

NAPIER UNIVERSITY
SCHOOL OF COMPUTING

FIRST DIET (SEMESTER TWO) EXAMINATION

SESSION 1999-2000

CS 71012: DATABASE SYSTEMS

DATE:

TIME ALLOWED: 2 HOURS

START TIME:

FINISH TIME:

EXAMINER:

A. CUMMING
K. CHISHOLM

QUESTION PAPER DATA

Number of Pages - TWELVE
Number of Questions – FORTY
Number of Sections - ONE

INSTRUCTIONS TO CANDIDATES

Answer ALL questions.

For each question select ONE from (a) to (e).

Answers must be inserted on the answer sheet provided.

PLEASE READ FULL INSTRUCTIONS BEFORE COMMENCING WRITING

NAPIER UNIVERSITY
SCHOOL OF COMPUTING

FIRST DIET (SEMESTER TWO) EXAMINATION

SESSION 1999-2000

CS 22004: DATABASE SYSTEMS

DATE:

TIME ALLOWED: 2 HOURS

START TIME:

FINISH TIME:

EXAMINER:

**G. RUSSELL
J. MURRAY**

QUESTION PAPER DATA

**Number of Pages - TWELVE
Number of Questions – FORTY
Number of Sections - ONE**

INSTRUCTIONS TO CANDIDATES

**Answer ALL questions.
For each question select ONE from (a) to (e).
Answers must be inserted on the answer sheet provided.**

PLEASE READ FULL INSTRUCTIONS BEFORE COMMENCING WRITING

1. In the Database Analysis Life Cycle, the "Testing and evaluation" stage
 - (a) comes after "Database Design" so that the design can be tested before implementation.
 - (b) comes after "Operation" so that the normal DBMS operation can be tested.
 - (c) is done before the feasibility study.
 - (d) is done before the "Operation" stage.
 - (e) is part of the "implementation and loading" stage. (1)

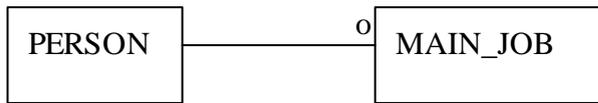
2. Consider the development process for the "Three level model". Select the statement which is most generally TRUE
 - (a) "Conceptual Design" is dependent on the target DBMS platform.
 - (b) "Data Model Mapping" concerns the external users' view of the data.
 - (c) "Physical Design" is used to produce an ER model.
 - (d) "Physical Design" may refer to indexes.
 - (e) An ER model is required prior to "Conceptual Design". (1)

3. In ER Modelling, a "Fan Trap"
 - (a) may be fixed using the Navathe algorithm.
 - (b) may occur when 1:m relationships fan in from a single entity.
 - (c) may occur when 1:m relationships fan out from a single entity.
 - (d) may occur when m:n relationships fan in from a single entity.
 - (e) may occur when m:n relationships fan out from a single entity. (1)

4. In ER Modelling, a "Chasm traps" can occur when entities are related via a relationship with
 - (a) no optionality
 - (b) no partial participation
 - (c) partial differentiation
 - (d) partial overhangs
 - (e) partial participation (1)

5. Which of the following is offered through the use of EER (Enhanced ER) Modelling?
 - (a) Canonisation
 - (b) Confirmation
 - (c) Many to many relation
 - (d) Partial participation
 - (e) Specialisation (1)

6. When producing relations for the diagram shown, where almost all PERSON entities relate to a MAIN_JOB:



Which of the following is most generally true?

- (a) it is best to keep PERSON and MAIN_JOB as separate relations.
 - (b) MAIN_JOB should be subsumed into PERSON.
 - (c) NULL attributes will be produced no matter what you do.
 - (d) PERSON should be subsumed into MAIN_JOB.
 - (e) None of the above. (1)
7. In Relational Algebra, a "tuple"
- (a) is a collection of attributes describing some real-world entity.
 - (b) is a collection of relations describing a mini-world view.
 - (c) is a set of atomic values.
 - (d) is effectively a column of a relation.
 - (e) is an index. (1)

Relation P	
ColW	ColX
A	4
B	5
C	6

Relation Q	
ColY	ColZ
B	7
D	4
C	6

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

P RIGHT OUTER JOIN_{ColX = ColZ} Q

- (a) 2
- (b) 3
- (c) 4
- (d) 7
- (e) None of the above. (1)

