Matric No: _____

NAPIER UNIVERSITY

SCHOOL OF COMPUTING

CO22001

DATABASE SYSTEMS

ACADEMIC SESSION: 2003-2004

EXAMINATION DIET: AUGUST

TRIMESTER: ONE

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

Mark: (1)

2. 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 foreign key attribute in a record always refers to another record which does not contain nulls
   e. a non-null foreign key attribute always refers to another record

Mark: (1)

3. Assume the relation R(<u>A, B</u>, C ,D, E) is in 2NF. Which of the following functional dependencies must be **TRUE**?

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

Mark: (1)

4. To transform a relation from first normal form to second normal form we must remove which one of the following?

   a. All repeating groups
   b. All partial-key dependencies
   c. All inverse partial-key dependencies
   d. All transitive dependencies
   e. None of the above

Mark: (1)

5. 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 best describes the relation R(<u>e,f,g</u>,h,I,j)?

   a. Boyce Codd Normal Form
   b. Second Normal Form
   c. Third Normal Form
   d. First Normal Form
   e. Forth Normal Form

Mark: (1)

6. Continuing from the previous question.

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

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

   Mark: (1)

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

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

   Mark: (1)

8. SELECT a,b FROM c,d
   where c has 10 records and d has 10 records results in
   a. a table with 0 records
   b. a table with 10 records
   c. a table with 100 records
   d. a table with 1000 records
   e. The number of records cannot be predicted.

   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(<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

Select the true statement

a. Sven has been entered in two events
b. Pierre does not compete in any event
c. There is a British competitor in every event.
d. Hilary has entered only the running event
e. Pat is competing in the jumping event

Mark: (1)

10. 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(_abelled_, 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 statement is intended to return the station id of the station which recorded the lowest value for sunshine. Select the simplest SQL statement which does this.

a. SELECT abelled , MIN(sunshine) FROM station;

b.
```
SELECT  abelled
FROM  station
WHERE  sunshine IN (
  SELECT MIN(sunshine)
  FROM  station
  );
```

c.
```
SELECT  abelled  FROM station
WHERE sunshine = MIN;
```

d.
```
SELECT  abelled
FROM  station AS A station AS B
WHERE  A.sunshine < B.sunshine;
```

e.
```
SELECT  abelled , MIN(sunshine)
FROM station
GROUP BY  abelled ;
```

Mark: (1)

11. 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 abelled ;
    b. SELECT AVG(name), rainfall
       FROM region, station
       WHERE regionid = station.region
       GROUP BY rainfall;
    c. SELECT name, AVG(rainfall)
       FROM region, station
       WHERE regionid = station.region
       GROUP BY name;
    d. SELECT AVG(rainfall)
       FROM station
       GROUP BY abelled ;
    e. SELECT AVG(name) ,AVG(rainfall)
       FROM region, station
       WHERE regionid = station.region
       GROUP BY rainfall;

    Mark: (1)

12. 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 is expensive, delete and seek are cheap.
    b. insert and delete are cheap, seek is expensive
    c. insert and delete are expensive, seek is cheap
    d. insert, delete and seek are all expensive
    e. insert, delete and seek are all cheap

    Mark: (1)

13. 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. consistency in data use
    c. reduced data redundancy
    d. the enforcement of standards.
    e. performance measurement

Mark: (1)

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

Mark: (1)

15.

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

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

a. job(<u>reference</u>, employer, salary) requirement(<u>job, skill</u>)
b. job(reference, employer, salary) requirement(<u>job</u>, skill)
c. job(reference, <u>employer</u>, salary) requirement(job, <u>skill</u>)
d. job(<u>reference, employer</u>, salary) requirement(<u>job</u>, skill)
e. job(<u>reference</u>, employer, salary) requirement(<u>job</u>, skill)

Mark: (1)

16. 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 a foreign key in both A and B.
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 the foreign key in the B relation.

Mark: (1)

17. 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.

    Select the most appropriate ER diagram for the above scenario:

