Matric No: _____

NAPIER UNIVERSITY

SCHOOL OF COMPUTING

CO22001

DATABASE SYSTEMS (MALAYSIA)

ACADEMIC SESSION: 2003-2004

EXAMINATION DIET: MAY 2004

EXAMINATION DURATION: 2 HOURS

READING TIME: NONE

EXAM PAPER INFORMATION

Number of pages – TWENTY-ONE

Number of questions – FORTY

Select ONE from (a) to (e)

Answers must be inserted on the EDPAC answer sheet provided using an HB pencil.

For full instructions see next page.

EXAMINER: DR. G. RUSSELL

PLEASE READ THE FULL INSTRUCTIONS BEFORE COMMENCING WRITING

Instructions to Candidates -

Write the following details in the top of the **Candidate Name** section **in this order**:

Your surname
Your Initials

In the machine readable part of the name section, make a **horizontal mark between the two brackets** on the letter of your choice to enter the following details in **machine readable** form **in this order:**

Your surname
Your initials

e.g. **[C] [H] [E] [S] [N] [E] [Y] [T]**

In the box named **Candidate Number** mark in your **matriculation number.**

In the box named **Subject Code**, mark in **001**

Leave the subject box blank.

At the end of the test, return **your answer sheet** to the invigilator.

Attempt **all** of the following questions. The test consists of 40 multiple choice questions.

All the questions offer five options. For each you are required to indicate which you consider the single most appropriate answer. Indicate your selection by making a mark in the row on the answer sheet corresponding to the question number. Use an HB pencil and make a mark the width of the column (A - E), which corresponds to your chosen answer. To change an answer put the mark in the new column and **circle** the correction.

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

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

Mark: (1)

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

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

Mark: (1)

3. 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: <u>Name</u>, DOB, Handicap, HomeClub
   b. Member: <u>Name, DOB</u>, Handicap, HomeClub
   c. Member: <u>Name, DOB</u>, Handicap, HomeClub, Yardage
   d. Member: <u>Name, DOB, Handicap, HomeClub, Yardage</u>
   e. Member: <u>Name</u>, DOB, Handicap, HomeClub, Yardage

   Mark: (1)

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

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

Mark: (1)

5. 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(<u>a, b</u>, c, d, e, f) ?

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

Mark: (1)

6. Consider the following functional dependencies:

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			

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

a. The schema R1(<u>a,b</u>,c) R2(<u>a,b</u>,d)
b. The schema R1(<u>a,b</u>) R2(<u>a</u>,c) R3(<u>b</u>,d)
c. The schema R1(<u>a</u>,b) R2(<u>b</u>,c) R3(<u>c</u>,d)
d. The schema R1(<u>a,b</u>) R2(<u>a,c</u>) R3(<u>a,d</u>)
e. The schema R(<u>a,b</u>,c,d)

Mark: (1)

7. 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(<u>cid</u>, name, nationality)
Event(<u>eid</u>, description, winner)
Competes(<u>*cid, eid*</u>)

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

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. 04
d. 03
e. None of the above

Mark: (1)

8. **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,casting,actor
  WHERE    movieid  = movie.id
  AND      actor.id = actorid
  AND      name     = 'Marilyn Monroe'
  AND      yr = 1959
;
```

b. The following SQL…

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

e. None of the above

Mark: (1)

9. A number of vats of chemical are monitored by an automatic system. Temperature and pressure readings are recorded for each vat at regular intervals.

   Currently there are 3 vats. The current procedure is to take readings four times a day at 01:00, 07:00 13:00 and 19:00.

   Each of the following schemes are being considered for storing data:

   A

   Temperature(theDate, vat1_0100, vat1_0700, vat1_1300, vat1_1900, vat2_0100, vat2_0700, vat2_1300, vat2_1900, vat3_0100, vat3_0700, vat3_1300, vat3_1900)
   Pressure(theDate, vat1_0100, vat1_0700, vat1_1300, vat1_1900, vat2_0100, vat2_0700, vat2_1300, vat2_1900, vat3_0100, vat3_0700, vat3_1300, vat3_1900)

   B

   Temperature(theDate, time, vat, value)
   Pressure(theDate, time, vat, value)

   The average pressure in vat 1 for the period 2 Feb 2003 to 5 Feb 2003 is required. The following SQL statement is an attempt at calculating this value based on schema A:

