

NAPIER UNIVERSITY
SCHOOL OF COMPUTING

SECOND DIET (Malaysia) - 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 FOUR

Number of questions – FORTY

INSTRUCTIONS TO CANDIDATES

Select ONE from (a) to (e)

1. 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
C	4	C	4

b.

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

c.

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

d.

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

e. None of the above

Mark: (1)

2. Continuing from the previous question

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

c.

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

d.

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

e. None of the above

Mark: (1)

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

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 is embedded to make C++ programs more efficient.
- d. SQL embedded in C++ provides facilities to extract data from a database.
- e. SQL cursors indicate the next line of code to be executed in a C++ program.

Mark: (1)

5. 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

Select the appropriate description of the Member relation:

- a. Member: Name, DOB, Handicap, HomeClub, Yardage
- b. Member: Name, DOB, Handicap, HomeClub
- c. Member: Name, DOB, Handicap, HomeClub
- d. Member: Name, DOB, Handicap, HomeClub, Yardage
- e. Member: Name, DOB, Handicap, HomeClub, Yardage

Mark: (1)

6. 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
- There is no means to represent a tie in a particular event.
- At least one of the tables is **NOT** in third normal form
- There is no means to represent a competitor taking part in more than one event

Mark: (1)

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

Mark: (1)

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

Mark: (1)

9. Select the TRUE statement which would indicate data in the database as "redundant".

- a. Secondary keys are not unique
- b. A VIEW has the same data as a TABLE.
- c. It can be derived from other data in the database.
- d. The data has not yet been COMMITTED to the database
- e. It is unique in the database

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

Mark: (1)

11. 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

- a. Keep them separate and put a foreign key in one of A or B.
- b. Combine A and B together into a single relationship.
- c. Use a primary key which is a composition of the primary keys of A and B.
- d. Combine A and B together into a single relation.
- e. Keep them separate and put a foreign key in both A and B.

Mark: (1)

12. A primary key in a table

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

Mark: (1)

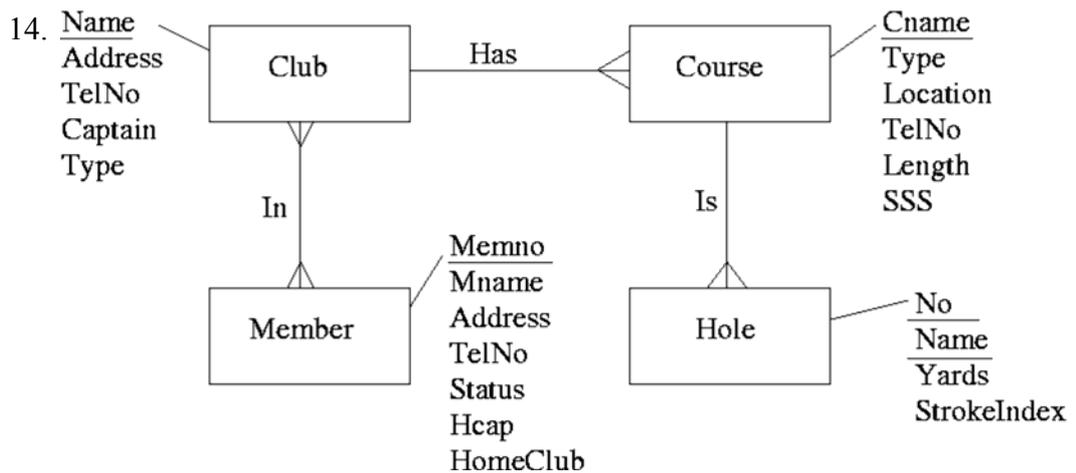
13.

Departments		Employees		WorkFor	
DeptNo	Depname	Empno	Empname	Empno	Depno
1	Computng	1	Gordon	1	1
2	Electrical	1	Ken	3	2
3	Geography	1	Brian	4	1
4	History	1	Colin	3	3
5	Business	1	George	1	2
				2	5

Using SQL Scenario 1, what is the CARDINALITY of the table "WorkFor"?

- a. 3
- b. 12
- c. 6
- d. 2
- e. none of the above.

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)
 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, 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)
- None of the above.

