

NAPIER UNIVERSITY
SCHOOL OF COMPUTING

SECOND DIET (UK) - SESSION 2002-03

DATABASE SYSTEMS

MODULE NO: CO22001

DATE: Aug 2003

EXAM TIME: 2 HOURS

START TIME: HOURS

FINISH TIME: HOURS

EXAMINERS:
G. RUSSELL
U. PRISS

QUESTION PAPER DATA

Number of pages – TWENTY

Number of questions – FORTY

INSTRUCTIONS TO CANDIDATES

Select ONE from (a) to (e)

1. In relational algebra, which of the options given is the relational algebra that is equivalent to the SQL shown below.

```
SELECT a
FROM x, y
WHERE x.a = y.b
```

- a. $\text{PROJECT}_a (x \text{ JOIN}_{x.a=y.b} y)$
- b. $\text{SELECT}_a (x \text{ JOIN}_{x.a=y.b} y)$
- c. $\text{SELECT}_a (x \text{ PROJECT}_{x.a=y.b} y)$
- d. $\text{PROJECT}_a (x \text{ SELECT}_{x.a=y.b} y)$
- e. None of the above

Mark: (1)

2. R S

ColA	ColB
A	1
C	4
D	6
E	6

ColC	ColD
C	2
D	1
G	6
J	7

Using the Relation Algebra Scenario, what is the cardinality of the resulting relation of:

$R \text{ JOIN}_{R.ColA = S.ColC} S$

- a. 6
- b. 8
- c. 0
- d. 2
- e. None of the above

Mark: (1)

3. Which of the following **best** describes the relationship between C and SQL?
- SQL is incompatible with programming language C
 - There exist mechanisms by which SQL statements can be embedded in C
 - There exist mechanisms by which C statements can be embedded in SQL program
 - C is compatible with the specialist database language SQL
 - SQL can be executed from within C programs by means of JDBC

Mark: (1)

4. Accessing data via embedded SQL requires which of the following.
- The cardinality of tables to be one.
 - Cursors to move from one column to the next.
 - The data dictionary to provide accurate library routines.
 - Pre-compiler support for the programming language.
 - The database to be fully normalised.

Mark: (1)

5. Given the following relation and dependencies, select the option that is the result of fully normalising the relation to BCNF.

$R(\underline{a}, b, c, d, e)$

$a \rightarrow c$

$d \rightarrow c, e$

- $R_1(\underline{a}, c)$
 $R_2(\underline{d}, e)$
 $R(\underline{a}, b, d)$
- $R_1(\underline{a}, c)$
 $R_2(\underline{d}, c, e)$
 $R(\underline{a}, b, d)$
- $R_1(\underline{d}, c, e)$
 $R(\underline{a}, b, d)$
- $R(\underline{a}, b, c, \underline{d}, e)$
- None of the above

Mark: (1)

6. To transform a relation from first normal form to second normal form we must remove which one of the following?
- a. All transitive dependencies
 - b. All inverse partial-key dependencies
 - c. All repeating groups
 - d. All partial-key dependencies
 - e. None of the above

Mark: (1)

7. There are two relations X and Y. Relation X has arity 1 and cardinality 2, relation Y has arity 3 and cardinality 4. Indicate the result of the SQL statement `SELECT COUNT(*) FROM X, Y.`
- a. 6
 - b. 4
 - c. 8
 - d. Cannot be calculated from the information given
 - e. None of the above.

Mark: (1)

8. Consider the following functional dependencies:

a, b	->	c, d
e	->	c
b	->	e, f

Given the functional dependencies shown above, what normal form is the following relation in: R(a, b, c, d, e, f)

- a. third normal form
- b. BCNF
- c. second normal form
- d. first normal form
- e. unnormalised

Mark: (1)

9. Consider the relational schema $R(\underline{A}, \underline{B}, C, D, E)$ with non-key functional dependencies $C, D \rightarrow E$ and $B \rightarrow C$.

Select the strongest statement that can be made about the schema R

- a. R is in first normal form
- b. R is in second normal form
- c. R is in BCNF normal form
- d. R is in third normal form
- e. None of the above

Mark: (1)

10. When implementing security in a DBMS, which of the following is NOT supported by the GRANT command.

- a. Providing SELECT privileges.
- b. Supporting the devolution of access control to non-DBAs.
- c. Changing passwords.
- d. Providing DELETE privileges.
- e. Removing privileges of other people.