```
SELECT SUM(vat1_0100+vat1_0700+vat1_1300+vat1_1900)/16
 FROM pressure
WHERE theDate BETWEEN '2 Feb 2003' AND '5 Feb 2003'
```

   Given that the reading for 3 Feb at 0700 is null, but that all other values are correct; select the statement that best describes the outcome:

   a. All values for 3 Feb are discarded, the remaining 12 values are summed and divided by 16
   b. The null value ropagates and zero rows are returned
   c. The 15 correct values are added – and then correctly divided by 15
   d. The 15 correct values are added – but then erroneously divided by 16
   e. The null value ropagates and a row with the value null is returned

   Mark: (1)

10.

Departments	
DepNo	**Depname**
1	Computng
2	Electrical
3	Geography
4	History
5	Business

Employees	
Empno	**Empname**
1	Gordon
1	Ken
1	Brian
1	Colin
1	George

WorkFor	
Empno	**Depno**
1	1
3	2
4	1
3	3
1	2
2	5

Using the SQL Scenario, what is the CARDINALITY of the relationship between the entity type departments and the entity type workfor?

a. 1:N
b. 2
c. N:1
d. M:N
e. 5

Mark: (1)

11. 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 the industry standard definition of ISO SQL
c. ORACLE SQL*PLUS is an attempt to implement a superset of ISO SQL
d. ORACLE SQL*PLUS is a commercial product, ISO SQL is freeware.
e. ORACLE SQL*PLUS is faster than ISO SQL

Mark: (1)

12. The role of a DBA includes which of the following?

    a. user interfaces, salary budgeting, performance monitoring.
    b. Loading data, evaluating new database systems, performance monitoring
    c. security, system testing, java programming
    d. Installing databases, C++ programming, user support.
    e. Supporting all programming languages which might be used with a database.

Mark: (1)

13. Which of the following best describes the internal level of the ANSI/SPARC three level architecture?

    a. The internal level is concerned with the layout of records and their locations within disk blocks.
    b. The internal level is concerned with the data as seen by individuals internal to the enterprise.
    c. The internal level is concerned with the users' view of the data.
    d. The internal level is concerned with the how stored fields are represented and which indices exist.
    e. The internal level provides a conceptual view of the data structure.

Mark: (1)

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

Mark: (1)

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

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

Mark: (1)

16. Which of the following can be found in a DSL?

    a. LDD
    b. DLD
    c. DLA
    d. PLA
    e. DML

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

Mark: (1)

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

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

Mark: (1)

19. 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:

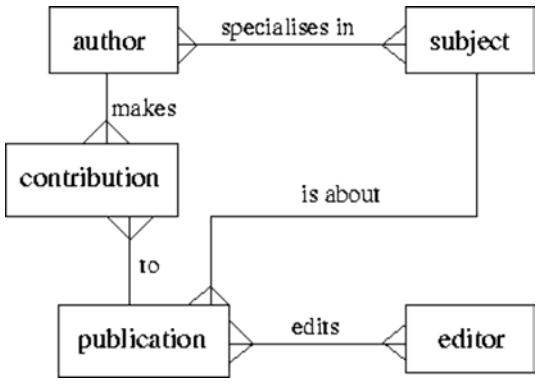

    Indicate the relation which has an incorrect cardinality shown:

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

Mark: (1)

20. In IMMEDIATE UPDATE, if the DBMS fails and is then restarted then the log file is parsed. Which one of the following statements is true?

    a. Never re-apply the old data.
    b. If a log entry DOES NOT end in COMMIT, apply the old data.
    c. Only apply the old data if the log entry ends in ABORT.
    d. Re-apply the old data ONLY if the entry ends in COMMIT.
    e. Always re-apply the new data.

Mark: (1)

21. Using DEFERRED UPDATE, modifications made by a transaction which has not yet aborted or committed

    a. are applied in the same way as in IMMEDIATE UPDATE.
    b. are applied to the log file and then the disk
    c. are applied only to the log file
    d. are applied only to the disk
    e. None of the above

Mark: (1)

22. 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 immediate updates
    d. Log-files with deferred updates
    e. None of the above

Mark: (1)

23. 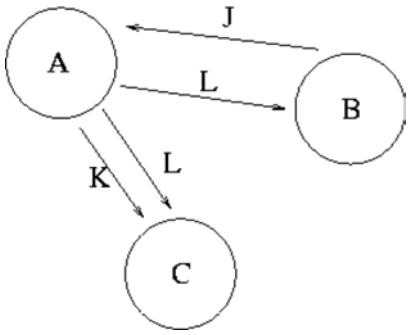

