

NAPIER UNIVERSITY
SCHOOL OF COMPUTING

FIRST DIET - SESSION 1998-99

DATABASE SYSTEMS 2

MODULE NO: CS22004

DATE: 1 JUNE 1999

EXAM TIME: 1½ HOURS

START TIME: 0930 HOURS

FINISH TIME: 1100 HOURS

EXAMINERS:

J. MURRAAY

G. RUSSELL

QUESTION PAPER DATA

Number of pages – FOURTEEN

Number of questions – FORTY

INSTRUCTIONS TO CANDIDATES

Select ONE from (a) to (e)

SECTION A

1. A given relation is known to be in third normal form. Select the statement which can be inferred from this:
 - (a) All attributes contribute to the primary key
 - (b) Each non-key attribute determines the primary key
 - (c) Each non-key attribute is determined by the primary key
 - (d) Every determinant is a candidate key
 - (e) The relation is not in fourth normal form.

(1)

2. Which of the following is generally a benefit of normalisation?
 - (a) Performance is improved
 - (b) Insertion anomalies are avoided
 - (c) Selection anomalies are avoided
 - (d) Number of tables is reduced
 - (e) None of the above.

(1)

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

(1)

4. Which of the following statements concerning relational databases is true?
 - (a) A foreign key field may be null
 - (b) A primary key field may be null
 - (c) All relations must be in at least third normal form
 - (d) Many to many mappings cannot be represented
 - (e) The primary key fields of a relation must be adjacent.

(1)

5. Given a relation country(name, continent, population) which of the following is a valid SQL statement?
 - (a) SELECT continent, population FROM country GROUP BY continent
 - (b) SELECT continent, SUM(population) FROM country GROUP BY continent
 - (c) SELECT name, population FROM country GROUP BY continent
 - (d) SELECT name, SUM(population) FROM country GROUP BY continent
 - (e) None of the above OR more than one of the above.

Questions 6 to 10 relate to the following database:

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.

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

Competitor		
cid	name	nationality
01	Pat	British
02	Hilary	British
03	Sven	Swedish
04	Pierre	French

Event		
eid	description	winner
01	running	
02	jumping	
03	throwing	

Competes	
cid	eid
01	01
02	01
03	02
04	02
04	03

6. Select the true statement:
- (a) There is a British competitor in every event
 - (b) Pierre does not compete in any event
 - (c) Sven has been entered in two events
 - (d) Pat is competing in the jumping event
 - (e) Hilary has entered only the running event.

(1)

7. Identify the result of the following SQL statement:

```
SELECT eid FROM Competes, Competitor
WHERE Competes.cid=Competitor.cid
AND nationality = 'Swedish'
```

- (a) 01
- (b) 02
- (c) 03
- (d) 04
- (e) None of the above.

(1)

8. Select the true statement:

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

(1)

9. The database is to be modified. A time in seconds is to be recorded for each runner, a distance in metres is to be entered for each jumper and for each thrower. Select the most appropriate modification:

- (a) A numeric attribute should be added to the Competes table
- (b) A numeric attribute should be added to the Event table
- (c) A numeric attribute should be added to the Competitor table
- (d) Three numeric attribute should be added to the Competitor table
- (e) Two numeric attributes should be added to the Competitor table.

(1)

10. 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:

- (i) A list of the names and nationalities of the competitors for a given event where the event description is given.
- (ii) 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 Competes
- (c) An index on name in Competitors
- (d) An index on description in Events
- (e) An index on eid in Event.

(1)

11. 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) 4
- (b) 6
- (c) 8
- (d) Cannot be calculated from the information given
- (e) None of the above.

(1)

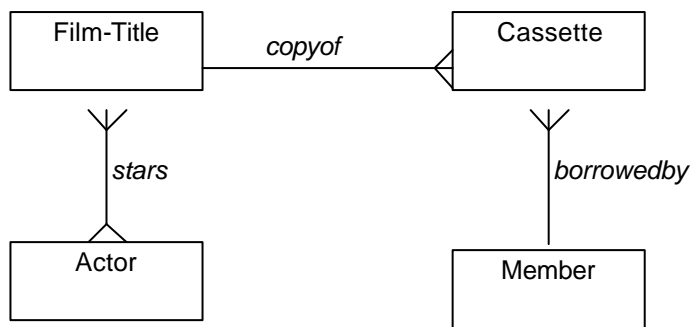
12. Which of the following is **not** usually part of the responsibilities of a database administrator?

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

(1)

13. Which of the following is the **most important** advantage of an "enterprise-wide" relational database management system compared against an ad-hoc arrangement?
- (a) Access to data will be faster.
 - (b) Confidentiality is assured.
 - (c) Disk space is used more efficiently.
 - (d) Inconsistencies are avoided.
 - (e) Network traffic is reduced.
- (1)
14. Which of the following **best** describes the internal level of the ANSI/SPARC three level architecture?
- (a) The internal level is concerned with the data as seen by individuals internal to the enterprise.
 - (b) The internal level is concerned with the how stored fields are represented and which indices exist.
 - (c) The internal level is concerned with the layout of records and their locations within disk blocks.
 - (d) The internal level is concerned with the users' view of the data .
 - (e) The internal level provides a conceptual view of the data structure.
- (1)
15. Which of the following **best** describes the costs of the operations insert, delete and seek on a table where records are stored in primary key order? Deleted records may be "flagged".
- (a) insert and delete are cheap, seek is expensive
 - (b) insert and delete are expensive, seek is cheap
 - (c) insert is expensive, delete and seek are cheap
 - (d) insert, delete and seek are all cheap
 - (e) insert, delete and seek are all expensive.
- (1)

The following entity relation diagram concerns a database to represent a video library.



Every video cassette has a unique bar code pasted to the spine. Every member of the video library has a membership card with a unique bar code.