Mark: (1)

11. A timetable database is required for a University Department. Each taught event is part of a module, each event will have exactly one member of staff associated and several individual students. Each event takes place in a single weekly time slot. Each time slot has a day of the week and a time of day associated. Staff and students can have more than one event to attend.

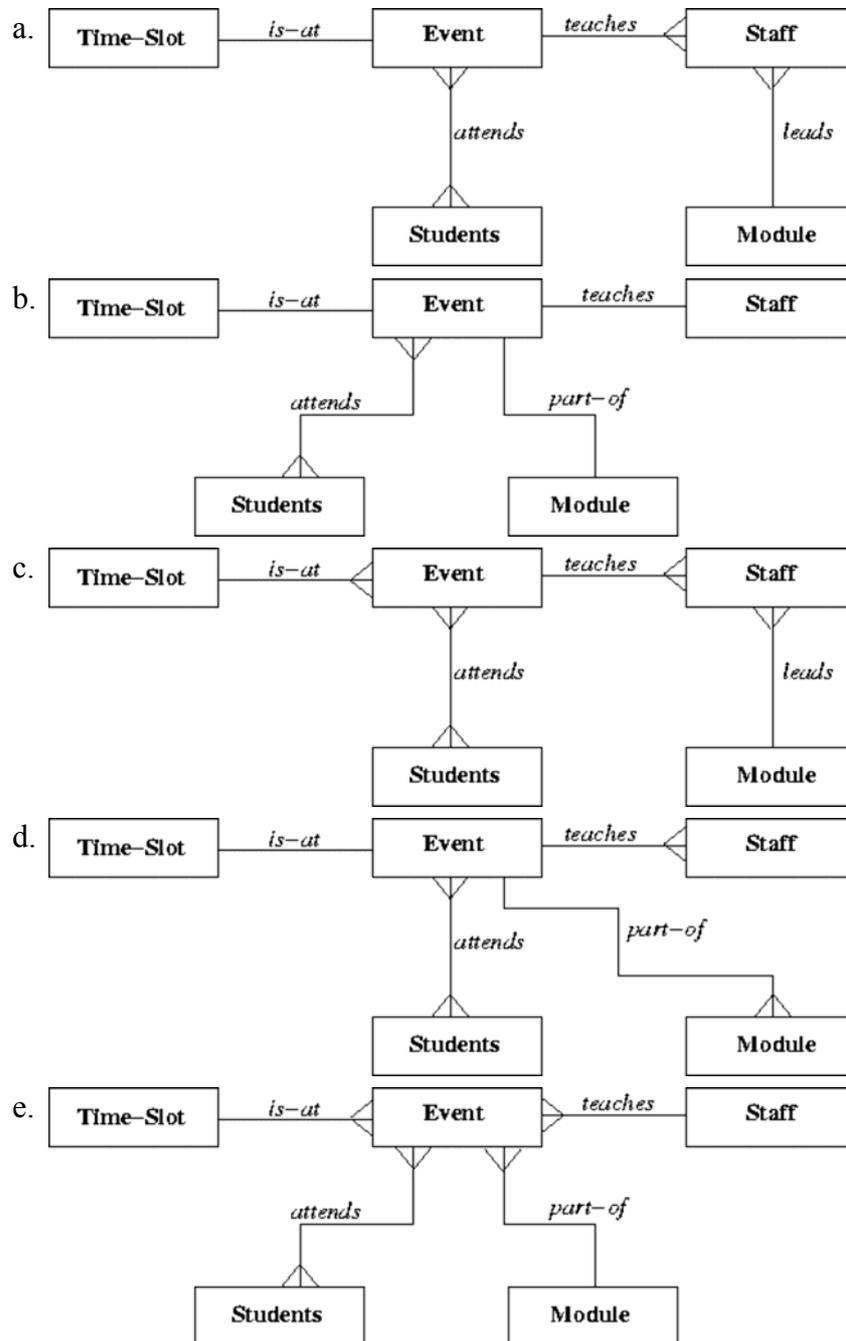
Which of the following is the **best** approach to implementing the *students that attend* relationship using a relational database system?