a.
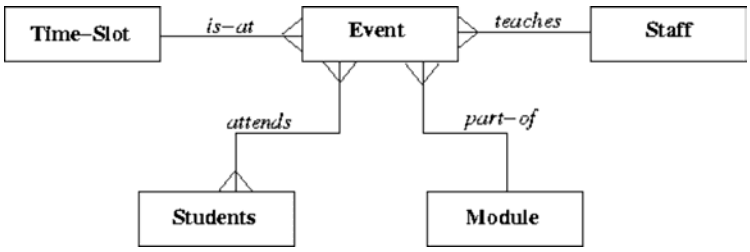

b.
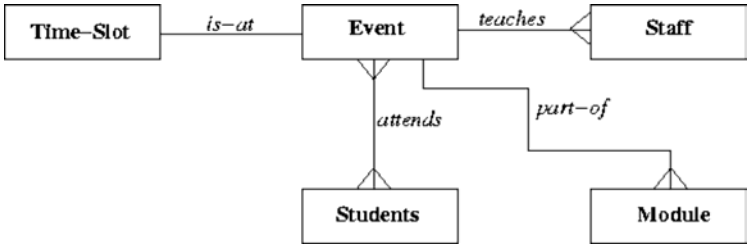

c.
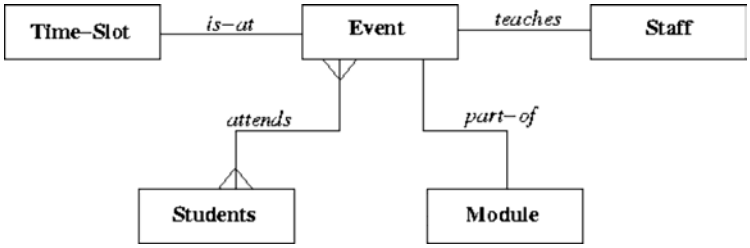

d.
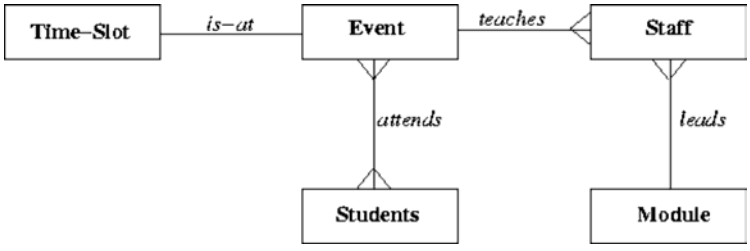

e.
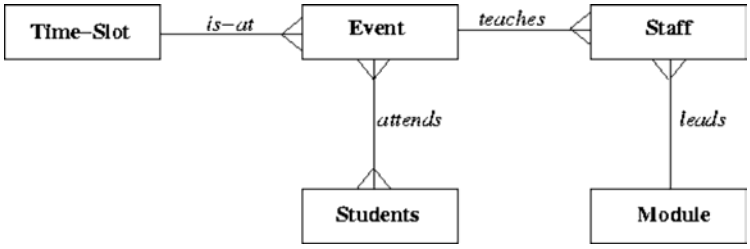

Mark: (1)

18. Continuing from the previous question.

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

    Mark: (1)

19. In the ANSI-SPARC three level architecture, which of the following is the name of one of the levels?

    a. constructional
    b. attribute
    c. internal
    d. index
    e. extra

    Mark: (1)

20. Which of the following is a good example of what is meant by serialisability?

    a. The situation where a cascade abort occurs.
    b. All transactions happen one after another.
    c. The result of the transactions is the same as if the transactions went one after another.
    d. The situation where the Lost Update problem exists.
    e. All disk access happens one after another.

    Mark: (1)

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

Mark: (1)

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

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

Mark: (1)

23. 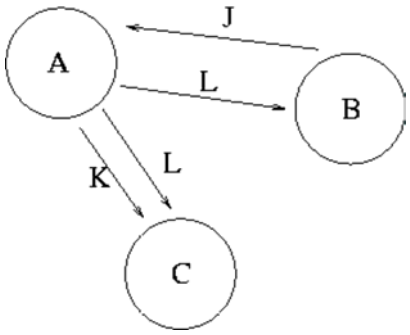

The graph in the Transaction Scenario was produced from the following schedule. Time

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:

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

Mark: (1)

24. R

ColA	ColB
A	1
C	4
D	6
E	6

S

ColC	ColD
C	2
D	1
G	6
J	7

Using the Relation Algebra Scenario, what is the cardinality of the resulting relation of:
R FULL OUTER $JOIN_{R.ColA = S.ColC}$ S

a. 6
b. 8
c. 0
d. 2
e. None of the above

Mark: (1)

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

Mark: (1)

26. 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
		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
C	4	C	4

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)

27. Continuing from the previous question.

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

    a.

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

    b.

Col1	Col2	Col3	Col4
A	1	A	1
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)

28. When comparing Main Memory of a computer to the data stored on a Hard Drive, select the TRUE statement.

    a. Main memory's transfer rate is lower than a hard drive
    b. Main memory's transfer rate is higher than a hard drive
    c. Transfer rates cannot be compared.
    d. Main memory's transfer rate is about the same as a hard drive
    e. Transfer rate is dependent on the seek time

Mark: (1)

29. DBMSs usually use a combination of main memory and secondary storage to hold the database. Select the FALSE statement.

    a. Main memory is usually erased when the power is switched off.
    b. Secondary storage is faster than main memory.
    c. Secondary storage is persistent.
    d. Main memory is too expensive to be able to hold large databases.
    e. Storing a database in secondary storage is a cost-effective approach.

Mark: (1)

30. Which of the following is TRUE when considering hash tables?

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

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. programs to be embedded in a database.
    d. SQL queries to be executed as part of a programming language.
    e. None of the above

Mark: (1)

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

    Mark: (1)

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

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

    Mark: (1)

34. With respect to Two-Phase Locking. Select the FALSE statement.

    a. A transaction may start to release locks at any time.
    b. Locks may be made at any point in the transaction schedule.
    c. Locks are released on an ABORT.
    d. The phases are "lock-acquisition" and "lock-release".
    e. Two-Phase Locking requires key pairs.

    Mark: (1)

35. Locking was introduced into databases so that

    a. reading and writing is possible.
    b. all simultaneous transactions are prevented.
    c. keys can be provided to maintain security.
    d. consistency can be enforced.
    e. passwords can be provided to maintain security

Mark: (1)

36. In the Database Analysis Life Cycle, what are the last two steps in the cycle before it begins to repeat?

    a. “Operation” and then “Payment”
    b. “Operation” and then “Training”
    c. “Operation” and then “Maintenance and Evolution”
    d. “Testing and Evaluation” and then “Operation”
    e. “Implementation” and then “Training”

Mark: (1)

37. In a relation, an instance is stored as

    a. a Many to many relationship
    b. an Attribute
    c. a Row
    d. a Column
    e. an index

Mark: (1)

38. The DOMAIN of an attribute is which of the following?

    a. The range of values that the attribute is allowed to take.
    b. The type of the attribute.
    c. Which tables the attribute is allowed to be in.
    d. More than one of the above.
    e. None of the above.

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 NO-UNDO/NO-REDO algorithm.
d. It is also called the UNDO/REDO algorithm.
e. None of the above.

Mark: (1)

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

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

Mark: (1)

Total marks [40]

END OF PAPER