

NAPIER UNIVERSITY
SCHOOL OF COMPUTING

FIRST DIET (UK) - SESSION 2002-03

DATABASE SYSTEMS

MODULE NO: CO22001

DATE: Jan 2003

EXAM TIME: 2 HOURS

START TIME: HOURS

FINISH TIME: HOURS

EXAMINERS:
G. RUSSELL
U. PRISS

QUESTION PAPER DATA

Number of pages – TWENTY THREE

Number of questions – FORTY

INSTRUCTIONS TO CANDIDATES

Select ONE from (a) to (e)

1.

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

A list of jobs together with the required skills is needed. It is important that jobs such as 03, which has no skills specified, are included. Which of the following operations is most appropriate?

- INNER JOIN
- UNION
- LEFT or RIGHT OUTER JOIN
- INTERSECTION
- CARTESIAN PRODUCT

Mark: (1)

2. Continuing from the previous question

The table requirement was created using the following SQL statement:

```
CREATE TABLE requirement (
    job INTEGER,
    skill VARCHAR(50),
    FOREIGN KEY job REFERENCES job(reference));
```

Select the true statement concerning the requirements table.

- referential integrity checks will prevent invalid values for job
- the value 'Windows 2000 Admin' is NOT permitted for the skill attribute
- it cannot store a field with a NULL value for job
- it must be created BEFORE the job table
- referential integrity checks will prevent identical rows being inserted

Mark: (1)

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

Mark: (1)

4. Select the TRUE statement.
- a. SQL is embedded within a C++ program to increase the speed of the C++ program.
 - b. SQL is embedded within C++ to handle sequential file processing.
 - c. SQL embedded in C++ provides facilities to extract data from a database.
 - d. SQL cursors indicate the next line of code to be executed in a C++ program.
 - e. SQL is embedded to make C++ programs more efficient.

Mark: (1)

5. Given the following relation and dependencies, state which normal form the relation is in.

$R(p, q, r, s, t)$
 $p, q \rightarrow r, s, t$
 $r, s \rightarrow p, q, t$
 $t \rightarrow s$

- a. Third normal form
- b. First normal form
- c. BCNF
- d. Second normal form
- e. Unnormalised

Mark: (1)

6. A golf club proposes to hold a database about members, instead of the current paper-based card system. (Please note that understanding of golf terms and/or any particular field is not assumed or indeed necessary). The current membership cards hold the following fields:

Member Details:

Name, DOB, Category, Handicap, BufferValue, Increment, Decrement, Home Club, Yardage, SSS

The following functional dependencies are identified:

FD1: Name, DOB => Handicap, HomeClub
FD2: DOB => Category
(i.e. Junior, Ordinary, Senior or Veteran etc.)
FD3: Handicap => BufferValue, Increment, Decrement
(The SGU Handicap system)
FD4: Yardage => SSS
(The Standard Scratch Score - an indication of the difficulty of the course based on its total length in yards.)

From this a third normal form of the relations has been produced which involves the following relations:

Member	concerns each individual member
Categories	concerns each type of membership
Handicap	concerns each class of golfer
Yardage	concerns each class of golf course

What type of functional dependency describes FD4?

- A whole-key dependency
- A P.K.D.
- A virtual dependency
- A transitive dependency
- None of the above

Mark: (1)

7. Consider the following functional dependencies:

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

Given the same functional dependencies as shown above, which option shows the relations normalised to third normal form of: R(a, b, c, d, e, f)

- a. R (a, b, c, d, e, f)
R (e, c)
R (b, e, f)
- b. R (a, b, c, d)
R (c, e)
R (e, f, b)
- c. R (a, b, c, d)
R (c, e)
R (b, e, f)
- d. R (a, b, d)
R (e, c)
R (b, e, f)
- e. R (a, b, c, d, e, f)

Mark: (1)

8. A lack of normalisation can lead to which one of the following problems :

- a. Deadlock
- b. Lost Updates
- c. Insertion problems
- d. Deferred updates
- e. Deletion of data

Mark: (1)

9. Assume the relation $R(\underline{A}, \underline{B}, C, D, E)$ is in at least 3NF. Which of the following functional dependencies must be FALSE?

- a. $A, B \rightarrow C$
- b. $A, C \rightarrow E$
- c. $A, B \rightarrow D$
- d. $C, D \rightarrow E$
- e. None of the above

Mark: (1)

10. If a system can enforce referential integrity, then this ensures that

- a. a record is always referred to from another record
- b. a foreign key attribute in a record always refers to another record which does not contain nulls
- c. a foreign key attribute in a record always refers to another record which contains nulls
- d. a non-null foreign key attribute always refers to another record
- e. a record can never contain a null value for a foreign key attribute.