The database is required to record which cassettes currently are on loan to whom and the dates they were loaned out. Members are charged fees when cassettes are returned.

One film may have many star actors. One actor may star in many films.

Each of the following questions consists of an assertion and a reason. Indicate your answer from the alternatives below.

	Assertion	Reason	
a	True	True	Reason is a correct explanation
b	True	True	Reason is NOT a correct explanation
c	True	False	
d	False	True	
e	False	False	

	Assertion		Reason
16.	It is not possible to establish which titles a given member has on loan	because	There is no direct relation between Member and Title (1)
17.	The “crows-foot” on the <i>borrowedby</i> relation is in the wrong position	because	At any time each cassette is being borrowed by only one member (1)
18.	The cardinality of the <i>stars</i> relation is incorrect	because	ER models may not have many-to-many relations (1)
19.	The “loan-date” should not be implemented as an attribute on the Members table	because	a given member may have several films on loan with different due dates (1)
20.	It is not possible to list all previous loans made to a particular member.	because	Only the current loans are recorded (1)

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

(1)






22. Most of the following are API calls defined in a particular database cursor system. Select the **least** plausible call.
- (a) **SQLExecute**: Executes a statement.
 - (b) **SQLConnect**: Connects to a specific driver by data source name, user ID, and password
 - (c) **SQLBindCol**: Binds an application variable to a column.
 - (d) **SQLBindRow**: Binds an application variable to a row.
 - (e) **SQLRowCount**: Returns the number of rows affected by an insert update, or delete request.

(1)

The following scenario applies to questions 23, 24 and 25.

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.

23. Select the most appropriate ER diagram for the above scenario:



(1)

24. Which of the following is the **best** approach to implementing the *attends* relation using a relational database system?

- (a) A repeating field *attends* is included as part of the *event* table
- (b) A repeating field *attends* is included as part of the *student* table
- (c) A secondary *attends* key is added to 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.

(1)

25. 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) A many-to-many relation between Events and Time-Slots is established
- (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 length of the event in minutes.
- (d) Each event has an attribute “start” which refers to Time-Slots and “duration” which gives the number of slots spanned.
- (e) Each event has two attributes “first” and “last” each of which refer to Time-Slots.

(1)

26. Which one of the following is **not** a method of implementing transactions at the physical level?

- (a) Differential files
- (b) Shadow-paging
- (c) Log-files with deferred updates
- (d) Log-files with immediate updates.
- (e) Branch and bound.

(1)

27. A lack of normalisation can lead to which one of the following problems?

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

(1)

28. Which one of the following problems can occur due to introducing locks in a concurrent transaction scenario?
- (a) Information overwrite
 - (b) Loss of information.
 - (c) Deadlock.
 - (d) Lack of integrity.
 - (e) None of the above.
- (1)
29. Which one of the following techniques is sometimes used to solve integrity problems in a concurrent transaction scenario?
- (a) First-come first-served.
 - (b) First-fit.
 - (c) Greedy algorithms.
 - (d) Strassens's algorithm.
 - (e) Two-phase locking.
- (1)
30. 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.
- (1)
31. 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.
- (1)

32. To transform a relation from first normal form to second normal form we must remove which one of the following?
- (a) All partial-key dependencies
 - (b) All inverse partial-key dependencies
 - (c) All repeating groups
 - (d) All transitive dependencies.
 - (e) None of the above.
- (1)
33. To transform a relation from second normal form to third normal form we must remove which one of the following?
- (a) All partial-key dependencies
 - (b) All inverse partial-key dependencies
 - (c) All repeating groups
 - (d) All transitive dependencies.
 - (e) None of the above.
- (1)
34. Which of the following **best** describes the relation between ISO SQL and ORACLE's SQL*PLUS?
- (a) ORACLE SQL*PLUS is faster than ISO SQL
 - (b) ORACLE SQL*PLUS may be installed on a wider range of platforms
 - (c) ORACLE SQL*PLUS is an attempt to implement a superset of ISO SQL
 - (d) ORACLE SQL*PLUS is the industry standard definition of ISO SQL
 - (e) ORACLE SQL*PLUS is a commercial product, ISO SQL is freeware.
- (1)
35. 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 perform too slowly
 - (b) Full normalisation may compromise existing applications/systems
 - (c) A fully normalised database may have too many tables
 - (d) Full normalisation may make some queries too complicated
 - (e) A fully normalised database may result in tables which are too large.
- (1)

The following functional dependencies apply to question 36-38:

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

36. Which if the following relational schemas might be the result of normalising R(s, q, t, u) ?

- (a) R1(s, q) R2(s, t) R3(q, u)
- (b) R1(s, q) R2(q, t) R3(t, u)
- (c) R1(s, q) R2(s, t) R3(q, u)
- (d) R1(s, q, t) R2(s, q, u)
- (e) R(s, q, t, u)

(1)

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

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

(1)

38. Which of the following best describes the relation R(e, f, g, h, i, j) ?

- (a) First Normal Form
- (b) Second Normal Form
- (c) Third Normal Form
- (d) Fourth Normal Form
- (e) Boyce Codd Normal Form.

(1)

39. Relation C is the join of relation A and relation B on condition p . Which of the following statements must be **true in all cases**?

- (a) The cardinality of C is greater than the cardinality of A
- (b) The cardinality of C is less than the cardinality of A
- (c) The arity of C is greater than the arity of A
- (d) The arity of C is less than the arity of A
- (e) None of the above.

(1)

40. Relation C is a projection of relation A. Which one 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 cardinality of C is less than the cardinality of A
- (c) The arity of C is greater than the arity of A
- (d) The arity of C is less than the arity of A
- (e) None of the above.

(1)

Total marks [40]

END OF PAPER