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

P RIGHT OUTER JOIN_{colX = colZ} Q

- (a) 2
 (b) 3
 (c) 4
 (d) 7
 (e) None of the above. (1)
10. 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) 6
 (c) 8
 (d) 9
 (e) None of the above. (1)
11. The following SQL produces a list of student names where the student had an "A" grade in the exam. Select the relational algebra statement which best represents the SQL.

```
SELECT name
FROM student, result
WHERE student.number = result.number
AND result.grade = 'A';
```

- (a) $\Pi_{\text{name}}(\sigma_{\text{grade} = 'A'}(\text{student} \bowtie \text{result}))$
 (b) $\sigma_{\text{name}}(\text{student} \bowtie (\Pi_{\text{grade} = 'A'}(\text{result})))$
 (c) $\sigma_{\text{name}}(\text{student} \bowtie (\sigma_{\text{grade} = 'A'}(\text{result})))$
 (d) $\sigma_{\text{name}}(\text{student} \bowtie_{\text{number}=\text{grade.number}} (\sigma_{\text{grade} = 'A'}(\text{result})))$
 (e) None of these. (1)
12. Select the problem which is demonstrated by the following schedule.

Time	Transaction A	Transaction B
T1	WRITE(X)	
T2		READ(X)
T3	ABORT	

- (a) Incoherency Analysis
 (b) Inconsistency Analysis
 (c) Inconsolable Update
 (d) Lost Update
 (e) Uncommitted Dependency. (1)

13. Locking can be fine grained or coarse grained. Select the TRUE statement.
- (a) Course grained locking improves parallelism.
 - (b) Fine grained locking is cheaper to implement.
 - (c) Fine grained locking is more expensive to implement.
 - (e) Only course grained locking results in enforcing serializability.
 - (e) Only fine grained locking results in enforcing serializability. (1)
14. With respect to Two-Phase Locking. Select the FALSE statement.
- (a) A transaction may start to release locks at any time.
 - (b) Locks are released on an ABORT.
 - (c) Two-Phase Locking requires key pairs.
 - (d) Locks may be made at any point in the transaction schedule.
 - (e) The phases are "lock-acquisition" and "lock-release". (1)
15. 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/NO-REDO algorithm.
 - (d) it is also called the UNDO/REDO algorithm.
 - (16) none of the above. (1)
16. With Immediate Update, writing to an attribute results in the DBMS
- (a) Releasing concomitant locks.
 - (b) Immediately performing a COMMIT.
 - (c) Immediately writing the change to disk.
 - (d) Writing the change to the log.
 - (e) Writing the old and new attribute value to the log. (1)
17. Select the situation which cannot result in a ROLLBACK.
- (a) A cascade abort is required.
 - (b) A committed transaction needs to be undone.
 - (c) A transaction is aborted.
 - (d) The DBMS application was corrupted by a power spike.
 - (e) The machine on which the DBMS was executing failed suddenly. (1)

18. DBMS's usually use a combination of main memory and secondary storage to hold the database. Select the FALSE statement.
- (a) Main memory is too expensive to be able to hold large databases.
 - (b) Main memory is usually erased when the power is switched off.
 - (c) Secondary storage is faster than main memory.
 - (d) Secondary storage is persistent.
 - (e) Storing a database in secondary storage is a cost-effective approach. (1)
19. With respect to a DBMS using a hard drive as secondary storage. Select the TRUE statement.
- (a) A block is the smallest unit which can be read from the disk by the operating system.
 - (b) Data can be read from a file in "attribute-size" amounts.
 - (c) Data should be stored in contiguous blocks to maintain consistency.
 - (d) The database should be backed up onto primary memory at regular intervals.
 - (e) The DBMS manages the disk blocks. (1)
20. With respect to the B⁺ tree index method, select the TRUE statement:
- (a) The index tree may become unbalanced as a result of updates.
 - (b) Records are physically stored in primary key order.
 - (c) B⁺ trees use a hashing algorithm.
 - (d) Only the primary key field may have a B⁺ tree index.
 - (e) None of the above. (1)
21. Indexes speed up data access. Select the TRUE statement:
- (a) Columns which are frequently modified are good candidates for indexing.
 - (b) Primary indexes can have duplicate keys.
 - (c) Primary indexes may have null values.
 - (d) Secondary indexes must have unique keys.
 - (e) None of the above. (1)
22. Select the activity that is most clearly part of the Database Administrator's responsibility.
- (a) Data deletion.
 - (b) Data entry.
 - (c) Data modification.
 - (d) Staff management.
 - (e) Performance tuning. (1)

