

**NAPIER UNIVERSITY**  
**SCHOOL OF COMPUTING**

**FIRST DIET (Malaysia) - SESSION 2002-03**

**DATABASE SYSTEMS**

**MODULE NO: CO22001**

**DATE: Feb 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. Relation C is a projection of relation A. Which of the following statements must be **true in all cases** where relation C is different from relation A?
- a. The cardinality of C is greater than the cardinality of A
  - b. The arity of C is less than the arity of A
  - c. The cardinality of C is less than the cardinality of A
  - d. The arity of C is greater than the arity of A
  - e. None of the above

Mark: (1)

2. Table A

Col1	Col2
A	1
B	3
C	4

Table B

Col3	Col4
A	1
C	4
D	5
E	3

Consider the tables A and B shown above and select the result of  
 $A \text{ JOIN}_{\text{col1} = \text{col3}} B$

a.

Col1	Col2	Col3	Col4
A	1	A	1
C	4	C	4
		D	5
		E	3

b.

Col1	Col2	Col3	Col4
A	1	A	1
B	3		
C	4	C	4
		D	5
		E	3

c.

Col1	Col2	Col3	Col4
A	1	A	1
C	4	C	4

d.

Col1	Col2	Col3	Col4
A	1	A	1
B	3		
C	4	C	4

e. None of the above

Mark: (1)

3. Accessing data via embedded SQL requires which of the following.

- a. The cardinality of tables to be one.
- b. The database to be fully normalised.
- c. Cursors to move from one column to the next.
- d. The data dictionary to provide accurate library routines.
- e. Pre-compiler support for the programming language.

Mark: (1)

4. A PHP script is required to return the number of rows in the table X. Two approaches are being considered:

T1            A cursor based on the SQL statement  
                      "SELECT \* FROM X" should be used.  
T2            A cursor based on the SQL statement  
                      "SELECT COUNT(\*) FROM X" should be used.

- a. Neither T1 nor T2 are feasible
- b. T2 is feasible, T1 is not feasible.
- c. T1 and T2 are feasible however T1 is better.
- d. T1 is feasible, T2 is not feasible.
- e. T1 and T2 are feasible however T2 is better.

Mark: (1)

5. Each of the following is an argument which might be used to support the use of relations which are not fully normalised. Select the **weakest** argument.

- a. A fully normalised database may have too many tables
- b. Full normalisation may compromise existing applications/systems
- c. Full normalisation may make some queries too complicated
- d. A fully normalised database may result in tables which are too large
- e. A fully normalised database may perform too slowly

Mark: (1)

6. 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. Cannot be calculated from the information given
- b. 4
- c. 6
- d. 8
- e. None of the above.

Mark: (1)

7. Consider the following functional dependencies

a,b  $\Rightarrow$  c,d    e,g,h  $\Rightarrow$  f,j  
a,c  $\Rightarrow$  b,d    p,q  $\Rightarrow$  r,s  
e,f,g  $\Rightarrow$  h,i    s  $\Rightarrow$  t  
f,g  $\Rightarrow$  j    q  $\Rightarrow$  u  
g,h  $\Rightarrow$  i

Which of the following relational schemas might be the result of normalising R(a,b,c,d)?

- a. The schema R(a,b,c,d)
- b. The schema R1(a,b) R2(a,c) R3(a,d)
- c. The schema R1(a,b) R2(b,c) R3(c,d)
- d. The schema R1(a,b,c) R2(a,b,d)
- e. The schema R1(a,b) R2(a,c) R3(b,d)

Mark: (1)

8. For a relation to be in third normal form, which of the following is NOT a requirement?
- Every non-key attribute is fully functionally dependent on the whole key
  - Every determinant is a candidate key
  - It must be in second normal form
  - There must be no transitive functional dependencies
  - more than one of the above.

Mark: (1)

9. An athletics meeting involves several competitors who participate in a number of events. The database is intended to record who is to take part in which event and to record the outcome of each event. As results become available the winner attribute will be updated with the cid of the appropriate competitor.

