

Matric No: _____

**NAPIER UNIVERSITY
SCHOOL OF COMPUTING**

**CO22001 S
DATABASE SYSTEMS
(SUNWAY COLLEGE)**

ACADEMIC SESSION: 2004-2005

EXAMINATION DIET: NOVEMBER 2005

EXAM DURATION: 2 HOURS

READING TIME: NONE

EXAM PAPER INFORMATION

Answer ALL questions.

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

For full instructions see next page.

Select ONE from (a) to (e)

Number of pages – TWENTY

Number of questions – FORTY

Number of sections – ONE

EXAMINERS: John Old, Ken Chisholm, Gordon Russell, Jessie Kennedy

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. [R] [U] [S] [S] [E] [L] [L] [G]

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

Mark: (1)

2. 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 transitive dependencies
 - d. All repeating groups
 - e. None of the above

Mark: (1)

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

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

$c \rightarrow a$

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

$c, d \rightarrow e, b$

- a. $R(b, c, \underline{d}, e)$
 $R1(\underline{c}, a)$
- b. $R(b, \underline{c}, \underline{d}, e)$
 $R1(\underline{c}, a)$
- c. $R(\underline{b}, c, d, e)$
 $R1(\underline{c}, a)$
- d. all of the above
- e. none of the above

Mark: (1)

4. Which of the following is NOT an example of data redundancy?
- a. a relationship between two entity types can be derived from another relationship
 - b. two attributes in the database have the same value
 - c. a value in the database can be derived by performing a calculation on other values.
 - d. a value in the database can be derived directly from another value.
 - e. none of the above

Mark: (1)

5. For a relation to be in third normal form, which of the following is NOT a requirement?
- a. There must be no transitive functional dependencies
 - b. Every non-key attribute is fully functionally dependent on the whole key
 - c. Every determinant is a candidate key
 - d. It must be in second normal form
 - e. more than one of the above.

Mark: (1)

6. Consider the relational schema $R(\underline{A}, \underline{B}, C, D, E)$ with non-key functional dependencies $C, D \rightarrow E$ and $B \rightarrow C$.

Select the strongest statement that can be made about the schema R

- a. R is in third normal form
- b. R is in BCNF normal form
- c. R is in first normal form
- d. R is in second normal form
- e. None of the above

Mark: (1)

7. Given a database:

```
Customer (Cust_no, Name, Address)
Order (Order_no, Cust_no, C_Date, Completed)
Make (Order_no, Maker_no, Dress_style, Colour)
```

Identify the SQL command which will return the customer's name and address for all orders that have been made or are being made from Red coloured materials.

- a. SELECT Name, Address
FROM Customer, Order, Make
WHERE Colour = 'Red'
AND Customer.Cust_no = Order.Cust_no
AND Completed in ('Y', 'N')
- b. SELECT Name, Address
FROM Customer, Order, Make
WHERE Customer.Cust_no = Order.Cust_no
AND Colour = 'Red'
AND Order.Order_no = Make.Order_no
- c. SELECT Name, Address
FROM Customer, Order, Make
WHERE Colour = 'Red'
AND Customer.Cust_no = Order.Cust_no
- d. SELECT Name, Address
FROM Customer, Make
WHERE Colour = 'Red'
- e. None of the above

Mark: (1)

8.

Departments		Employees		WorkFor	
DepNo	Depname	Empno	Empname	Empno	Depno
1	Computng	1	Gordon	1	1
2	Electrical	1	Ken	3	2
3	Geography	1	Brian	4	1
4	History	1	Colin	3	3
5	Business	1	George	1	2
				2	5

Using the tables shown, which of the following SQL queries gives a list of departments and their employees?

- ```
SELECT depname,empname
FROM departments, employees, workfor
WHERE departments.depno = workfor.depno
 AND workfor.empno = employees.empno
;
```
- ```
SELECT depname,empname
FROM departments, employees
WHERE departments.depno = employees.empno
;
```
- ```
SELECT depname,empname
FROM departments, workfor
WHERE departments.depno = workfor.depno
;
```
- ```
SELECT depno,empno
FROM workfor
;
```
- None of the above.

Mark: (1)

9. `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)

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

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

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

Mark: (1)

11. Select the activity that is most clearly part of the Database Administrator's responsibility.

- a. Data modification.
- b. Data entry.
- c. Staff management.
- d. Performance tuning.
- e. Data deletion.