Mark: (1)

11. When producing relations for the diagram shown, where almost all PERSON entities relate to a MAIN_JOB:



Which of the following is best?

- a. it is best to keep PERSON and MAIN_JOB as separate relations.
- b. PERSON should be subsumed into MAIN_JOB.
- c. MAIN_JOB should be subsumed into PERSON.
- d. NULL attributes will be produced no matter what you do.
- e. None of the above.

Mark: (1)

12. Two entity types A and B are related by a 1:1 relationship which is optional at both ends. In the process of implementing A and B as relations you would
- Combine A and B together into a single relationship.
 - Keep them separate and put a foreign key in one of A or B.
 - Use a primary key which is a composition of the primary keys of A and B.
 - Combine A and B together into a single relation.
 - Keep them separate and put a foreign key in both A and B.

Mark: (1)

13. When mapping ER models into relations, which of the following is NOT true?
- Each m-n relationship is mapped into a new relation.
 - Each attribute is mapped into a column of the corresponding relation.
 - Each 1-m relationship is mapped into a new relation.
 - Each individual entity is mapped into a row of the corresponding relation.
 - Each entity type is mapped into a relation.

Mark: (1)

14. When mapping from an ER diagram into relations, some relationships may have optionality. Select the TRUE statement.
- If the relationship is 1:n with one end optional, you must put the foreign key in the many end.
 - If the relationship is 1:n with one end optional, you must put the foreign key in the end which is not optional.
 - If the relationship is 1:n with one end optional, you must put the foreign key in the optional end
 - If the relationship is 1:1 with one end optional, you must subsume.
 - If there is both optionality and a cardinality greater than one, then you must use a composite foreign key.

Mark: (1)

15. A back-up and recovery regime should protect an organisation against:

- a. incorrect data
- b. inconsistent data
- c. data validation
- d. data corruption
- e. insecure data

Mark: (1)

16. Select the activity that is most clearly part of the Database Administrator's responsibility.

- a. Data entry.
- b. Staff management.
- c. Data modification.
- d. Data deletion.
- e. Performance tuning.

Mark: (1)

17. Referential Integrity is where

- a. references in the database are identical to references in an ER diagram.
- b. foreign keys always refer to a candidate key or are NULL.
- c. primary keys in one table always refer to primary keys in another table.
- d. foreign keys in one table always refer to foreign keys in another table.
- e. foreign keys cannot be changed by queries.

Mark: (1)

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

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

Mark: (1)

19. Within EER diagram techniques, which of the following could be the result of Generalising

```
superclass - card(cardnumber, issuer, cardholder)
subclass   - visa(expiryDate, creditLimit)
subclass   - switch(issueDate, colour)
```

- a. superclass - card(cardnumber, issuer, cardholder)
subclass - visa(expiryDate, creditLimit)
subclass - switch(issueDate, colour)
- b. card(cardnumber, issuer, cardholder, expiryDate, creditLimit, issueDate, colour)
- c. superclass - card(cardnumber, issuer, cardholder)
subclass - visa(expiryDate, creditLimit)
subclass - switch(expiryDate, issueDate, colour)
- d. subclass - card(cardnumber, issuer, cardholder)
subclass - visa(expiryDate, creditLimit)
subclass - switch(issueDate, colour)
- e. subclass - card(cardnumber, issuer, cardholder)
superclass - visa(expiryDate, creditLimit)
superclass - switch(issueDate, colour)

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

22. 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.

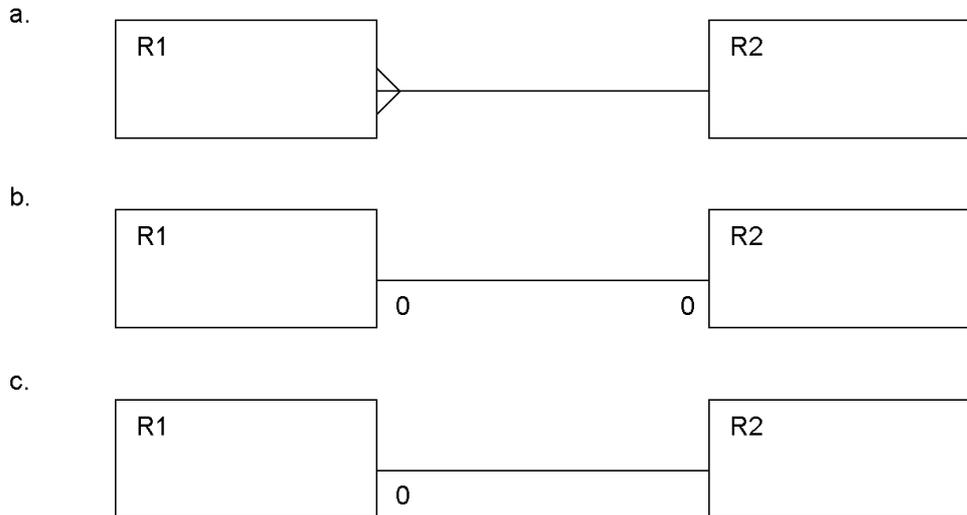
Each of the weekly time slots is exactly one hour long, however we wish to represent the fact that some events take more than one hour. Which of the following does **not** represent a possible solution.

