ("C" means cancelled, NULL means not cancelled).

Ecancel - indicates if the enrolment is cancelled ("C" when enrolment was cancelled, NULL if not cancelled)

18) Which of the following SQL statements would produce a list of all Courses which have at least two Training sessions?

- (a) `SELECT Cid, COUNT(*) FROM Training session WHERE Scancel IS NULL AND COUNT(*) >= 2;`
- (b) `SELECT Cid, Count(*) FROM Training session GROUP BY Cid HAVING COUNT(Scancel) =0 ;`
- (c) `SELECT Cid, COUNT(Cid) FROM Training session WHERE Scancel IS NULL GROUP BY Cid HAVING COUNT(*) >=2;`
- (d) `SELECT Cid, COUNT(*) FROM Course GROUP BY Cid HAVING COUNT(*) >= 2;`
- (e) `SELECT Cid, COUNT(*) FROM Training session WHERE Scancel IS NULL GROUP BY Cid HAVING COUNT(Cid) >=2;`

- 19) Which of the following SQL statements produce the course ids of all courses along with corresponding Training session numbers and the date on which each Training session starts? The course ids for which no Training session has yet been planned should also be listed.

(a) SELECT Cid, Sno, Date FROM Course c, Training session s WHERE c.Cid = s.Cid;
(b) SELECT Cid, Sno, Date FROM Course c LEFT OUTER JOIN Training session s ON c.Cid = s.Cid;
(c) SELECT Cid, Sno, Date FROM Course c RIGHT OUTER JOIN Training session s ON c.Cid = s.Cid;
(d) SELECT c.Cid, s.Sno, s.Date FROM Course c, Training session s WHERE c.Cid = s.Cid UNION ALL SELECT c.Cid, 0, NULL FROM Course c WHERE c.Cid NOT IN (SELECT Cid FROM Training session);
(e) SELECT Cid, Sno, Date FROM Course c, Training session s WHERE c.Cid = s.Cid UNION ALL SELECT Cid, Sno, Date FROM Training session WHERE Cid IS NULL;

- 20) Consider the following query.

```
SELECT Lname
From Lecturer
WHERE Salary=(SELECT MAX(Salary) FROM Lecturer);
```

Which of the following querie(s) is/are equivalent to the above query?

(a) SELECT Lname FROM Lecturer WHERE Salary >= ANY (SELECT Salary FROM Lecturer);
(b) SELECT Lname FROM Lecturer WHERE Salary IN (SELECT MAX(Salary) From Lecturer);
(c) SELECT Lname, MAX(Salary) FROM Lecturer GROUP BY Lname;
(d) SELECT Lname FROM Lecturer WHERE Salary >= All (SELECT Salary FROM Lecturer);
(e) SELECT L1.Lname FROM Lecturer As L1 LEFT OUTER JOIN Lecturer AS L2 ON L1.Salary < L2.Salary WHERE L2.salary IS NULL;

- 21) Which of the following SQL statements produce the Training_session number (Sno) of all Training_sessions for which none of the Enrolment has been cancelled?

(a) SELECT Distinct s.Sno FROM Training_session s, Enrolment e WHERE s.Sno = e.Sno AND e.Ecancel IS NULL;
(b) SELECT Distinct s.Sno FROM Training_session s, Enrolment e WHERE s.Sno = e.Sno AND e.Ecancel IS NOT NULL;
(c) SELECT s.Sno FROM Training_session s WHERE s.Sno NOT IN (SELECT e.Sno FROM Enrolment e WHERE e.Ecancel IS NOT NULL);
(d) SELECT s.Sno FROM Training_session s WHERE s.Sno In (SELECT e.Sno FROM Enrolment e WHERE e.Ecancel IS NULL);
(e) SELECT s.Sno FROM Training_session s WHERE NOT EXISTS (SELECT * FROM Enrolment e WHERE e.Sno = s.Sno AND e.Ecancel IS NOT NULL);

Answer the question from (22) – (26) considering a database which keeps track of people and their share holdings. Assume that there are no records in the tables other than the ones given here.