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

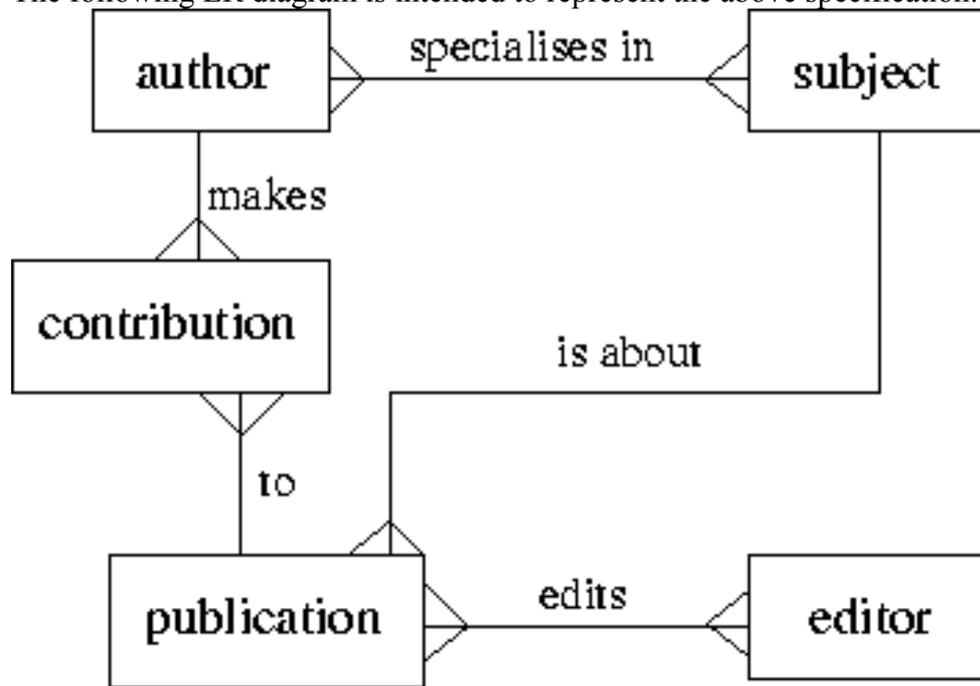
17. The relationship between two entity types A and B is 1:1, and the relationship is optional at the A end. Only 50% of B entities are related to an A entity. Now consider mapping these entity types into relations. Select the best statement from the following list:

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

Mark: (1)

18. 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:

- a. specialises in
- b. is about
- c. to
- d. makes
- e. None of the above

Mark: (1)

19. Continuing from the previous question

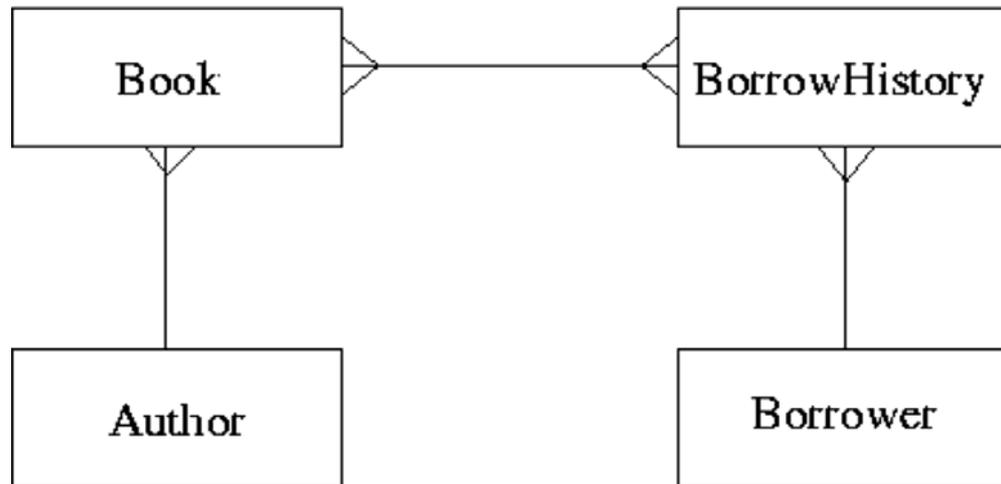
The specification is to be changed so that an author can develop a publication covering more than one subject area and that the schema must be able to store the percentage of the pages concerned with each of the subjects. Select an appropriate change to the ER diagram:

- a. author-publication becomes many to many
- b. publication-subject becomes many to many
- c. author-subject becomes many to many
- d. more than one of the above
- e. none of the above

Mark: (1)

20. 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)

Considering the information in ER Scenario, which of the following statements is TRUE?