- a. A repeating field *attends* is included as part of the *student* table
- b. A secondary *attends* key is added to the event table
- c. A repeating field *attends* is included as part of the *event* table
- d. A table *attends* contains an event/student pair for every instance of a student attending an event
- e. None of the above

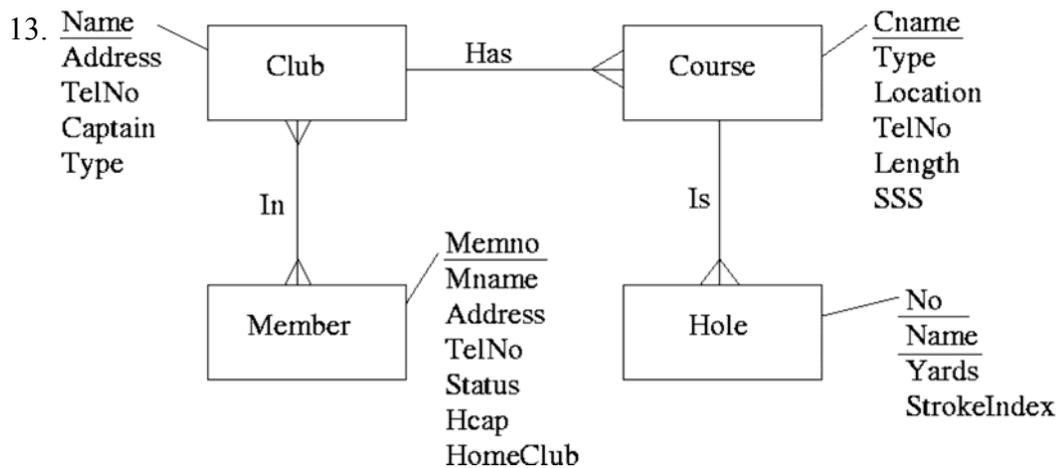
Mark: (1)

12. Continuing from the previous question

Select the most appropriate ER diagram for the above scenario:



Mark: (1)



After mapping the above ERD to a relational schema which of the following set of relations would be obtained?

- Club (Name, Address, TelNo, Captain, Type)
 Course (Cname, Type, Location, TelNo, Length, SSS, Club)
 Member (MemNo, Mname, Address, TelNo, Status, Hcap, HomeClub)
 Hole (No, Name, Yards, StrokeIndex, Cname)
 In (MemNo, Name)
- Club (Name, Address, TelNo, Captain, Type)
 Course (Cname, Type, Location, TelNo, Length, SSS)
 Member (MemNo, Mname, Address, TelNo, Status, Hcap, HomeClub)
 Hole (No, Name, Yards, StrokeIndex, Cname)
 In (MemNo, Name)
- Club (Name, Address, TelNo, Captain, Type)
 Course (Cname, Type, Location, TelNo, Length, SSS, Club)
 Member (MemNo, Mname, Address, TelNo, Status, Hcap, HomeClub)
 Hole (No, Name, Yards, StrokeIndex, Cname)
- Club (Name, Address, TelNo, Captain, Type)
 Course (Cname, Type, Location, TelNo, Length, SSS, Club)
 Member (MemNo, Mname, Address, TelNo, Status, Hcap, HomeClub)
 Hole (No, Name, Yards, StrokeIndex)
 In (MemNo, Name)
- None of the above.

Mark: (1)

14. Select the TRUE statement.

- a. For referential integrity, all primary keys should be non null.
- b. For entity integrity, all primary keys should be null or unique.
- c. For entity integrity, all foreign keys should be null.
- d. For referential integrity, all foreign keys should equal a primary key in another table.
- e. For referential integrity, each foreign key should be null or equal to a primary key in another table.

Mark: (1)

15. The use of a Data Dictionary produces many benefits. Select the benefit which is NOT due to proper use of a data dictionary.

- a. consistency in data use
- b. reduced data redundancy
- c. performance measurement
- d. improved documentation
- e. the enforcement of standards.

Mark: (1)

16. The Data Dictionary of a DBMS can be used for a variety of tasks. Which of the following is NOT ONE OF THESE?

- a. Data analysis
- b. Costing Change
- c. Improved Documentation
- d. Invoicing
- e. Enforcement of standards

Mark: (1)

17.

job		
reference	employer	salary
01	Napier	£20000
02	GCHQ	£22000
03	Napier	£24000

requirement	
job	skill
01	Unix Admin
01	Oracle Admin
02	Unix Admin
02	Number Theory

Select the term which best describes the cardinality of the relationship between the table job to the table requirement.

- many to one
- some to many
- one to many
- many to many
- one to one

Mark: (1)

18. Continuing from the previous question

Which of the following show appropriate primary keys for the tables?

- job(reference, employer, salary) requirement(job, skill)

Mark: (1)

19. Continuing from the previous question

Choose the SQL statement which will return details of the jobs of interest to a candidate with experience of Unix Administration.