Person

| <u>NIC</u> | Name | Cname | Salary |
|------------|------------|-----------|--------|
| 657324567v | A.J.Perera | Amazon | 65000 |
| 784589012v | S.A.Dias | Yahoo | 70000 |
| 725678234v | R.T.Zoysa | Microsoft | 50000 |
| 683456785v | P.M.Costa | Google | 55000 |
| 754567890v | W.A.Yapa | Yahoo | 75000 |

Holding

| <u>NIC</u> | <u>Cname</u> | Numshares |
|------------|--------------|-----------|
| 657324567v | Yahoo | 3000 |
| 754567890v | Microsoft | 4500 |
| 683456785v | Google | 4500 |
| 725678234v | Yahoo | 5000 |
| 784589012v | Google | 6000 |
| 683456785v | Amazon | 3500 |
| 725678234v | Microsoft | 4000 |

Company

| <u>Cname</u> | Location |
|--------------|------------|
| Amazon | California |
| Yahoo | New Jersy |
| Google | Newyork |
| Microsoft | California |

A person is uniquely identified by a national identity card number (NIC). A company is uniquely identified by its name. Each person is employed by exactly one company (denoted by Cname in Person table), but may hold any number of different shares from different companies. The Holding table was defined as given below:

```

Create table Holding {
    NIC char(10) ,
    Cname Varchar(25),
    Numshares Integer CHECK (Numshares >= 2500),
    Primary key (NIC, Cname)
};

```

22) Which SQL statement defines the FOREIGN KEY constraints of the Holding table?

- (a) ALTER TABLE Holding ADD CONSTRAINTS person_fk FOREIGN KEY (NIC) REFERENCES Person(NIC), company_fk FOREIGN KEY (Cname) REFERENCES Company(Cname);
- (b) ALTER TABLE Holding ADD FOREIGN KEY (NIC) REFERENCES Person(NIC);
ALTER TABLE Holding ADD FOREIGN KEY (Cname) REFERENCES Company(Cname);
- (c) ALTER TABLE Holding Add CONSTRAINT person_fk FOREIGN KEY (NIC) REFERENCES Person(NIC), ADD CONSTRAINT company_fk FOREIGN KEY (Cname) REFERENCES Company(Cname);
- (d) ALTER TABLE Holding ADD CONSTRAINT person_fk FOREIGN KEY (NIC) REFERENCES Person(NIC);
ALTER TABLE Holding ADD CONSTRAINT company_fk FOREIGN KEY (Cname) REFERENCES Company(Cname);
- (e) ALTER TABLE Holding ADD CONSTRAINT person_fk FOREIGN KEY (NIC) REFERENCES Person(NIC), CONSTRAINT company_fk FOREIGN KEY (Cname) REFERENCES Company(Cname);

23) Suppose we wish to find the NIC numbers of the persons who do not own shares of their employers. Which of the following queries will return the correct set of NIC numbers?

(I) SELECT p.NIC FROM Person p, Holding h
WHERE p.NIC = h.NIC and p.Cname <> h.Cname;

(II) SELECT p.NIC FROM Person p
WHERE (p.NIC,p.Cname) NOT IN
(SELECT h.NIC, p.Cname FROM Holding h);

(III) SELECT NIC FROM Person p
WHERE Cname <> ALL
(SELECT Cname FROM Holding h
WHERE p.NIC = h.NIC);

- | | | |
|---------------------|-------------------|--------------------|
| (a) II only | (b) I and II only | (c) I and III only |
| (d) II and III only | (e) III only | |

24) Suppose we wish to find the average salary of the persons who own more than 3500 shares of Microsoft or more than 3500 shares of Yahoo. Which of the following queries will correctly compute the desired average?

(I) SELECT AVG(salary) FROM Person
WHERE NIC IN (SELECT NIC FROM Holding WHERE (Cname = 'MicroSoft' OR
Cname = 'Yahoo') AND Numshares > 3500);

(II) SELECT AVG(salary) FROM Person p, Holding h
WHERE p.NIC = h.NIC AND ((Cname = ' MicroSoft' AND Numshares > 3500) OR
(Cname = 'Yahoo' AND Numshares > 3500));

(III) SELECT AVG(salary) FROM Person p
WHERE Exists (SELECT * FROM Holding h WHERE (h.Cname = ' MicroSoft' OR
h.Cname = 'Yahoo') AND (Numshares > 3500) AND (p.NIC = h.NIC));

- | | | |
|------------|---------------------|-------------------|
| (a) I only | (b) II only | (c) I and II only |
| (d) All | (e) II and III only | |

25) Consider Holding table with the enforcement of integrity constraints as stated in question (22).

