



UNIVERSITY OF COLOMBO, SRI LANKA



UCSC UNIVERSITY OF COLOMBO SCHOOL OF COMPUTING

(Successor to the Institute of Computer Technology (ICT), University of Colombo)



DEGREE OF BACHELOR OF INFORMATION TECHNOLOGY (EXTERNAL)

Academic Year 2003/2004 – 2nd Year Examination – Semester 4

IT4501: Advanced Database Management Systems

PART 1 - Multiple Choice Question Paper

25th July, 2004

(ONE AND A HALF HOURS)

Important Instructions:

- The duration of the paper is **1 ½ (one and a half) hours**.
- The medium of instruction and questions is English.
- The paper has **35 questions** and **09 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 -1 (*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.**
- **Non-programmable Calculators may be used.**

Use the following description to answer questions 1-4.

A company uses a relational database to store data about their employees and the projects which they are working on. The database relations of this system are given below.

Employee(empno, ename, address, salary, designation)

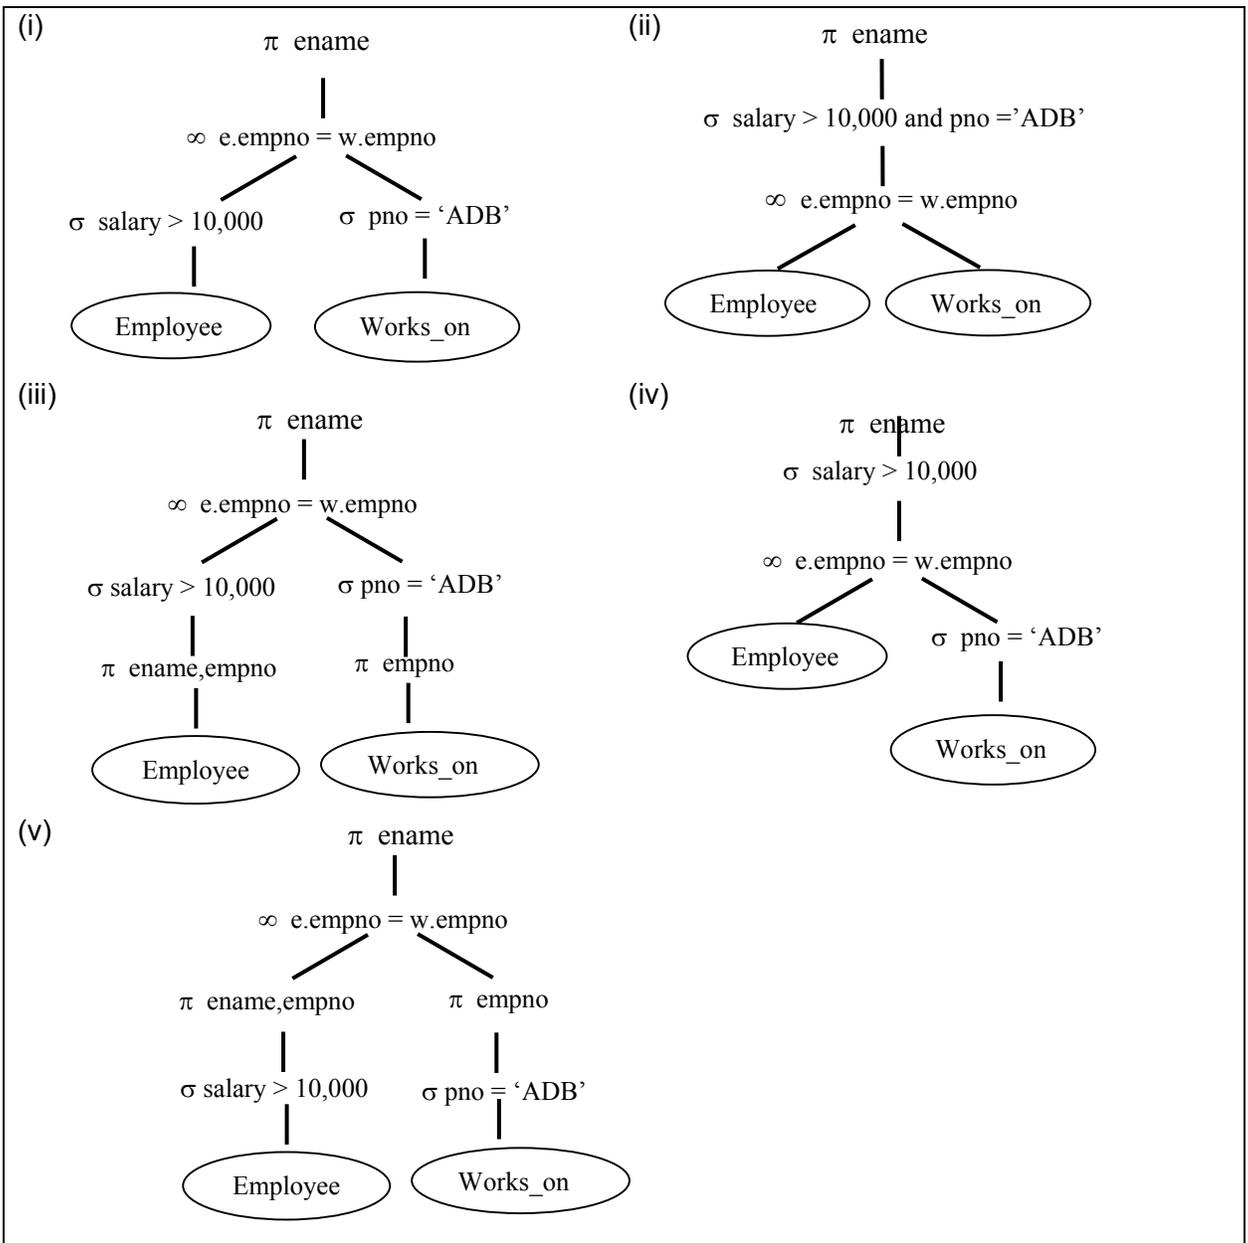
Project(pno, pname, plocation)

