

Matric No: \_\_\_\_\_

**NAPIER UNIVERSITY**  
**SCHOOL OF COMPUTING**

**MODULE NO: CO22001 J**  
**DATABASE SYSTEMS (JAMES WATT COLLEGE)**

ACADEMIC SESSION: 2004-05

EXAMINATION DIET: JANUARY

TRIMESTER: ONE

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-TWO

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

2. Select the TRUE statement which would indicate data in the database as "redundant".
  - a. It is unique in the database
  - b. A VIEW has the same data as a TABLE.
  - c. The data has not yet been COMMITTED to the database
  - d. It can be derived from other data in the database.
  - e. Secondary keys are not unique

Mark: (1)

3. Consider the relational schema  $R(\underline{A}, B, C, D, E)$  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 second normal form
- b. R is in BCNF normal form
- c. R is in first normal form
- d. R is in third normal form
- e. None of the above

Mark: (1)

4. 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)

5. Given the following relation and dependencies, state which normal form the relation is in.

```
R(p,q, r, s, t)
p,q -> r,s,t
r,s -> p,q,t
t -> s
```

- a. Unnormalised
- b. First normal form
- c. Second normal form
- d. Third normal form
- e. BCNF

Mark: (1)

6. 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

What type of functional dependency describes FD4?

- A transitive dependency
- A P.K.D.
- A virtual dependency
- A whole-key dependency
- 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 most likely return the date of all orders that have completed and are made of Red or Blue coloured materials.

- a. SELECT Date  
FROM order, Make  
WHERE completed = 'Y'  
AND (colour = 'Red' OR colour = 'Blue')
- b. SELECT C\_Date  
FROM order, Make  
WHERE Order.Order\_no = Make.Order\_no  
AND completed = 'Y'  
AND colour = 'Red'  
OR colour = 'Blue'
- c. SELECT C\_Date  
FROM order, Make  
WHERE Order.Order\_no = Make.Order\_no  
AND completed = 'Y'  
AND colour = 'Red' OR 'Blue'
- d. SELECT Date  
FROM order, Make  
WHERE completed = 'Y'  
AND colour = 'Red'  
OR colour = 'Blue'
- e. SELECT Date  
FROM order, Make  
WHERE Order.Order\_no = Make.Order\_no  
AND completed = 'Y'  
AND (colour = 'Red'  
OR colour = 'Blue')

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. 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 4 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 B:

```
SELECT AVG(value)
  FROM pressure
 WHERE theDate BETWEEN '2 Feb 2003' AND '5 Feb 2003'
    AND vat = 1
```

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:

- The null value propogates and a row with the value null is returned
- The 15 correct values are added - but then erroneously divided by 16
- The 15 correct values are added - and then correctly divided by 15
- The null value propogates and zero rows are returned
- All values for 3 Feb are discarded, the remaining 12 values are summed and divided by 16

Mark: (1)



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

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

Mark: (1)

11. The role of a DBA includes which of the following topics

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

Mark: (1)

12. Which of the following is the most important advantage of an "enterprise-wide" relational database management system compared against an ad-hoc arrangement?

- a. Network traffic is reduced.
- b. Access to data will be faster.
- c. Disk space is used more efficiently.
- d. Inconsistencies are avoided.
- e. Confidentiality is assured.


Mark: (1)

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

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

Mark: (1)

14. 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)

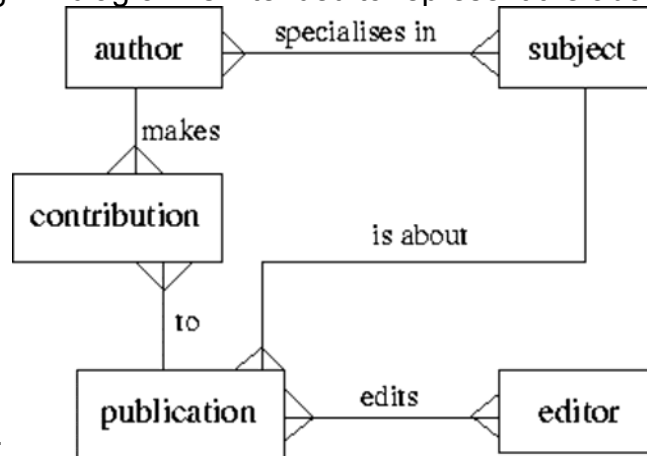
15. In the Database Analysis Life Cycle, what are the last two steps in the cycle before it begins to repeat?