The following sequence of modifications is executed in the given order on Holding table.

- (i) INSERT INTO Holding VALUES ('754567890v', 'Toshiba', 4500);
- (ii) UPDATE Holding SET Numshares = 2000 WHERE NIC = '683456785v';
- (iii) INSERT INTO Holding VALUES ('767497235v', 'Yahoo', 3500);
- (iv) DELETE FROM Holding WHERE Cname = 'Microsoft';

At the end of these statements, the sum of the shares over all the tuples in Holding table is:

- | | | |
|------------|------------|-----------|
| (a) 25,500 | (b) 30,000 | (c) 22000 |
| (d) 26000 | (e) 26,500 | |

26) Which of the following statement(s) is/are true about views?

- | |
|---|
| (a) View can be created as read only. |
| (b) A view can be created as a join on two or more tables. |
| (c) A view cannot have an ORDER BY clause in the SELECT statement. |
| (d) A view cannot be updated when it is created with a GROUP BY clause in the SELECT statement. |
| (e) A view must have aliases defined for the column names in the SELECT statement. |

27) Consider that the values for relation R(a,b,c) is as given below.

| a | b | c |
|---|---|---|
| 1 | 1 | 3 |
| 1 | 2 | 3 |
| 2 | 1 | 4 |
| 2 | 3 | 5 |
| 2 | 4 | 1 |
| 3 | 2 | 4 |

A view called Test is defined as:

CREATE VIEW Test AS SELECT a+b AS d, c FROM R;

Which of the following tuples would be in the result of the query:

SELECT d, SUM(c) FROM Test GROUP BY d HAVING COUNT(*) <> 1;

- | | | |
|-----------|------------|-----------|
| (a) (2,5) | (b) (3,12) | (c) (5,9) |
| (d) All | (e) None | |

Consider a part of the company database which has the two tables Employee and

Department to answer questions from (28) – (30).

Employee (emp_id CHAR(05), ename VARCHAR(25), salary NUMBER NOT NULL, dept_id NUMBER FOREIGN KEY (dept_id) REFERENCES Department(dept_id) mgr_id CHAR(05));

Department (dept_id NUMBER, dept_name VARCHAR (20))

- 28) A view called **Emp_Dept** that contains three columns from the Employee and Department tables: *emp_id*, *ename* and *dept_name* have been created. It is required to modify the view by adding a fourth column, *mgr_id* from the Employee tables. Which of the following SQL statement(s) would one execute to accomplish this task?

- (a) ALTER VIEW Emp_Dept (ADD mgr_id char(05));
- (b) MODIFY VIEW Emp_Dept (ADD mgr_id char(05));
- (c) ALTER VIEW Emp_Dept AS SELECT emp_id, ename, dept_name, mgr_id FROM employee e, department d WHERE e.dept_id = d.dept_id;
- (d) CREATE OR REPLACE VIEW Emp_Dept AS SELECT emp_id, ename, dept_name, mgr_id FROM employee e, department d WHERE e.dept_id = d.dept_id;
- (e) DROP AND REPLACE VIEW Emp_Dept AS SELECT emp_id, ename, dept_name, mgr_id FROM employee e, department d WHERE e.dept_id = d.dept_id;

- 29) It is necessary to create a view called **Salary_Range** to show the department id, minimum salary, and maximum salary paid in that department, only if the minimum salary is less than Rs.40000 and the maximum salary is more than Rs. 85000. Which of the following SQL statement(s) would one execute to create this view?