Works_on(empno, pno, hours)

The following SQL query has been written to retrieve the names of the employees who are working on the project 'ADB' and are paid more than Rs: 10,000 as salary.

```
SELECT ename
FROM Employee e, Works_on w
WHERE e.empno = w.empno and w.pno = 'ADB' and e.salary > 10000;
```

The following query trees have been drawn for the above query.



1) Which of the following query trees will correctly produce the results given in the above query?

- | | | |
|----------|----------|-----------|
| (a) (i) | (b) (ii) | (c) (iii) |
| (d) (iv) | (e) (v) | |

2) Which of the following best represent(s) the optimal query tree for the above query?

- | | | |
|----------|----------|-----------|
| (a) (i) | (b) (ii) | (c) (iii) |
| (d) (iv) | (e) (v) | |

3) Which of the following represent(s) the least optimised query tree for the above query?

- | | | |
|----------|----------|-----------|
| (a) (i) | (b) (ii) | (c) (iii) |
| (d) (iv) | (e) (v) | |

4) Which of the following represent(s) the query trees in the order of least optimised to most optimised?

- | | | |
|---------------------|----------------------|---------------------|
| (a) (i), (v), (iii) | (b) (iii), (iv), (i) | (c) (ii), (iv), (v) |
| (d) (iv), (ii), (i) | (e) (ii), (iv), (i) | |

5) Select the query operations which are in the order of **most** expensive to **least** expensive?

- | | |
|--|---|
| (a) cartesian product, union, select | (b) intersection, cartesian product, project |
| (c) cartesian product, intersection, project | (d) set difference, cartesian product, select |
| (e) union, select, intersection | |

Use the following description to answer questions 6-8.