Mark: (1)

12. The Data Dictionary of a DBMS can be used for a variety of tasks. Which of the following is NOT ONE OF THESE?

- a. Enforcement of standards
- b. Improved Documentation
- c. Invoicing
- d. Data analysis
- e. Costing Change

Mark: (1)



13. 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 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 two attributes "first" and "last" each of which refer to Time-Slots
- d. Each event has an attribute "start" which refers to Time-Slots and "duration" which gives the length of the event in minutes
- e. Each event has an attribute "start" which refers to Time-Slots and "duration" which gives the number of slots spanned

Mark: (1)

14. In relational database evolution, "Conceptual Design" is the stage where we map

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

Mark: (1)

15. 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 and B should be kept separate with the foreign key in the B relation.
- c. B should be subsumed by A
- d. A should be subsumed by B
- e. A and B should be kept separate with the foreign key in the A relation.

Mark: (1)

16. 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)

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

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


Mark: (1)

18. At the Physical design stage, select the TRUE statement.

- a. ER diagrams are mapped into relationships
- b. ER diagrams are mapped into tables.
- c. Indices are identified and implemented for tables
- d. ER diagrams are mapped into relations
- e. Relationships are mapped into tables

Mark: (1)

19. Given the following portion of an ER diagram with a 1:n relationship which of the following is the correct rule for mapping it in to a relational schema?



- a. Take the primary key from person and add it to the car relation as a foreign key.
- b. Take the primary key from car and add it to the person relation as a foreign key.
- c. Subsume one of the relations into the other.
- d. All of the above will work.
- e. Take the primary key from both car and person and put them into a new relation called "owns".

Mark: (1)


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

- a. partial differentiation
- b. no partial precipitation
- c. no optionality
- d. partial participation
- e. partial overhangs

Mark: (1)

## 21. ER Scenario

The scenario described here is that of a book library. Books in the library can be borrowed by a borrower, and a complete history of all the books a borrower has borrowed is held in the BorrowHistory entity set. All books must have an author.



The attributes of each entity set are listed below:

Author (name, country)  
Book (title, publisher)  
BorrowHistory (when-borrowed, when-due-back)  
Borrower (name, address, date-of-birth)

Considering the information in ER Scenario, which of the following statements is TRUE?

- a. A borrower can only borrow the same book once
- b. A book can only have a single author
- c. A borrower must borrow multiple books.
- d. A book can only be borrowed once
- e. Many-to-many relationships should never appear in an ER diagram

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

Mark: (1)

23. Transactions are described as supporting the ACID model. What does the "C" stand for in ACID?

- a. Concurrency control
- b. Correctness
- c. Consistency preservation
- d. Completeness
- e. Computation

Mark: (1)

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

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

Mark: (1)

25. Transactions are often referred to in terms of ACID. Which one of the following is not part of the term ACID.

- a. Atomic
- b. Complete
- c. Durable
- d. Isolation
- e. More than one of the above

Mark: (1)

26. In relational algebra, which of the options given is the relational algebra that is equivalent to the SQL shown below.

```
SELECT a
FROM x, y
WHERE x.a = y.b
```

- a.  $\text{PROJECT}_a (x \text{ SELECT}_{x.a=y.b} y)$
- b.  $\text{SELECT}_a (x \text{ JOIN}_{x.a=y.b} y)$
- c.  $\text{PROJECT}_a (x \text{ JOIN}_{x.a=y.b} y)$
- d.  $\text{SELECT}_a (x \text{ PROJECT}_{x.a=y.b} y)$
- e. None of the above

Mark: (1)

27. In relational algebra, the  $\rho$  operator represents

- a. PROJECT
- b. UNION
- c. SELECT
- d. RENAME
- e. JOIN

Mark: (1)

28. Which of the following is TRUE when considering hash tables.

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

Mark: (1)

29. Using SERIAL data organisation -

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

Mark: (1)

30. 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 expensive, seek is cheap.
- b. Insert is expensive, delete and seek are cheap.
- c. Insert and delete are cheap, seek is expensive.
- d. Insert, delete and seek are all expensive.
- e. Insert, delete and seek are all cheap.