- (a) CREATE VIEW Salary_Range (dept_id, min_salary, max_salary) AS SELECT dept_id, MIN(salary), MAX(salary) FROM employees WHERE MIN(salary) < 40000 AND MAX(salary) > 85000;
- (b) CREATE VIEW Salary_Range (dept_id, min_salary, max_salary) AS SELECT dept_id, MIN(salary), MAX(salary) FROM employees WHERE MIN(salary) < 40000 AND MAX(salary) > 85000 GROUP BY dept_id;
- (c) CREATE VIEW Salary_Range (dept_id, min_salary, max_salary) AS SELECT dept_id, MIN(salary), MAX(salary) FROM employees HAVING MIN(salary) < 40000 AND MAX(salary) > 85000;
- (d) CREATE VIEW Salary_Range (dept_id, min_salary, max_salary) AS SELECT dept_id, MIN(salary), MAX(salary) FROM employees GROUP BY dept_id HAVING MIN(salary) < 40000 AND MAX(salary) > 85000;
- (e) CREATE VIEW Salary_Range (dept_id, min_salary, max_salary) AS SELECT dept_id, MIN(salary), MAX(salary) FROM employees GROUP BY dept_id, salary HAVING MIN(salary) < 40000 AND MAX(salary) > 85000;

- 30) It is necessary to create a view called Temp that allows the user to insert rows through the view. Which Statement(s), when used to create the Temp view, allows the user to insert rows?

- (a) CREATE VIEW Temp AS SELECT emp_id, emp_name, dept_id
FROM employees WHERE mgr_id IN (102, 120);
- (b) CREATE VIEW Temp AS SELECT emp_id, emp_name, salary, dept_id
FROM employees WHERE mgr_id IN (102, 120);
- (c) CREATE VIEW Temp AS SELECT dept_id, SUM(sal) TOTALSAL
FROM employees WHERE mgr_id IN (102, 120) GROUP BY dept_id;
- (d) CREATE VIEW Temp AS SELECT DISTINCT dept_id, emp_name, job_id,
FROM employees;
- (e) CREATE VIEW Temp AS SELECT emp_id, emp_name, salary, dept_name
FROM employees e, Department d WHERE e.dept_id = d.dept_id;

- 31) Suppose relation $R(A,B,C,D,E)$ has the following functional dependencies:

$A \rightarrow B$ $B \rightarrow C$ $BC \rightarrow A$
 $A \rightarrow D$ $E \rightarrow A$ $D \rightarrow E$

Which of the following is **not** a key of R?

- | | | |
|-------|----------|-------|
| (a) A | (b) E | (c) B |
| (d) D | (e) B, C | |

Consider the relation $R(A,B,C,D,E)$ with the given Functional Dependencies to answer questions from (32) – (34).

$A \rightarrow B$ $BC \rightarrow D$ $A \rightarrow C$

- 32) What is/are the candidate key(s) of the relation R?

- | | | |
|--------------|-----------|-------|
| (a) (C,D) | (b) A | (c) B |
| (d) (B,C, D) | (e) (B,C) | |

- 33) Which of the following statements is/are true?

- | |
|---|
| (a) The Functional dependency $A \rightarrow C$ violates 2 NF. |
| (b) The best normal form the relation R satisfies is 1 NF. |
| (c) The Functional dependency $BC \rightarrow D$ violates 2 NF. |
| (d) The best normal form the relation R satisfies is 2 NF. |
| (e) The Functional dependency $BC \rightarrow D$ violates 3 NF. |

34) Which of the following statements is/are true with respect to the decomposition of R?

- (a) R1(A,B), R2(B,C,D) is lossless and dependency preserving.
- (b) R1(A,B), R2(A,C,D) is lossless and dependency preserving.
- (c) R1(A,B,C) , R2(B,C,D) is lossless and dependency preserving.
- (d) R1(A,C), R2(B,C,D) is lossless and dependency preserving.
- (e) There is no lossless and dependency preserving decomposition for R.

35) Consider the relation Channelling (Doctor#, Patient#, Date, Treat_code, Charge) with functional dependencies

(Doctor#, Patient#, Date)→(Diagnosis, Treat_code, Charge)

(Treat_code)→(Charge)

Which of the following statements is/are true?

- (a) The Functional dependency (Treat_code)→(Charge)violates 2NF.
- (b) The Functional dependency (Treat_code)→(Charge) violates 3NF.
- (c) The best normal form the relation Channelling satisfies is 3 NF.
- (d) The best normal form the relation Channelling satisfies is 2NF.
- (e) The decomposition R1(Doctor#, Patient#, Date, Diagnosis}, R2 (Treat_code, Charge) is dependency preserving

36) Consider the following relations for an order-processing application database.