Due to the evolutionary nature of the engineering design environment, solutions of the design objects are stored as versions. Version data relevant to each design object is *versionNo*, *createdTime*, *parentversionNo* and *status*. In this scenario, a bicycle is considered as the design object with the attributes *typeName*, *weight*, *material*, *colour* and *speed*. Version data can be defined in a class called *Version*. Bicycle information can be defined in a class called *Bicycle*. Construction of bicycle versions using a class called *BicycleVersion* requires combining bicycle data and version data together. Five suggested class definitions for this purpose are as given below:

- | | |
|--|---|
| (i) Class <code>BicycleVersion</code> extends <code>Version</code> {
<code>Bicycle bicycle</code> ;
} | (ii) Class <code>BicycleVersion</code> extends <code>Bicycle</code> {
<code>Version version</code> ;
} |
| (iii) Class <code>BicycleVersion</code> {
<code>Version version</code> ;
<code>Bicycle bicycle</code> ;
} | (iv) Class <code>BicycleVersion</code> {
attribute <code>String versionNumber</code> ;
attribute <code>Time createdTime</code> ;
attribute <code>String parentversionNo</code> ;
attribute <code>Boolean status</code> ;
<code>Bicycle bicycle</code> ;
} |
| (v) Class <code>BicycleVersion</code> extends <code>Bicycle</code> , <code>Version</code> {
} | |

6) Which of the above class definitions will enable one to combine version data and bicycle data to create `BicycleVersion`?

- | | | |
|---------------------|-----------------------|--------------|
| (a) (ii) & (v) only | (b) (iii) & (iv) only | (c) (v) only |
| (d) (i) & (ii) only | (e) all of them | |

7) Which of the above class definitions will enable one to retrieve bicycle design data (object) easily from BicycleVersion?

- | | | |
|----------|----------|-----------|
| (a) (i) | (b) (ii) | (c) (iii) |
| (d) (iv) | (e) (v) | |

8) Which of the following statements is/are correct?

- | |
|---|
| <p>(a) Class definitions (i) and (ii) use an inheritance relationship.</p> <p>(b) Multiple inheritance in option (v) describes the situation where the sub-classes version and bicycle inherit the superclass BicycleVersion.</p> <p>(c) Class definitions (i) and (ii) use association.</p> <p>(d) Class definition (iii) uses aggregation.</p> <p>(e) Class definition (iv) will cause re-definition of version data for creation of any new versionable design object class.</p> |
|---|

Consider the following schedule S1 for questions 9 to 13. Here the transaction may either commit or rollback.

S1	
<u>T1</u>	<u>T2</u>
Read(X); S = X + 5;	Read(X);
Read(X); X = S + X; Write(X);	X = X+10; Write(X);
Commit/Rollback	Commit/Rollback

9) If the schedule S1 was executed in read uncommitted mode, what problem(s) could one face?

- | | | |
|-----------------|----------------|-----------------------|
| (a) Lost update | (b) Dirty read | (c) Incorrect Summary |
| (d) Deadlock | (e) Phantom | |

10) If the schedule S1 was executed in read committed mode, what problem(s) could one face?

- | | | |
|-----------------|----------------|-----------------------|
| (a) Lost update | (b) Dirty read | (c) Incorrect Summary |
| (d) Deadlock | (e) Phantom | |

11) If the schedule S1 was executed in repeatable read mode, what problem(s) could one face?

- | | | |
|-----------------|----------------|-----------------------|
| (a) Lost update | (b) Dirty read | (c) Incorrect Summary |
| (d) Deadlock | (e) Phantom | |

12) If the schedule S1 was executed in serializable mode, what problem(s) could one face?

- | | | |
|-----------------|----------------|-----------------------|
| (a) Lost update | (b) Dirty read | (c) Incorrect Summary |
| (d) Deadlock | (e) Phantom | |

13) Which of the following is a/are property/properties of a transaction?

- | | | |
|-----------------|-----------------|----------------|
| (a) concurrency | (b) persistency | (c) durability |
| (d) sharing | (e) multi-user | |

14) Which protocols or set of rules is/are used to guarantee serializability?

- | | | |
|---------------|-------------------|-------------------|
| (a) two-phase | (b) three-phase | (c) time-stamping |
| (d) locking | (e) multi-version | |

15) A transaction T_i is assigned a unique timestamp $TS(T_i)$ before starting the execution of the transaction. $WRITE-TS(X)$ denotes the largest timestamp of any transaction that executes $WRITE(X)$ successfully and $READ-TS(X)$ denotes the largest timestamp of any transaction that executes $READ(X)$ successfully. When transaction T_i issues $READ(X)$ or $WRITE(X)$ under timestamp protocol, rollback could occur for T_i if the timestamp satisfies (a) certain condition(s). Which of the following is/are the condition(s)?