Competitor(cid, name, nationality)  
 Event(eid, description, winner)  
 Competes(cid, eid)

Competitor			Event			Competes	
cid	name	nationality	eid	description	winner	cid	eid
01	Pat	British	01	running		01	01
02	Hilary	British	02	jumping		02	01
03	Sven	Swedish	03	throwing		03	02
04	Pierre	French				04	02
						04	03

Select the true statement

- The Event table has a composite key
- Competes is the Cartesian product of Competitor and Event
- At least one of the tables is **NOT** in third normal form
- There is no means to represent a tie in a particular event.
- There is no means to represent a competitor taking part in more than one event

Mark: (1)

10. Continuing from the previous question

The actual database is to contain several million competitors and nearly one hundred thousand events. The database must be indexed for producing fast results for two particular queries:

- (X) A list of the names and nationalities of the competitors for a given event where the event description is given.
- (Y) A list of event descriptions for a given competitor where the cid is given.

Which of the following is not required?

- a. An index on cid in Competes
- b. An index on eid in Event
- c. An index on name in Competitors
- d. An index on description in Events
- e. An index on eid in Competes

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

Mark: (1)

12. An ER diagram has two entity sets A and B exists which are linked by a relationship 1:N. The A side of the relationship is optional. When mapping this into relations...
- a. the foreign key is placed in relation A.
  - b. the foreign key is placed in relation B.
  - c. you should subsume A into B.
  - d. the foreign key is placed in both relation A and relation B
  - e. you have to use create another relation first.

Mark: (1)

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

Mark: (1)

14. Which of the following is the highest normal form by which the Appointment relation can be classified?

Appointment (patient, consultant, hospital, address, date, time)

Given

patient, consultant -> hospital, address, date, time  
hospital -> address

- a. 1NF
- b. Unnormalised
- c. 2NF
- d. BCNF
- e. 3NF

Mark: (1)



15. The Data Dictionary provides which of the following features.

- a. Support for backup procedures.
- b. Transaction deadlock detection.
- c. Support for the SQL interface to the database.
- d. Costings for future database changes.
- e. None of the above.

Mark: (1)

16. Which of the following is not usually part of the responsibilities of a database administrator?

- a. Approving structural changes to the database
- b. Ensuring that an adequate back-up regime is in place
- c. Designing data entry screens
- d. Monitoring the performance of the system.
- e. Issuing accounts to users

Mark: (1)

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

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

Mark: (1)

18.

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.

- a. many to many
- b. one to one
- c. many to one
- d. one to many
- e. some to many

Mark: (1)

19. Continuing from the previous question

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

- a. job(reference, employer, salary) requirement(job, skill)
- b. job(reference, employer, salary) requirement(job, skill)
- c. job(reference, employer, salary) requirement(job, skill)
- d. job(reference, employer, salary) requirement(job, skill)
- e. job(reference, employer, salary) requirement(job, skill)

Mark: (1)

20. 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 reference='Unix Admin'
AND skill = 'Unix Admin';
```
- b. 

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

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