ORDER (Ord#, CustId#, Ord_date, Total_amount)

ORDER-ITEM (Ord#, Item#, Qty_ordered, Total_price, Discount%)

Assume that each item has a different discount. The Total_price refers to the total value of one item. Ord_date is the date on which the order was placed, and the Total_amount is the amount of the order. If natural join is applied on the relations Order-Item and Order in this database, which of the following statements is/are true with respect to the resulting relation schema R?

- (a) The key of R is Ord#, Item#, Cust#
- (b) The key of R is Ord#, Item#.
- (c) The best normal form the relation R satisfies is 1NF.
- (d) The best normal form the relation R satisfies is 2NF.
- (e) The functional dependency Ord# → Ord_date, Total_amount violates 3NF.

Consider the following scenario to answer the question from (37) – (40).

A bank maintains **Bank** (**CustId**, **Name**, **Address**, **AcctType**, **InterestRate**, **Balance**) **table** as a part of its database. The DBA of the bank creates five users namely *User1*, *User2*, *User3*, *User4* and *User5*. He initially grants privileges to all the other users except *User5* and to this end the following SQL statements are issued by the DBA :

- **CREATE VIEW Temp AS
SELECT * FROM Bank
WHERE Balance > 50000
WITH CHECK OPTION;**
- **GRANT CREATE VIEW, CREATE TABLE TO user1;**
- **GRANT INSERT, SELECT ON Temp TO user2;**
- **GRANT SELECT, UPDATE(AcctType, InterestRate) ON Bank TO user3
WITH GRANT OPTION;**
- **GRANT UPDATE (Balance) ON Temp TO User4;**

37) Consider the following SQL statements issued by the following users as in (I) – (III):

- (I) **User1:** CREATE VIEW Temp_2
 AS SELECT *
 FROM Temp WHERE AcctType = 'Savings';
- (II) **User3:** GRANT SELECT ON Temp TO User5;
- (III) **User4:** UPDATE Temp SET Balance = Balance * 1.1
 WHERE CustId IN (SELECT Custid FROM Temp
 WHERE AcctType = 'Savings');

Which of the above SQL statements will get executed successfully?

- | | | |
|-------------------|-------------|--------------------|
| (a) I Only | (b) II Only | (c) I and III Only |
| (d) I and II Only | (e) None | |

38) Which of the following statement(s) will justify the correct answer for (37) above?

- | |
|--|
| (a) User1 is unable to execute (I) as sufficient privileges have not been granted to create the view. |
| (b) User3 is unable to execute the statement (II) as select privilege on Bank is not sufficient to perform select on Temp. |
| (c) User4 is able to execute (III) with the update privilege on Balance. |
| (d) User4 is unable to execute (III) as sufficient privileges have not been granted to User4 for this update. |
| (e) User3 is able to execute (II) as select privilege on Bank will in turn enable select privilege on its view Temp. |

39) Consider the following SQL statements issued by the given user as in (I) – (III):

- (I) **User1** : CREATE TABLE Loan (Loanid CHAR(05) PRIMARY KEY,
Amount INTEGER NOT NULL, CustId CHAR(05)
REFERENCES Bank);
- (II) **User3**: UPDATE Bank
SET InterestRate = InterestRate + 0.75
WHERE AcctType = 'Savings';
- (III) **User2** : INSERT INTO Temp
VALUES ('10012885', 'J.K.Perera', No.43, Main Street, Kandy',
'Savings', 10%, 45000);

Which of the above SQL statements will get executed successfully?

- | | | |
|------------|-------------------|---------------------|
| (a) I Only | (b) I and II Only | (c) II and III Only |
| (d) All | (e) II Only | |

40) DBA executes the following SQL statement for User5.

GRANT CREATE TABLE, CREATE VIEW TO user5;

User5 in turn issues the following Statements in the given order.

- (I) CREATE TABLE Employee(Empid CHAR(05) PRIMARY KEY, Name
VARCHAR(25),Designation VARCHAR(15), Salary NUMBER);
- (II) CREATE VIEW Emp AS SELECT Empid, Name, Designation FROM
Employee ;
- (III) GRANT SELECT ON Emp TO user2;
- (IV) GRANT INSERT ON employee TO User1;

Which of the above SQL statements will get executed successfully?

- | | | |
|------------|-------------------|-------------------|
| (a) I Only | (b) I and II Only | (c) I and IV Only |
| (d) All | (e) None | |