- a.

```
SELECT * FROM job
WHERE skill = 'Unix Admin';
```
- b.

```
SELECT * FROM job
WHERE reference='Unix Admin'
AND skill = 'Unix Admin';
```
- c.

```
SELECT * FROM job, requirement
WHERE reference=job AND skill='Unix Admin';
```
- d.

```
SELECT * FROM job
WHERE employer = 'Napier'
OR employer = "'GCHQ';
```
- e.

```
SELECT 'Unix Admin' FROM job, requirement;
```

Mark: (1)

20. In relational database evolution, "Conceptual Design" is the stage where we map

- a. ER diagrams into relations
- b. ER diagrams into tables
- c. Specification into ER diagrams
- d. Specification into relations
- e. Specification into Marketing Ideas

Mark: (1)

21. Enhanced ER Models offer a few advantages over normal ER diagrams. Which of the following is one of these?

- a. Cataloguing
- b. Allegation
- c. Aggregation
- d. Mitigation
- e. Capitalisation

Mark: (1)

22. Which of the following is part of the ANSI/SPARC three level architecture model.

- a. conceptual
- b. coaxial
- c. contactable
- d. contextual
- e. client

Mark: (1)

23. The relationship between two entity types A and B is 1:1, and the relationship is optional at the A end. Only 50% of B entities are related to an A entity. Now consider mapping these entity types into relations. Select the best statement from the following list:

- a. A and B should be kept separate with the foreign key in the A relation.
- b. A and B should be kept separate with a foreign key in both A and B.
- c. B should be subsumed by A
- d. A should be subsumed by B
- e. A and B should be kept separate with the foreign key in the B relation.

Mark: (1)

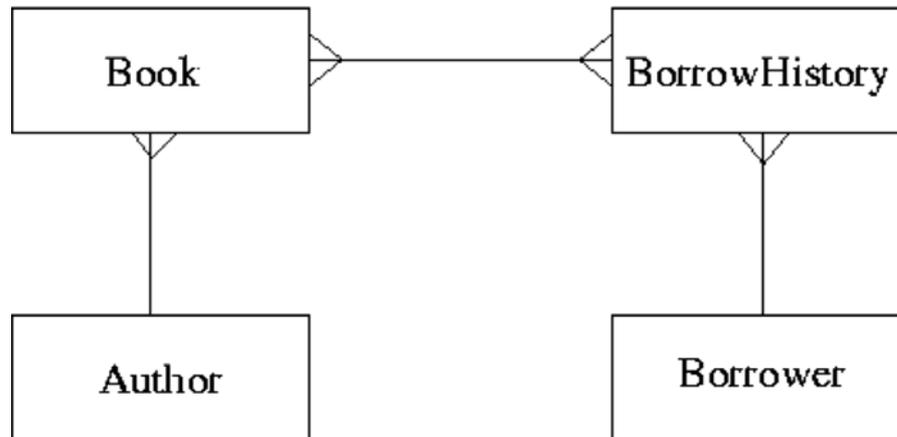
24. A primary key in a table

- a. must not contain NULL
- b. must be a number
- c. must be different from foreign keys in another table
- d. must be different from a primary key in another table
- e. can be the same value as other primary key values in the same table

Mark: (1)

25. ER Scenario

The scenario described here is that of a book library. Books in the library can be borrowed by a borrower, and a complete history of all the books a borrower has borrowed is held in the BorrowHistory entity set. All books must have an author.



The attributes of each entity set are listed below:

Author (name, country)
Book (title, publisher)
BorrowHistory (when-borrowed, when-due-back)
Borrower (name, address, date-of-birth)

ER Scenario 1 is a good example of:

- A Cliff Trap
- A Chasm Trap
- A Fan Trap
- A Tsunami
- None of the above

Mark: (1)

26. Deadlock occurs when

- the user requests a cascade abort
- the DBMS cannot decide what to do next.
- when transactions compete for the same resource.
- the state of a lock in the DBMS changes from live to dead.
- a transaction cannot decide what to do next.