23. The use of a Data Dictionary produces many benefits. Select the benefit which is NOT due to proper use of a data dictionary.
- (a) consistency in data use
 - (b) improved documentation
 - (c) reduced data redundancy
 - (d) simplified performance measurement
 - (e) the enforcement of standards. (1)

24. Select the TRUE statement related to the following SQL:

```
SELECT X.a, X.b
FROM foo X, foo Y
WHERE X.b = Y.b AND Y.a='bar' ;
```

- (a) a and b are both tables.
- (b) foo is an alias for X.
- (c) the result MUST include row with 'bar' in it.
- (d) the statement is NOT valid SQL.
- (e) this is an example of a self-join. (1)

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

25. 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

- (a) The result will include unwanted rows.
- (b) The result will include unwanted columns.
- (c) The LIKE clause must include a wild card.
- (d) The result will be as intended.
- (e) Zero row will be returned. (1)

26. Select the SQL statement which results in a list of the average rainfall for each region:

- (a)

```
SELECT AVG(rainfall)
FROM station
GROUP BY stationid;
```
- (b)

```
SELECT name, AVG(rainfall)
FROM region, station
WHERE regionid = station.region
GROUP BY name;
```
- (c)

```
SELECT AVG(name) ,AVG(rainfall)
FROM region, station
WHERE regionid = station.region
GROUP BY rainfall;
```
- (d)

```
SELECT AVG(name), rainfall
FROM region, station
WHERE regionid = station.region
GROUP BY rainfall;
```
- (e)

```
SELECT name, AVG(rainfall)
FROM region, station
WHERE regionid = station.region
GROUP BY stationid;
```

(1)

27. 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 = MIN;
```
- (b)

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

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

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

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

(1)

28. Select the SQL statement which will list the station identity and rainfall for any station wetter than station 6.

(a)

```
SELECT stationid, rainfall
FROM station
WHERE rainfall >
  (SELECT stationid, rainfall
   FROM station
   WHERE stationid=6);
```

(b)

```
SELECT stationid, rainfall
FROM station
WHERE stationid=6;
```

(c)

```
SELECT b.stationid, a.rainfall
FROM station a, station b
WHERE a.rainfall < b.rainfall
AND a.stationid = 6;
```

(d)

```
SELECT a.stationid, a.rainfall
FROM station a, station b
WHERE a.rainfall < b.rainfall
AND b.stationid = 6;
```

(e)

