

Matric No: _____

NAPIER UNIVERSITY
SCHOOL OF COMPUTING

CO22001J
DATABASE SYSTEMS
(JAMES WATT COLLEGE)

ACADEMIC SESSION: 2003-2004

EXAMINATION DIET: AUGUST

TRIMESTER: ONE

EXAMINATION DURATION: 2 HOURS

READING TIME: NONE

EXAM PAPER INFORMATION

Number of pages – TWENTY

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] [F]

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

Mark: (1)

2. If a system can enforce referential integrity, then this ensures that
- a. a foreign key attribute in a record always refers to another record which does not contain nulls.
 - b. a record can never contain a null value for a foreign key attribute.
 - c. a record is always referred to from another record.
 - 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 contains nulls.

Mark: (1)

3. Which normalisation transformation corresponds to "Eliminating partial key dependencies"?
- a. unnormalised to 1NF
 - b. 1NF to 2NF
 - c. 2NF to 3NF
 - d. 3NF to BCNF
 - e. None of the above.

Mark: (1)

4. Assume the relation $R(\underline{A}, \underline{B}, C, D, E)$ is in 2NF. Which of the following functional dependencies must be **TRUE**?

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

Mark: (1)

5. Sometimes it is best NOT to fully normalise a database. Select the strongest argument which would support this statement.

- a. A fully normalised database may perform too slowly.
- b. A fully normalised database may result in tables which are too large.
- c. Normalisation will increase the number of attributes in each table.
- d. An un-normalised database may make some queries too complicated.
- e. Normalisation will reduce the number of tables.

Mark: (1)

6. Given the following relation and dependencies, select the option that is the result of fully normalising the relation to BCNF.

$R(\underline{a}, \underline{b}, c, d, e)$

$c \rightarrow a$

$d, e \rightarrow a, b, c$

$c, d \rightarrow e, b$

- a. $R(\underline{b}, \underline{c}, \underline{d}, e)$
 $R_1(\underline{c}, a)$
- b. $R(\underline{b}, c, \underline{d}, \underline{e})$
 $R_1(\underline{c}, a)$
- c. $R(\underline{b}, \underline{c}, d, e)$
 $R_1(\underline{c}, a)$
- d. all of the above
- e. none of the above

Mark: (1)

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

Mark: (1)

8.

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

Choose the SQL statement which will return details of the jobs of interest to a candidate with experience of Unix Administration.

- ```
SELECT * FROM job
WHERE reference='Unix Admin'
AND skill = 'Unix Admin';
```
- ```
SELECT 'Unix Admin' FROM job, requirement;
```
- ```
SELECT * FROM job
WHERE employer = 'Napier'
OR employer = 'GCHQ';
```
- ```
SELECT * FROM job, requirement
WHERE reference=job AND skill='Unix Admin';
```
- ```
SELECT * FROM job
WHERE skill = 'Unix Admin';
```

Mark: (1)

## 9. 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,casting,actor
WHERE movieid = movie.id
AND name = 'Marilyn Monroe'
;
```

c. The following SQL...

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

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)

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

Mark: (1)

11. Continuing from the previous question.

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

```
SELECT stationid
FROM station AS A station AS B
WHERE A.sunshine < B.sunshine;
```
- c. 

```
SELECT stationid, MIN(sunshine)
FROM station
GROUP BY stationid;
```
- d. 

```
SELECT stationid, MIN(sunshine) FROM station;
```
- e. 