- a. Each event has an attribute "start" which refers to Time-Slots and "duration" which gives the length of the event in minutes
- b. A one-to-many relation between Events and Time-Slots is established
- c. Each event has an attribute "start" which refers to Time-Slots and "duration" which gives the number of slots spanned
- d. Each event has two attributes "first" and "last" each of which refer to Time-Slots
- e. A many-to-many relation between Events and Time-Slots is established

Mark: (1)

23. Given the following relation select which of the ER diagrams could describe the relation.

R1 (a, b, c, d) d is a foreign key
 R2 (d, x, y)



- a. Diagram c.
- b. Diagram b.
- c. Diagram a.
- d. two of the above
- e. All three of a,b, and c.

Mark: (1)

24. 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 place where the users interface to the DBMS.
- b. It is the link between users and the storage structures.
- c. It is not part of the model.
- 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)

25. In ER Modelling, "Chasm traps" can occur when entities are related via a relationship with

- a. Partial Differentiation
- b. Partial Participation
- c. No optionality
- d. Partial Overhangs
- e. No Partial Precipitation

Mark: (1)

26. In the acronym ACID, used to describe the important properties of transactions, what does the letter A stand for?

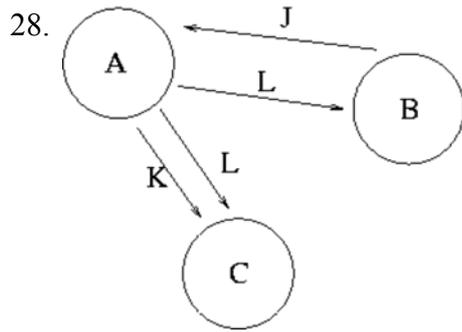
- a. After
- b. Action
- c. Also
- d. Atomic
- e. None of the above

Mark: (1)

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

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

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:

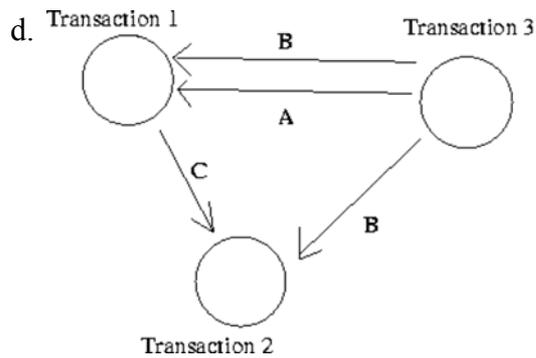
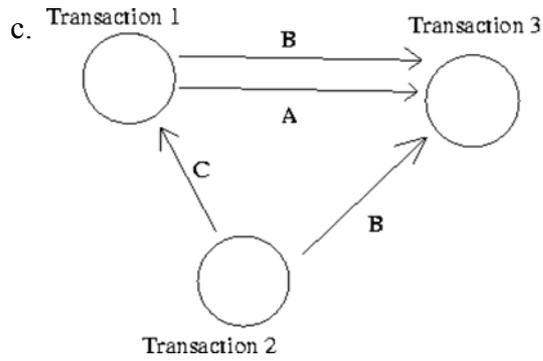
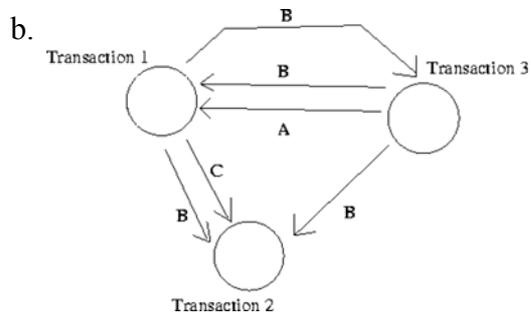
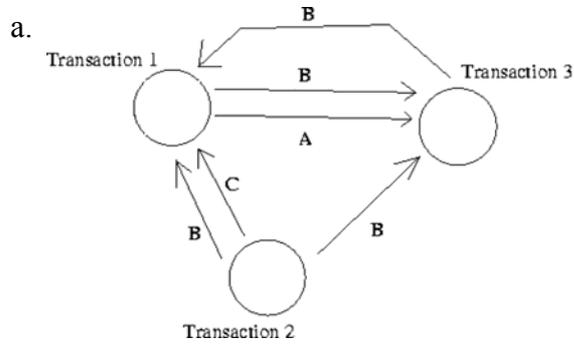
- One or more of the arrows are pointing the wrong way
- One or more of the loops are missing
- Errors in BOTH loops and arrow direction
- The circles should contain the attributes, and the loops labeled with the transaction names.
- The precedence graph is accurately drawn