```
SELECT a.stationid, a.rainfall
FROM station a, station b
WHERE a.rainfall > b.rainfall
AND b.stationid = 6;
```

(1)

29. Select the TRUE statement.

- (a) For referential integrity, all foreign keys should equal a primary key in another table.
- (b) For entity integrity, all foreign keys should be null.
- (c) For entity integrity, all primary keys should be null or unique.
- (d) For referential integrity, all primary keys should be non null.
- (e) For referential integrity, each foreign key should be null or equal to a primary key in another table.

(1)

30. Select the TRUE statement.

- (a) SQL is embedded within a C++ program to increase the speed of the C++ program.
- (b) SQL cursors indicate the next line of code to be executed in a C++ program.
- (c) SQL is embedded within C++ to handle sequential file processing.
- (d) SQL is embedded to make C++ programs more efficient.
- (e) SQL embedded in C++ provides facilities to extract data from a database. (1)

31. Which of the following is the highest normal form by which the Appointment relation can be classified?

Appointment (patient, consultant, hospital, address, date, time)

Given

patient, consultant → hospital, address, date, time

hospital → address

- (a) Unnormalised
- (b) 1NF
- (c) 2NF
- (d) 3NF
- (e) BCNF (1)

32. What is the highest normal form by which the following relation can be classified?

Delivery(Itemno, CustomerId, Quantity, CustomerTelno)

Given

Itemno, CustomerId → Quantity, CustomerTelno

CustomerId → CustomerTelno

- (a) Unnormalised
- (b) 1NF
- (c) 2NF
- (d) 3NF
- (e) BCNF (1)

33. Consider a table where records are stored in primary key order. Select the statement that **best** describes the cost of the operations: insert, delete and seek.

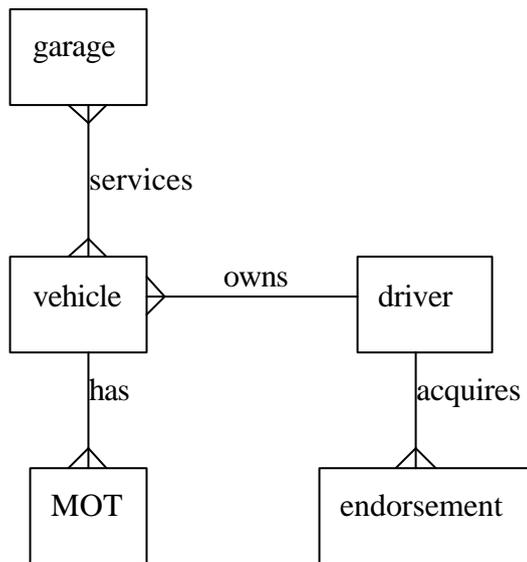
Deleted records are NOT “Flagged”.

- (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)

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

- (a) Hash key clash
- (b) Performance degradation
- (c) Loss of integrity
- (d) Transaction rollover
- (e) None of the above. (1)

The following entity relationship diagram concerns a database storing national records for vehicle and driver registration:



The “MOT” is a certificate to show that a car is mechanically safe. All vehicles are required to have regular MOT inspections at a garage.
 Cars less than 3 years old are not required to have an MOT certificate
 A driver may acquire endorsements for traffic offences such as speeding.

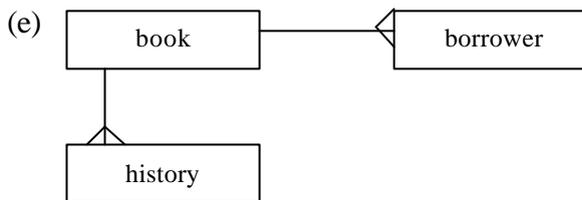
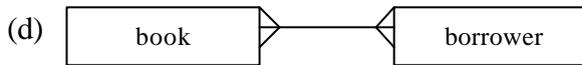
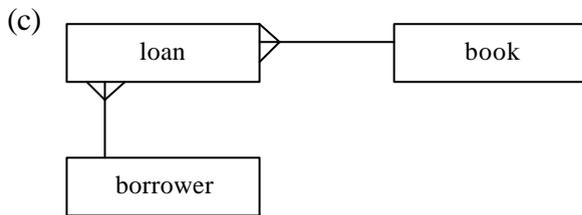
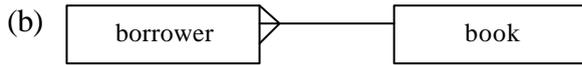
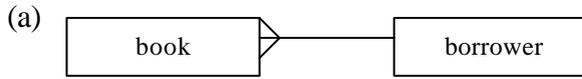
Details of all MOT certificate, including expired certificates, are included.

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 | |
|-----|-----------------------------------------------------------------------------------------|------------------------------------------------------------------|-----|
| 35. | A vehicle can have only one driver specified | The relationship is marked as many to one. | (1) |
| 36. | It is not possible to determine which cars are owned by drivers with endorsements | There is no direct relationship between endorsement and vehicle. | (1) |
| 37. | It is not possible to determine which garage has issued which certificate | The services relation is one to many. | (1) |
| 38. | It is not possible to associate MOT certificates with the relevant driver | Ownership of vehicles at a previous date is not recorded. | (1) |
| 39. | The “owns” relationship may be represented by an “owner” attribute in the vehicle table | A vehicle may have only one recorded owner. | (1) |

40. A library includes books and borrowers. At any one time a book may be borrowed by a single borrower. A record is kept of the current location of each book and its borrowing history. Select the ER diagram which **best** represents this scheme.



(1)

END OF PAPER