From Transaction Scenario 1, given the precedence graph, which of the following is TRUE?

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

Mark: (1)

24. 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 arity of C is greater than the arity of A
c. The cardinality of C is less than the cardinality of A
d. The arity of C is less than the arity of A
e. None of the above

Mark: (1)

25.

Relation P	
ColW	**ColX**
A	4
B	5
C	6

Relation Q	
ColY	**ColZ**
B	7
D	4
C	6

Consider the relations P and Q above. The number of rows in the unconditional join, or Cartesian product of P and Q is

a. 6
b. 8
c. 4
d. 9
e. None of the above.

Mark: (1)

26. π is the projection operator. Σ is the selection operator. R is a relation. Select the relational expression which could possibly return the following result:

a	**c**
1	2
2	3

a. $\sigma_{a,\,c}\,R$
b. $\pi_{a,c}(\sigma_{a=c}\,R)$
c. $\pi_{a<c}\,(\pi_{a,\,c}\,R)$
d. $\pi_{a<2}\,R$
e. $\sigma_{a<c}\,(\pi_{a,\,c}\,R)$

Mark: (1)

27.

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?

a. UNION
b. INNER JOIN
c. LEFT or RIGHT OUTER JOIN
d. CARTESIAN PRODUCT
e. INTERSECTION

Mark: (1)

28. 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 seek operations may be faster.
d. More disk space may be required
e. Certain updates may be slower

Mark: (1)

29. With respect to the B+ tree index method, select the TRUE statement:

    a. Records are physically stored in primary key order.
    b. B+ trees use a hashing algorithm.
    c. The index tree may become unbalanced as a result of updates.
    d. Only the primary key field may have a B+ tree index.
    e. None of the above.

Mark: (1)

30. Which of the following is TRUE concerning TRANSFER TIME for hard drives?

    a. Must be measured as a ratio of seek time.
    b. Is the time to move data from the disk surface to the hard drive.
    c. Transfer time is measured in Mbytes.
    d. Seek time is greater than transfer time.
    e. Depends on the number of files being transferred.

Mark: (1)

31. The purpose of Embedded SQL is to allow

    a. Databases to be embedded in SQL
    b. Programming language to be embedded in SQL
    c. SQL queries to be executed as part of a programming language.
    d. Programs to be embedded in a database.
    e. None of the above

Mark: (1)

32. Accessing data via embedded SQL requires which of the following?

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

Mark: (1)

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

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

Mark: (1)

34. Which of the following is a type of lock which cannot be obtained in Oracle?

    a. Exclusive Lock
    b. Write Lock
    c. Shared Lock
    d. Read lock
    e. Insert lock

Mark: (1)

35. 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. keys are aquired in the first phase.
    d. locks can only be acquired on primary keys.
    e. None of the above.

Mark: (1)

36. 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 both relation A and relation B
    b. the foreign key is placed in relation A.
    c. you have to use create another relation first.
    d. the foreign key is placed in relation B.
    e. you should subsume A into B.

Mark: (1)

37. When mapping EER superclasses/subclasses, which of the following options is not a valid possibility?

    a. Map each subclass onto separate relations
    b. Use one relation for the superclass
    c. Map each subclass using class inheritance
    d. Use a relation for the superclass and a separate relation for each of the subclasses.
    e. All of the above are valid options.

Mark: (1)

38. A primary key in a table

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

Mark: (1)

39. Which of the following is not used to support recovery and consistency in a database system?

    a. A dump
    b. Two phase commit
    c. Checkpointing
    d. A journal
    e. Access logging

Mark: (1)

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

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

Mark: (1)

Total marks [40]

END OF PAPER