Mark: (1)

time	TRANSACTION 1	TRANSACTION 2	TRANSACTION 3
t1	read(A)		
t2		read(B)	read(B)
t3	read(B)		
t4		write(C)	
t5			write(B)
t6			read(A)
t7			write(A)
t8	write(A)		
t9	write(C)		

29. Given the transaction schedule shown on the previous page:

Which of the following precedence graphs depicts the above scenario?



e. None of the above.

Mark: (1)

30. The option of dropping a secondary index to a table is being considered. Which of the following is most likely to be a consequence of dropping the index?

- a. Certain foreign key relations may not be maintained
- b. Certain insertions may be faster
- c. Certain updates may be slower
- d. More disk space may be required
- e. Certain seek operations may be faster.

Mark: (1)

31. Indexes speed up data access. Select the TRUE statement:

- a. Columns which are frequently modified are good candidates for indexing.
- b. Secondary indexes must have unique keys.
- c. Primary indexes can have duplicate keys.
- d. Primary indexes may have null values.
- e. None of the above.

Mark: (1)

32. Using SERIAL data organisation -

- a. new records can only overwrite old records
- b. new records are written after the last record
- c. new records are written in primary key order
- d. new records are written randomly
- e. new records are written in candidate key order

Mark: (1)

33. Which of the following is TRUE for two-phase locking.

- a. lock acquisition is the second phase.
- b. locks can be acquired at any time.
- c. locks can only be acquired on primary keys.
- d. keys are acquired in the first phase.
- e. None of the above.

Mark: (1)

34. Which one of the following techniques is sometimes used to solve integrity problems in a concurrent transaction scenario?

- a. First-come first-served.
- b. Greedy algorithms.
- c. Strassens's algorithm.
- d. Two-phase locking.
- e. First-fit.

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 results in a list of the average rainfall for each region:

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

Mark: (1)

36. Continuing from the previous question

The following SQL is intended to return the details of just the weather station in FIFE.

```
SELECT *  
FROM region, station  
WHERE name LIKE 'FIFE';
```

Select the TRUE statement

- a. The result will include unwanted stations.
- b. The result will be as intended.
- c. Zero row will be returned.
- d. The result will include unwanted rows.
- e. The LIKE clause must include a wild card.

Mark: (1)

37. 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 LEFT OUTER JOIN_{col1 = col3} B

a.

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

b.

Col1	Col2	Col3	Col4
A	1	A	1
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)

38. **Films Database**

Consider the following database:

MOVIE(id,title,yr)

ACTOR(id,name)

CASTING(movieid,actorid)

Identify the SQL command which will return the titles of all 1959 Marilyn Monroe films.

a. The following SQL...

```
SELECT title FROM movie,actor
WHERE name = 'Marilyn Monroe'
AND yr = 1959
;
```

b. The following SQL...

```
SELECT title FROM movie,casting,actor
WHERE movieid = movie.id
AND actor.id = actorid
AND movie.yr = casting.yr
AND name = 'Marilyn Monroe'
AND yr = 1959
;
```

c. The following SQL...

```
SELECT title FROM movie,casting,actor
WHERE movieid = movie.id
AND name = 'Marilyn Monroe'
;
```

d. The following SQL...

```
SELECT title FROM movie,casting,actor
WHERE movieid = movie.id
AND actor.id = actorid
AND name = 'Marilyn Monroe'
AND yr = 1959
;
```

e. None of the above

Mark: (1)

39. During Rollforward, which of the following is not true

- a. Committed transactions are reapplied in the same order as they appear in the log.
- b. Transactions which are in the log but not yet committed are restarted where they left off.
- c. Committed transactions in the log are reapplied.
- d. Only once the rollforward has been completed can new transactions be performed.
- e. Transactions which are reapplied are reapplied to the image as it was since the last checkpoint.

Mark: (1)

40. A database can be left in an inconsistent state due to

- a. A transaction fails and its changes are applied to the database.
- b. Roll-forward after a failure.
- c. Deadlock
- d. Inaccurate data is entered into the database.
- e. Transactions being aborted.

Mark: (1)

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