```
SELECT stationid FROM station
WHERE sunshine = MIN;
```

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

14. In Enhanced ER diagrams, a subclass

- a. is part of Chun's notation.
- b. may contain only one superclass.
- c. can only exist in Chen's notation.
- d. is contained in one superclass.
- e. may contain many superclasses.

Mark: (1)

15. 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. A one-to-many relation between Events and Time-Slots is established.
- b. A many-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.

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

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

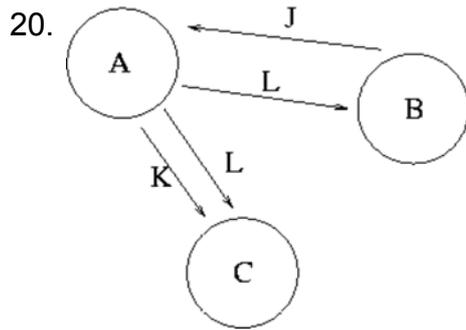
- a. attribute
- b. internal
- c. constructional
- d. index
- e. extra

Mark: (1)

19. Consider the development process for the "Three level model". Select the statement which is most generally TRUE.

- a. "Conceptual Design" is dependent on the target DBMS platform.
- b. "Data Model" concerns the external users' view of the data.
- c. "Physical Design" may refer to indexes.
- d. "Physical Design" is used to produce an ER model.
- e. An ER model is required prior to "Conceptual Design".

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 loops are missing.
- The precedence graph is accurately drawn.
- 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.
- Errors in BOTH loops and arrow direction.

Mark: (1)

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

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

Mark: (1)

22. Select the problem which is demonstrated by the following schedule.

| Time | Transaction A | Transaction B |
|------|---------------|---------------|
| T1   | Write(X)      |               |
| T2   |               | READ(X)       |
| T3   | Abort         |               |

- a. Inconsistency Analysis
- b. Uncommitted Dependency.
- c. Incoherency Analysis
- d. Inconsolable Update
- e. Lost Update

Mark: (1)

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

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

Mark: (1)

24.

| Relation P |      | Relation Q |      |
|------------|------|------------|------|
| CoIW       | CoIX | CoIY       | CoIZ |
| A          | 4    | B          | 7    |
| A          | 5    | D          | 4    |
| B          | 4    | C          | 6    |
| B          | 5    |            |      |
| C          | 6    |            |      |

Consider the relations P and Q above. Select the number of rows in the table resulting from the following join.

P RIGHT OUTER JOIN<sub>CoIX = CoIZ</sub> Q

- a. 4
- b. 2
- c. 5
- d. 3
- e. None of the above.

Mark: (1)

25. In Relational Algebra, a "tuple"

- a. is a collection of relations describing a mini-world view.
- b. is a collection of attributes describing some real-world entity.
- c. is a set of atomic values.
- d. is an index.
- e. is effectively a column of a relation.

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    |
| C    | 4    | C    | 4    |

c.

| Col1 | Col2 | Col3 | Col4 |
|------|------|------|------|
| A    | 1    | A    | 1    |
| B    | 3    |      |      |
| 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.

| Relation P |      | Relation Q |      |
|------------|------|------------|------|
| CoIW       | CoIX | CoIY       | CoIZ |
| A          | 4    | B          | 7    |
| B          | 5    | D          | 4    |
| C          | 6    | 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. 4
- b. 8
- c. 6
- d. 9
- e. None of the above.

Mark: (1)

28. When accessing a disk block, the seek time

- a. is insignificant in comparison to transfer times.
- b. is about the same as transfer times.
- c. is measured in nanoseconds.
- d. is the time taken to search for data in a sorted list of database rows.
- e. greatly exceeds transfer times.

Mark: (1)

29. 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. A good hash function is one which results in many hash collisions.
- d. Hash tables speed up sequential record scanning.
- e. DBMS systems can only hash on numbers.

Mark: (1)

30. Using SERIAL data organisation -

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

Mark: (1)

31. Select the TRUE statement.

- a. SQL is embedded within C++ to handle sequential file processing.
- b. SQL embedded in C++ provides facilities to extract data from a database.
- c. SQL is embedded within a C++ program to increase the speed of the C++ program.
- 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)

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

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. Two-phase locking.
- c. Strassens's algorithm.
- d. First-fit.
- e. Greedy algorithms.

Mark: (1)

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

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

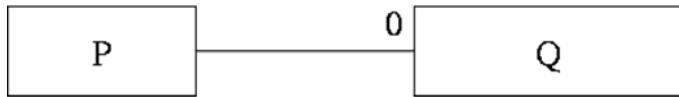
Mark: (1)

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

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

Mark: (1)

36. When producing relations for the diagram shown, where almost all P entities relate to a Q entity:



Which of the following is usually best?

- a. It is best to keep P and Q as separate relations.
- b. This cannot be mapped into relations.
- c. Q should be subsumed into P.
- d. Have an intermediate table containing the primary keys from both P and Q.
- e. P should be subsumed into Q.

Mark: (1)

37. When mapping ER models into relations, which of the following is NOT true?

- a. Each entity type is mapped into a relation.
- b. Each attribute is mapped into a column of the corresponding relation.
- c. Each m-n relationship is mapped into a new relation.
- d. Each individual entity is mapped into a row of the corresponding relation.
- e. Each 1-m relationship is mapped into a new relation.

Mark: (1)

38. In a relation, an instance is stored as

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

Mark: (1)

39. Aborting a transaction

- a. removes changes made in a transaction after it has committed.
- b. results in deadlock.
- c. removes changes made so far in the current transaction.
- d. is only possible in Microsoft Access.
- e. deletes the database for security reasons.

Mark: (1)

40. During Rollforward, which of the following is not true?

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

Mark: (1)

Total marks [40]

**END OF PAPER**