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

Mark: (1)

16. 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:

- to
- specialises in
- is about
- makes
- None of the above

Mark: (1)

17. What is the main problem with the following portion of an ER diagram?




- It has a chasm trap
- It has a fan trap
- It should not have a optional relationship
- It should have a many to many (m:n) relationship
- None of the above

Mark: (1)

18. 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)

ER Scenario 1 is a good example of:

- a. A Chasm Trap
- b. A Fan Trap
- c. A Tsunami
- d. A Cliff Trap
- e. None of the above

Mark: (1)

19. When following the Database Analysis Life Cycle

- a. Evaluation is directly after Loading.
- b. Operation follows directly after Implementation.
- c. User training is performed as part of evolution.
- d. You test the database before loading in the user data
- e. You first design the database and then write the specification

Mark: (1)

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

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

Mark: (1)

21. In the ANSI/SPARC three level database model, the external view is best described by which one of the following options:

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

Mark: (1)

22. Which one of the following problems can occur due to introducing locks in a concurrent transaction scenario?

- a. Deadlock
- b. Loss of information
- c. Lack of integrity
- d. Information overwrite
- e. None of the above

Mark: (1)

23. Deadlock occurs when

- a. the state of a lock in the DBMS changes from live to dead.
- b. a transaction cannot decide what to do next.
- c. the user requests a cascade abort
- d. when transactions compete for the same resource.
- e. the DBMS cannot decide what to do next.

Mark: (1)

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

Mark: (1)

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

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

Mark: (1)

26.  $\pi$  is the projection operator.  $\sigma$  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.  $\sigma_{a < c} (\pi_{a, c} R)$
- b.  $\pi_{a < c} (\pi_{a, c} R)$
- c.  $\pi_{a < 2} R$
- d.  $\sigma_{a, c} R$
- e.  $\pi_{a, c} (\sigma_{a=c} R)$

Mark: (1)

27. A database includes two relations S and P

S

Matric_No	F_Name	L_Name	Prog_Code
04009991	Alicia	Smith	0001
04009992	Alan	Smith	0002
04009993	John	Bush	NULL

P

Prog_Code	P_Name
0001	Computing
0002	Soft. Eng.

The result of the cross product of S and P is:

a.

Matric_No	F_Name	L_Name	Prog_Code	P_Name
04009991	Alicia	Smith	0001	Computing
04009992	Alan	Smith	0002	Soft. Eng.
04009993	John	Bush	NULL	NULL

b.

Matric_No	F_Name	L_Name	Prog_Code	Prog_Code	P_Name
04009991	Alicia	Smith	0001	0001	Computing
04009991	Alicia	Smith	0001	0002	Soft. Eng.
04009992	Alan	Smith	0002	0001	Computing
04009992	Alan	Smith	0002	0002	Soft. Eng.
04009993	John	Bush	NULL	NULL	NULL
04009993	John	Bush	NULL	NULL	NULL

c.

Matric_No	F_Name	L_Name	Prog_Code	P_Name
04009991	Alicia	Smith	0001	Computing
04009992	Alan	Smith	0002	Soft. Eng.

d.

Matric_No	F_Name	L_Name	Prog_Code	Prog_Code	P_Name
04009991	Alicia	Smith	0001	0001	Computing
04009991	Alicia	Smith	0001	0002	Soft. Eng.
04009992	Alan	Smith	0002	0001	Computing
04009992	Alan	Smith	0002	0002	Soft. Eng.
04009993	John	Bush	NULL	0001	Computing
04009993	John	Bush	NULL	0002	Soft. Eng.

e. None of the above

Mark: (1)



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

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

Mark: (1)

29. Which of the following is always TRUE about indexing.

- a. You cannot query tables without building an index.
- b. indexing always speeds up table access
- c. secondary indexing must also be based on unique secondary keys.
- d. Primary indexing can handle duplicate keys.
- e. Widely varying fields are good indexing candidates.

Mark: (1)

30. Indexes speed up data access. Select the TRUE statement:

- a. Primary indexes can have duplicate keys.
- b. Primary indexes may have null values.
- c. Columns which are frequently modified are good candidates for indexing.
- d. Secondary indexes must have unique keys.
- e. None of the above.

Mark: (1)

31. The following is taken from the documentation of the PHP scripting language. It refers to a "result identifier" which identifies a cursor into a MySQL query. MySQL is a relational database which is independent of PHP.

**mysql\_fetch\_row**

mysql\_fetch\_row -- Get a result row as an enumerated array

**Description**

array mysql\_fetch\_row (int result)

Returns: An array that corresponds to the fetched row, or false if there are no more rows. Mysql\_fetch\_row() fetches one row of data from the result associated with the specified result identifier. The row is returned as an array. Each result column is stored in an array offset, starting at offset 0.

Subsequent call to mysql\_fetch\_row() would return the next row in the result set, or false if there are no more rows.

Select the true statement:

- a. The number of remaining rows may always be determined by a call to mysql\_fetch\_row.
- b. If the result identifier refers to an empty result then mysql\_fetch\_row will cause an error
- c. A single call to mysql\_fetch\_row returns a column
- d. mysql\_fetch\_row can be used to create a new table
- e. A call to mysql\_fetch\_row "advances" the cursor.