- | | |
|---|--|
| (a) For $READ(X)$ if $TS(T_i) < WRITE-TS(X)$ | (b) For $READ(X)$ if $TS(T_i) \geq WRITE-TS(X)$ |
| (c) For $WRITE(X)$ if $TS(T_i) < READ-TS(X)$ | (d) For $WRITE(X)$ if $TS(T_i) \geq WRITE-TS(X)$ |
| (e) For $WRITE(X)$ if $TS(T_i) < WRITE-TS(X)$ | |

16) The main techniques for recovery from failures are

- | | | |
|----------------------|-----------------------|----------------------|
| (a) deferred update. | (b) immediate update. | (c) serializability. |
| (d) timestamping. | (e) shadow paging. | |

Consider the following schedule S1 for questions 17 and 18. Here initial values of A, B, C are 1000, 2000 and 700 respectively.

S1	
T1	T2
Read(A)	Read(C)
A = A-50	
Write(A)	C = C-100
Read(B)	
B = B+50	
Write(B)	
Commit	Write(C)
	Commit

17) If schedule S1 was executed in differed update mode, what are the database values for A, B and C immediately before executing the Commit statement for transaction T1?

- | | | |
|---------------------------|--------------------------|---------------------------|
| (a) A=950, B=2000, C=700 | (b) A=950, B=2050, C=700 | (c) A=1000, B=2000, C=700 |
| (d) A=1000, B=2050, C=700 | (e) A=950, B=2050, C=600 | |

18) If schedule S1 was executed in immediate update mode, what are the database values of A, B and C immediately before executing the Commit statement for transaction T1?

- | | | |
|---------------------------|--------------------------|---------------------------|
| (a) A=950, B=2000, C=700 | (b) A=950, B=2050, C=700 | (c) A=1000, B=2000, C=700 |
| (d) A=1000, B=2050, C=700 | (e) A=950, B=2050, C=600 | |

Use the following description to answer questions 19 and to 210.

The following 4 relations are part of a relational database schema of a company created by its Database Administrator (DBA).

Employee(empno, ename, address, sex, salary, designation, managerno, deptno)

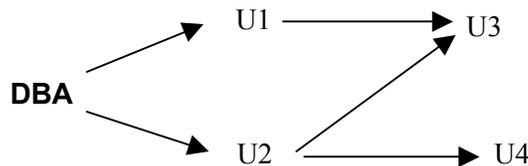
Department(deptno, dname, managerno, dlocation)

Project(pnumber, pname, plocation, deptno)

Works_on(empno, pnumber, hours)

Employee1(empno, ename, address, salary, designation, deptno) is a view that is to be created to represent a vertical subset of *employee* relation which has been created by an authorised user.

Privileges are granted to user accounts U1, U2, U3 and U4 as described below and shown in the authorisation diagram.



- (A) U1 and U2 are granted authorisation by DBA, U3 is granted authorisation by both U1 and U2 while U4 is granted authorisation by U2.
- (B) U1 can retrieve all the attributes from Employee and update and retrieve all the attributes from Project relation.
- (C) U2 can retrieve the attributes empno, ename, address, salary, designation and deptno from Employee, and update address, salary and designation on Employee.
- (D) Later on, the privileges of U2 are revoked by the DBA.

19) Based on the privileges granted in (B) above, which of the following operations can be performed by U1?

- (a) CREATE VIEW emp_project AS
SELECT empno, ename, designation, pname
FROM Employee e, Project p, Works_on w
WHERE e.empno = w.empno and p.pnumber = w.pnumber;
- (b) SELECT * FROM emp_project;
- (c) DELETE FROM Project WHERE pname = 'ADB';
- (d) SELECT COUNT(*) FROM Employee WHERE designation = 'clerk';
- (e) GRANT DELETE ON Project TO U3;

20) Which of the following could assist one to achieve the privileges of the user U2?