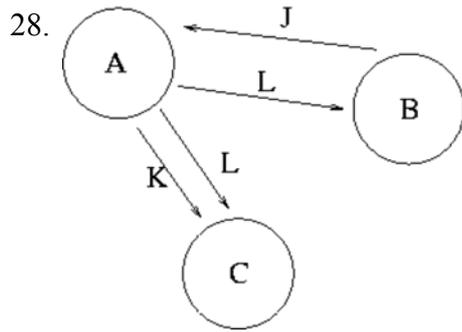
Mark: (1)

27. Select the problem which is demonstrated by the following schedule.

Time	Transaction A	Transaction B
T1	Write(X)	
T2		READ(X)
T3	Abort	

- a. Lost Update
- b. Uncommitted Dependency.
- c. Inconsolable Update
- d. Incoherency Analysis
- e. Inconsistency Analysis

Mark: (1)



The graph in the Transaction Scenario was produced from the following schedule. time

time	Transaction A	Transaction B	Transaction C
1		WRITE(J)	
2	READ(J)		
3	WRITE(K)		
4			READ(K)
5	WRITE(L)		
6		READ(L)	
7			READ(L)
8	COMMIT		
9		COMMIT	
10			COMMIT

Compare the precedence graph to the transaction schedule, and select the TRUE statement from the following:

- One or more of the loops are missing
- Errors in BOTH loops and arrow direction
- The circles should contain the attributes, and the loops labeled with the transaction names.
- The precedence graph is accurately drawn
- One or more of the arrows are pointing the wrong way

Mark: (1)

29. In transaction, cascade rollback

- a. Can occur in systems which use immediate writeback
- b. Occurs in systems which use the "waterfall" transaction management system
- c. Can occur in systems which use deferred writeback
- d. Is a result of simultaneous transaction commits.
- e. None of the above

Mark: (1)

30. Select the statement which best characterises the performance of the insert, seek and sort operations when comparing a hash index against a B+ tree index.

- a. sort and insert are similar, but seek will be slower with the hash index
- b. insert and seek are similar, but sort will be slower with the hash index
- c. seek will be similar, insert and sort will be slower
- d. seek and sort are similar, but insert will be slower with the hash index
- e. insert, seek and sort will all be similar

Mark: (1)

31. With respect to a DBMS using a hard drive as secondary storage. Select the TRUE statement.

- a. Data should be stored in contiguous blocks to maintain consistency.
- b. Data can be read from a file in "attribute-size" amounts.
- c. The DBMS manages the disk blocks.
- d. A block is the smallest unit which can be read from the disk by the operating system.
- e. The database should be backed up onto primary memory at regular intervals.

Mark: (1)

32. DBMS's usually use a combination of main memory and secondary storage to hold the database. Select the FALSE statement.
- a. Secondary storage is persistent.
 - b. Storing a database in secondary storage is a cost-effective approach.
 - c. Main memory is too expensive to be able to hold large databases.
 - d. Secondary storage is faster than main memory.
 - e. Main memory is usually erased when the power is switched off.

Mark: (1)

33. Select the true statement from the following statements comparing fine and coarse grain locking granularity.
- a. fine grained locking required less processing by the DBMS
 - b. the level of concurrent access is unaffected by grain size
 - c. the grain size is best selected using Dykras' Algorithm
 - d. coarse grained locking has less concurrency
 - e. None of the above

Mark: (1)

34. Which of the following is TRUE for two-phase locking.
- a. locks can only be acquired on primary keys.
 - b. locks can be acquired at any time.
 - c. keys are aquired in the first phase.
 - d. lock acquisition is the second phase.
 - e. None of the above.

Mark: (1)

35. The following database contains weather measurements for a number of stations around the UK.. Each station is in a region, each station records a value for rainfall in cm and sunshine in hours.

```
region(regionid, name)
station(stationid, rainfall, sunshine, region)
```