- A book can only have a single author
- A borrower must borrow multiple books.
- Many-to-many relationships should never appear in an ER diagram
- A borrower can only borrow the same book once
- A book can only be borrowed once

Mark: (1)

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

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

Mark: (1)

22. At the Physical design stage, select the TRUE statement.

- a. ER diagrams are mapped into relationships
- b. Relationships are mapped into tables
- c. Indices are identified and implemented for tables
- d. ER diagrams are mapped into relations
- e. ER diagrams are mapped into tables.

Mark: (1)

23. When following the Database Analysis Life Cycle

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

Mark: (1)

24. Referential Integrity is where

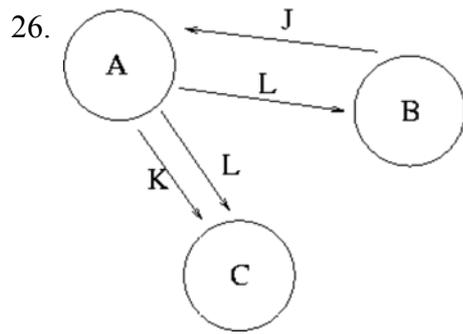
- a. references in the database are identical to references in an ER diagram.
- b. primary keys in one table always refer to primary keys in another table.
- c. foreign keys cannot be changed by queries.
- d. foreign keys always refer to a candidate key or are NULL.
- e. foreign keys in one table always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

Mark: (1)

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

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

Mark: (1)



From Transaction Scenario 1, given the precedence graph, which of the following is TRUE?

- a. Nothing as there is not enough information in the graph
- b. That the transaction schedule is unserialisable
- c. Nothing, as precedence graphs do not work for more than two transactions
- d. That the transaction schedule could be both serialisable and unserialisable
- e. That the transaction schedule is serialisable

Mark: (1)

27. Continuing from the previous question

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:

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

Mark: (1)

28. In an DBMS without concurrency control, what consistency problem does the following transaction schedule depict?

Time Transaction A Transaction B

t1	read R	
t2		read R
t3	write R	
t4		write R

- a. Lost Update
- b. Deadlock
- c. Inconsistent Analysis
- d. Uncommitted Dependency
- e. Dirty Read

Mark: (1)

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

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

Mark: (1)

30. Indexes speed up data access. Which of the following are TRUE?

- a. Primary indexes can have duplicate keys.
- b. An attribute which only has a limited number of possible values will still have access performance improved using an index.
- c. Secondary indexes must have unique keys.
- d. Columns which are frequently modified are good candidates for indexing.
- e. None of the above.

Mark: (1)

31. Consider each of the following properties of an index which uses ordered trees. Select the feature which characterises POOR performance.

- a. The tree contains long branches
- b. The branching factor is high
- c. The maximum depth of the index tree is small
- d. The attributes being indexed are short
- e. The index tree is balanced

Mark: (1)

32. A given relation is known to be in third normal form. Select the statement which can be inferred from this:

- a. The relation is not in fourth normal form.
- b. Each non-key attribute is determined by the primary key
- c. All attributes contribute to the primary key
- d. Each non-key attribute determines the primary key
- e. Every determinant is a candidate key

Mark: (1)

33. Select the problem that can occur due to introducing locks in a concurrent transaction scenario.

- a. Loss of integrity
- b. Transaction rollover
- c. Performance degradation
- d. Hash key clash
- e. None of the above.

Mark: (1)

34. With respect to Two-Phase Locking, select the TRUE statement.

- a. Before accessing an item a lock must first be acquired
- b. If a needed lock cannot be acquired then the transactions are deadlocked
- c. Locks can be acquired at any point in a transaction
- d. Locks are only required when accessing keys
- e. None of the above

Mark: (1)

35. Which of the following **best** describes the relation between ISO SQL and ORACLE's SQL*PLUS?

- a. ORACLE SQL*PLUS may be installed on a wider range of platforms
- b. ORACLE SQL*PLUS is faster than ISO SQL
- c. ORACLE SQL*PLUS is a commercial product, ISO SQL is freeware.
- d. ORACLE SQL*PLUS is an attempt to implement a superset of ISO SQL
- e. ORACLE SQL*PLUS is the industry standard definition of ISO SQL

Mark: (1)

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

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

Mark: (1)

37. 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 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

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

Mark: (1)

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

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

Mark: (1)

39. With Immediate Update, writing to an attribute results in the DBMS

- a. Releasing concomitant locks.
- b. Writing the new value to the log.
- c. Immediately performing a COMMIT.
- d. Writing the old and new attribute value to the log.
- e. Immediately writing the change to the database stored on disk.

Mark: (1)

40. What is the main advantage of immediate update.

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

Mark: (1)

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