- (a) GRANT SELECT(empno, ename, address, salary, designation, deptno), UPDATE (address, salary, designation) ON Employee TO U2;
- (b) CREATE VIEW Employee1 AS SELECT empno, ename, address, salary, designation, deptno FROM Employee;
- (c) GRANT SELECT(empno, ename, address, salary, designation, deptno), UPDATE (address, salary, designation) ON Employee1 TO U2 WITH GRANT OPTION;
- (d) GRANT SELECT, UPDATE (address, salary, designation) ON Employee1 TO U2 WITH GRANT OPTION;
- (e) GRANT SELECT(empno, ename, address, salary, designation, deptno), UPDATE (address, salary, designation) ON Employee TO U2 WITH GRANT OPTION;

21) In order to avoid abandoned privileges, which of the following would prevent the revoke of privileges on U2?

- (a) REVOKE SELECT (empno, ename, address, salary, designation, deptno), UPDATE(address, salary, designation) ON Employee FROM U2 CASCADE;
- (b) REVOKE SELECT ,UPDATE(address, salary, designation) ON Employee1 FROM U2 CASCADE;
- (c) REVOKE SELECT (empno, ename, address, salary, designation, deptno), UPDATE(address, salary, designation) ON Employee FROM U2 RESTRICT;
- (d) REVOKE SELECT, UPDATE(address, salary, designation) ON Employee1 FROM U2 RESTRICT;
- (e) REVOKE GRANT OPTION FOR SELECT, UPDATE(address, salary, designation) ON Employee FROM U2 RESTRICT;

22) Which of the following statements is/are correct about a logic programming language such as Prolog/Datalog?

- (a) Facts in Prolog declare things which are always true unconditionally.
- (b) Prolog can indeed very easily use recursive definitions.
- (c) The declarative approach makes programming in Prolog more difficult than in typical procedural programming languages such as C or Pascal.
- (d) Prolog specifies a method for distinguishing between atoms and variables where atoms start with upper-case letters and variables start with lower-case letters.
- (e) Datalog is a procedural query language that is based on Prolog programming language.

23) The following facts and rules have been defined in logic programming?

parent(nanda, sunil).
parent(sarath, sunil).
parent(jeeva, nanda).
parent(sunil, soma).
female(nanda).
male(sarath).
male(sunil).
female(jeeva).
female(soma).
mother(X,Y) :- parent(X,Y), female(X). // X is Y's mother
father(X,Y) :- parent(X,Y), male(X). // X is Y's father

Which of the following rules in logic programming correctly define(s) the grandmother relation *grandmother(X,Y)* where X is the grandmother of Y?

- (a) grandmother(X,Y) :- parent (X,Z), parent(Y,Z), female(X).
- (b) grandmother(X,Y) :- mother(X,Z), parent(Z,Y).
- (c) grandmother(X,Y) :- parent (X,Z), parent(Z,Y), female(X).
- (d) grandmother(X,Y) :- mother(X,Z), mother(Z,Y) ; father(Z,Y).
- (e) grandmother(X,Y) :- parent(X,Z), mother(Z,Y).

24) Data partition/relation fragmentation can be applied when different relations are distributed amongst different disks, or different sections of a relation are distributed amongst different disks. Which data partitioning scheme(s) is/are generally used to evenly distribute the data across fragments?

- (a) horizontal partitioning
- (b) mixed partitioning
- (c) round-robin
- (d) range partitioning
- (e) hash partitioning

Consider the following information to answer questions 25-33.

A university Faculty has five departments (D1-D5) offering practical classes. They are D1 - Physics, D2 - Chemistry, D3 - Statistics, D4 - Zoology and D5 - Botany.

Data of four relations (R1-R4) namely,

R1 - Lecturer(lno, name, department), R2 - Student(sno, sname),
R3 - Course(cno, cname, lno) and R4 - Attendance(cno, sno, lab, day)

are to be managed using a DBMS for use by the five departments. The database administrator wants to allocate tables to suitable sites (departments) without fragmenting them. The allocation strategy is to increase the amount of local access for each site and minimise the network traffic. Three transactions (T1-T3) where

T1 (prepare an attendance sheet for a given course),

T2 (list of labs in use on a given day) and

T3 (list of departments and for a given department, the list of courses it offers)

are issued by each of the specified departments to access data of specified tables as read only on a specified frequency of occurrence as given below.