```
SELECT 'Unix Admin' FROM job, requirement;
```
- e. 

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

Mark: (1)

21. In Enhanced ER diagrams, a subclass

- a. Is part of Chun's notation.
- b. may contain many superclasses.
- c. may contain only one superclass.
- d. is contained in one superclass.
- e. can only exist in Chen's notation.

Mark: (1)

22. In the ANSI/SPARC three level database model, the external view is best described by which one of the following options:

- a. It is the link between users and the storage structures.
- b. It is not part of the model.
- c. It is the place where the users interface to the DBMS.
- d. It is dependent on the underlying DBMS product used (e.g. Oracle, DBASE).
- e. It is the place where the storage structures link to the database.

Mark: (1)

23. The external view of the ANSI-SPARC architecture chiefly concerns:

- a. the way individual users see the data
- b. the formal description of the data
- c. the interface to other applications
- d. the data that users outside the company are permitted to view
- e. the way the data is actually stored

Mark: (1)


24. When following the Database Analysis Life Cycle

- a. You test the database before loading in the user data
- b. User training is performed as part of evolution.
- c. Evaluation is directly after Loading.
- d. You first design the database and then write the specification
- e. Operation follows directly after Implementation.

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

Mark: (1)

26. Which one of the following requires some hardware assistance to be implemented in an efficient manner?

- a. Differential files
- b. Shadow-paging
- c. Log-files with deferred updates
- d. Log-files with immediate updates
- e. None of the above

Mark: (1)

27. Which one of the following algorithms is best suited for long-lived transactions with relatively few roll-backs?

- a. Differential files
- b. Shadow-paging
- c. Log-files with deferred updates
- d. Log-files with immediate updates
- e. None of the above

Mark: (1)

28. When a transaction aborts

- a. all users must be notified
- b. all changes it has made are immediately available to other transactions
- c. it can abort transactions which have already committed
- d. the modifications of all transactions currently running are also aborted
- e. it releases all of its locks

Mark: (1)

29. Deadlock occurs when

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

Mark: (1)

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

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

Mark: (1)

31. Which of the following is TRUE when considering hash tables.

- a. Hash collisions cannot be handled.
- b. The hash function only operates on the index field.
- c. A good hash function is one which results in many hash collisions.
- d. DBMS systems can only hash on numbers.
- e. Hash tables speed up sequential record scanning.

Mark: (1)

32. Consider a table where records are stored in primary key order. Select the statement that best describes the cost of the operations: insert, delete and seek. Deleted records are NOT "Flagged".
- a. Insert, delete and seek are all expensive.
  - b. Insert, delete and seek are all cheap.
  - c. Insert is expensive, delete and seek are cheap.
  - d. Insert and delete are expensive, seek is cheap.
  - e. Insert and delete are cheap, seek is expensive.

Mark: (1)

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

Mark: (1)

34. Locking was introduced into databases so that
- a. Consistency can be enforced.
  - b. All simultaneous transactions are prevented.
  - c. Passwords can be provided to maintain security
  - d. Keys can be provided to maintain security.
  - e. Reading and writing is possible.

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

Select the SQL statement which will list the station identity and rainfall for any station with more rainfall than station 6.

- ```
SELECT stationid, rainfall
FROM station
WHERE rainfall >
(SELECT stationid, rainfall
FROM station
WHERE stationid=6);
```
- ```
SELECT a.stationid, a.rainfall
FROM station a, station b
WHERE a.rainfall > b.rainfall
AND b.stationid = 6;
```
- ```
SELECT a.stationid, a.rainfall
FROM station a, station b
WHERE a.rainfall < b.rainfall
AND b.stationid = 6;
```
- ```
SELECT stationid, rainfall
FROM station
WHERE stationid=6;
```
- ```
SELECT b.stationid, a.rainfall
FROM station a, station b
WHERE a.rainfall < b.rainfall
AND a.stationid = 6;
```

Mark: (1)

36. Continuing from the previous question

Select the SQL statement which results in a list of the average rainfall for each region:

- a. 

```
SELECT name, AVG(rainfall)
FROM   region, station
WHERE  regionid = station.region
GROUP BY stationid;
```
- b. 

```
SELECT AVG(name), rainfall
FROM   region, station
WHERE  regionid = station.region
GROUP BY rainfall;
```
- c. 

```
SELECT AVG(name) ,AVG(rainfall)
FROM   region, station
WHERE  regionid = station.region
GROUP BY rainfall;
```
- d. 

```
SELECT AVG(rainfall)
FROM   station
GROUP BY stationid;
```
- e. 

```
SELECT name, AVG(rainfall)
FROM   region, station
WHERE  regionid = station.region
GROUP BY name;
```

Mark: (1)

37. Select the TRUE statement related to the following SQL:

```
SELECT X.a, X.b
FROM foo X, foo Y
WHERE X.b = Y.b AND Y.a='bar';
```

- a. a and b are both tables.
- b. this is an example of a self-join.
- c. the result MUST include row with 'bar' in it.
- d. the statement is NOT valid SQL.
- e. foo is an alias for X.

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

40. Select the situation which cannot result in a ROLLBACK.

- a. The machine on which the DBMS was executing failed suddenly.
- b. A transaction is aborted.
- c. A cascade abort is required.
- d. The DBMS application was corrupted by a power spike.
- e. A committed transaction needs to be undone.

Mark: (1)

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