Mark: (1)

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

Mark: (1)

32. Select the TRUE statement.

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

Mark: (1)

33. In a database with a strict locking mechanism deadlock can occur. Select the TRUE statement.

- a. Deadlock can be avoided using a Precedence Graph
- b. Deadlock can be avoided using Two-Phase Locking
- c. Deadlock can be avoided using better primary keys
- d. Deadlock can be avoided using Timestamping
- e. None of the above

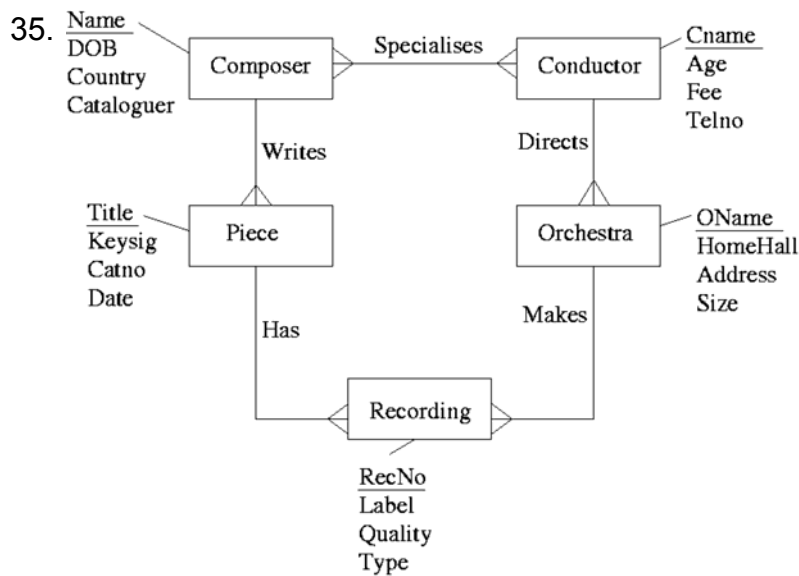
Mark: (1)



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

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

Mark: (1)



After mapping the above ERD to a relational schema which of the following set of relations would be obtained?

- Composer(Name, DOB, Country, Cataloguer, *Cname*)  
 Conductor(Cname, Age, Fee, TelNo, *Name*)  
 Piece(Title, KeySig, CatNo, Date, *Name*)  
 Orchestra(OName, HomeHall, Adress, Size, *Cname*)  
 Recording(RecNo, Label, Quality, Type, *Title*, *OName*)
- Composer(Name, DOB, Country, Cataloguer)  
 Conductor(Cname, Age, Fee, TelNo, *OName*)  
 Piece(Title, KeySig, CatNo, Date, *Name*)  
 Orchestra(OName, HomeHall, Adress, Size)  
 Recording(RecNo, Label, Quality, Type, *Title*, *OName*)  
 Specialises(Name, Cname)
- Composer(Name, DOB, Country, Cataloguer)  
 Conductor(Cname, Age, Fee, TelNo)  
 Piece(Title, KeySig, CatNo, Date, *Name*, *RecNo*)  
 Orchestra(OName, HomeHall, Adress, Size, *Cname*, *RecNo*)  
 Recording(RecNo, Label, Quality, Type)  
 Specialises(Name, Cname)
- Composer(Name, DOB, Country, Cataloguer)  
 Conductor(Cname, Age, Fee, TelNo)  
 Piece(Title, KeySig, CatNo, Date, *Name*)  
 Orchestra(OName, HomeHall, Adress, Size, *Cname*)  
 Recording(RecNo, Label, Quality, Type, *Title*, *OName*)  
 Specialises(Name, Cname)
- None of the above.

Mark: (1)

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

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

Mark: (1)

37. 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. 2NF
- b. 3NF
- c. 1NF
- d. BCNF
- e. Unnormalised

Mark: (1)

38. Select the TRUE statement.

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

Mark: (1)

39. What is the main advantage of immediate update.

- a. Changes can be stored on the disk before a commit.
- b. The locking strategy used can be timestamp based.
- c. Transactions that are long are given priority over short transactions.
- d. All transaction data is held in memory making aborts faster.
- e. Transactions that are short are given priority over long transactions.

Mark: (1)

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

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

Mark: (1)

Total Marks [40]

**End of Paper**