REGIONID	NAME
1	BORDERS
2	FIFE
3	LOTHIAN

STATIONID	RAINFALL	SUNSHINE	REGION
1	10	2	1
2	11	4	1
3	55	0	3
4	23	1	3
5	17	6	2
6	11	4	2
7	41	3	2

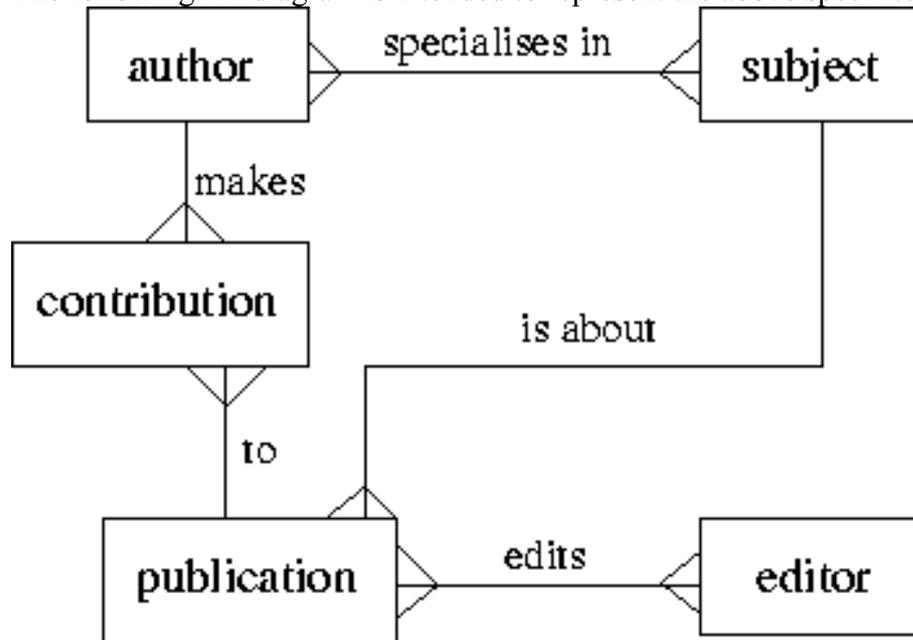
The following SQL statement is intended to return the station id of the station which recorded the lowest value for sunshine. Select the simplest SQL statement which does this.

- ```
SELECT stationid
FROM station
WHERE sunshine IN (
 SELECT MIN(sunshine)
 FROM station
);
```
- ```
SELECT stationid FROM station
WHERE sunshine = MIN;
```
- ```
SELECT stationid
FROM station AS A station AS B
WHERE A.sunshine < B.sunshine;
```
- ```
SELECT stationid, MIN(sunshine) FROM station;
```
- ```
SELECT stationid, MIN(sunshine)
FROM station
GROUP BY stationid;
```

Mark: (1)

36. A publishing company produces academic books on various subjects. Books are written by authors who specialise in one or more particular subject. The company employs a number of editors who do not have particular specialisations but who take sole responsibility for editing one or more publications. A publication covers a single subject area but may be written by one or more author - the contribution of each author is recorded as a percentage for the purposes of calculating royalties.

The following ER diagram is intended to represent the above specification:



Indicate the relation which has an incorrect cardinality shown:

- to
- makes
- specialises in
- is about
- None of the above

Mark: (1)

37. The standard language SQL contains features to perform which of the following functions:

- a. detecting redundant data
- b. specifying user passwords
- c. specifying user access rights
- d. specifying disk geometry
- e. specifying frequency of backups

Mark: (1)

38. Which of the following is TRUE about a foreign key?

- a. It can have a value which does not relate to a primary key.
- b. It can relate to multiple rows in another table.
- c. It is only used in multi-language database implementations.
- d. It can relate to columns which are not primary keys.
- e. It maintains a relationship between tables.

Mark: (1)

39. What is the main advantage of immediate update.

- a. The locking strategy used can be timestamp based.
- b. Transactions that are short are given priority over long transactions.
- c. Changes can be stored on the disk before a commit.
- d. Transactions that are long are given priority over short transactions.
- e. All transaction data is held in memory making aborts faster.

Mark: (1)

40. With references to Deferred Update, which of the following is TRUE?

- a. it is also called the UNDO/REDO algorithm.
- b. it is also called the NO-UNDO/REDO algorithm.
- c. it is also called the UNDO/NO-REDO algorithm.
- d. it is also called the NO-UNDO/NO-REDO algorithm.
- e. none of the above.

Mark: (1)

|    |    |    |    |    |    |    |    |
|----|----|----|----|----|----|----|----|
| 1  | a. | 11 | d. | 21 | c. | 31 | d. |
| 2  | d. | 12 | e. | 22 | a. | 32 | d. |
| 3  | b. | 13 | a. | 23 | a. | 33 | d. |
| 4  | d. | 14 | e. | 24 | a. | 34 | e. |
| 5  | a. | 15 | c. | 25 | e. | 35 | a. |
| 6  | d. | 16 | d. | 26 | c. | 36 | e. |
| 7  | c. | 17 | c. | 27 | b. | 37 | c. |
| 8  | d. | 18 | a. | 28 | d. | 38 | e. |
| 9  | a. | 19 | c. | 29 | a. | 39 | c. |
| 10 | c. | 20 | c. | 30 | b. | 40 | b. |