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

Mark: (1)

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

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

Mark: (1)

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

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

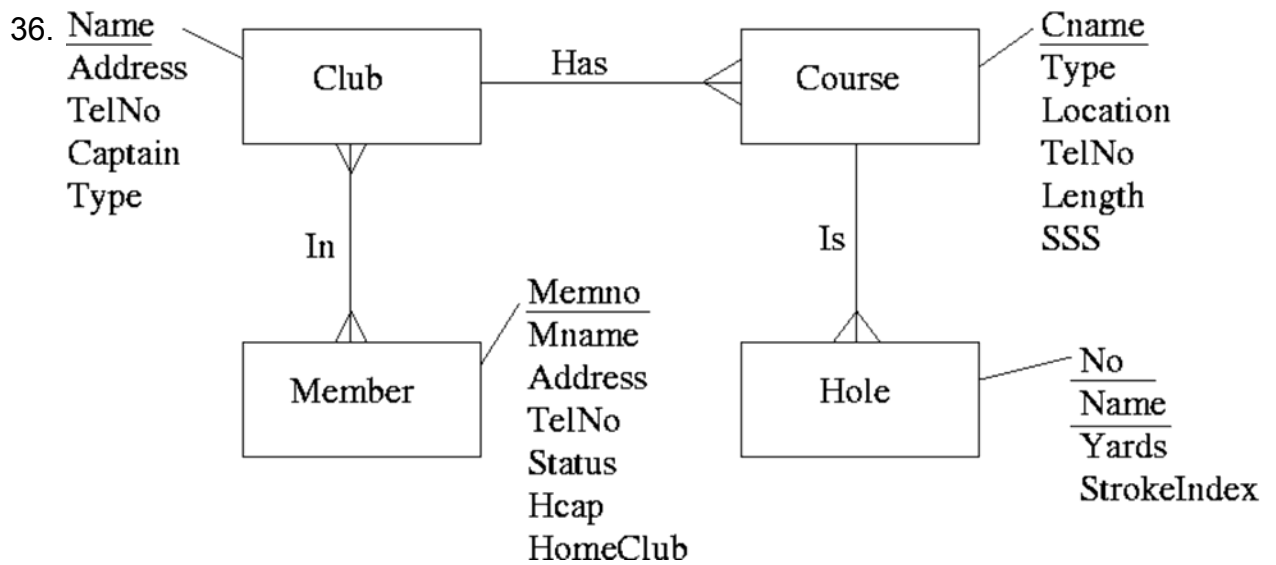
Mark: (1)

35. 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. 1NF
- b. 3NF
- c. BCNF
- d. 2NF
- e. Unnormalised

Mark: (1)



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

- Club(Name, Address, TelNo, Captain, Type)  
 Course(Cname, Type, Location, TelNo, Length, SSS, Club)  
 Member(MemNo, Mname, Address, TelNo, Status, Hcap, HomeClub)  
 Hole(No, Name, Yards, StrokeIndex)  
 In(MemNo, Name)
- Club(Name, Address, TelNo, Captain, Type)  
 Course(Cname, Type, Location, TelNo, Length, SSS, Club)  
 Member(MemNo, Mname, Address, TelNo, Status, Hcap, HomeClub)  
 Hole(No, Name, Yards, StrokeIndex, Cname)
- Club(Name, Address, TelNo, Captain, Type)  
 Course(Cname, Type, Location, TelNo, Length, SSS, Club)  
 Member(MemNo, Mname, Address, TelNo, Status, Hcap, HomeClub)  
 Hole(No, Name, Yards, StrokeIndex, Cname)  
 In(MemNo, Name)
- Club(Name, Address, TelNo, Captain, Type)  
 Course(Cname, Type, Location, TelNo, Length, SSS)  
 Member(MemNo, Mname, Address, TelNo, Status, Hcap, HomeClub)  
 Hole(No, Name, Yards, StrokeIndex, Cname)  
 In(MemNo, Name)
- None of the above.

Mark: (1)

37. 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.

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

Mark: (1)

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

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

Mark: (1)

39. Aborting a transaction

- a. Removes changes made in a transaction after it has committed.
- b. Is only possible in Microsoft Access
- c. Results in deadlock
- d. Deletes the database for security reasons
- e. Removes changes made so far in the current transaction.

Mark: (1)

40. A database can be left in an inconsistent state due to

- a. Transactions being aborted.
- b. Inaccurate data is entered into the database.
- c. Roll-forward after a failure.
- d. Deadlock
- e. A transaction fails and its changes are applied to the database.

Mark: (1)

Total Marks [40]

**End of Paper**