Transaction	Sites	Frequency	Table Accessed
T1	D4, D5	3	read R2 (once), R3 (once) & R4 (once)
T2	D1, D2	5	read R4 (once)
T3	D1, D2, D3	1	read R1 (twice) & R3 (once)

25) If a centralised approach is to be used, the most suited site(s) to allocate all the relations is/are

- | | | |
|----------------------------|----------------------|----------------------|
| (a) only D1. | (b) only D2. | (c) either D1 or D2. |
| (d) either D1 or D2 or D3. | (e) either D4 or D5. | |

26) If a distributed approach is used without replication, the most suited site(s) to allocate the relation R1 is/are

- | | | |
|----------------------------|--------------------|--------------|
| (a) only D1. | (b) only D2. | (c) only D3. |
| (d) either D1 or D2 or D3. | (e) D1, D2 and D3. | |

27) If a distributed approach is used with replication, the most suited site(s) to allocate the relation R1 is/are

- | | | |
|----------------------------|--------------------|--------------|
| (a) only D1. | (b) only D2. | (c) only D3. |
| (d) either D1 or D2 or D3. | (e) D1, D2 and D3. | |

28) If a distributed approach is used without replication, the most suited site(s) to allocate the relation R2 is/are

- | | | |
|----------------|------------------------|----------------------|
| (a) only D4. | (b) only D5. | (c) either D4 or D5. |
| (d) D4 and D5. | (e) neither D4 nor D5. | |

29) If a distributed approach is used with replication, the most suited site(s) to allocate the relation R2 is/are

- | | | |
|----------------|------------------------|----------------------|
| (a) only D1. | (b) only D2. | (c) either D4 or D5. |
| (d) D4 and D5. | (e) neither D4 nor D5. | |

30) If a distributed approach is used without replication, the most suited site(s) to allocate the relation R3 is/are

- | | | |
|----------------------------|--|----------------------|
| (a) only D4. | (b) only D5. | (c) either D4 or D5. |
| (d) either D1 or D2 or D3. | (e) either D1 or D2 or D3 or D4 or D5. | |

- 31) If a distributed approach is used with replication, the most suited site(s) to allocate the relation R3 is/are
- | | | |
|----------------------|----------------------------|----------------------------|
| (a) either D4 or D5. | (b) D4 and D5. | (c) either D1 or D2 or D3. |
| (d) D1, D2 and D3. | (e) D1, D2, D3, D4 and D5. | |
- 32) If a distributed approach is used without replication, the most suited site(s) to allocate the relation R4 is/are
- | | | |
|--------------|----------------------|----------------------|
| (a) only D1. | (b) only D2. | (c) either D1 or D2. |
| (d) only D4. | (e) either D4 or D5. | |
- 33) If a distributed approach is used with replication, the most suited site(s) to allocate the relation R4 is/are
- | | | |
|----------------------|------------------------|----------------------|
| (a) either D1 or D2. | (b) D1 and D2. | (c) either D4 or D5. |
| (d) D4 and D5. | (e) D1, D2, D4 and D5. | |
- 34) Which of the following statements is/are correct?
- | |
|--|
| <p>(a) <i>Polyinstantiation</i> is a technique used by database security professionals to avoid inference capabilities.</p> <p>(b) <i>Update</i> authorisation allows modification and deletion of data.</p> <p>(c) Creation of a view requires <i>resource</i> authorization.</p> <p>(d) <i>Data pollution</i> is a security technique used in statistical databases.</p> <p>(e) <i>Delete</i> authorisation and <i>drop</i> authorisation are similar since they both allow deletion of tuples only.</p> |
|--|
- 35) Which of the following statements is/are correct about a Temporal Database?
- | |
|---|
| <p>(a) A Temporal Database contains time varying data including historical and current data.</p> <p>(b) Queries on time-varying data can efficiently be performed using SQL.</p> <p>(c) Previous states in a temporal database are preserved and no modification to the past is allowed.</p> <p>(d) Each tuple in a temporal database must include at least one timestamp.</p> <p>(e) A Temporal database can be considered as a snapshot database.</p> |
|---|
