

Matric No: _____

**NAPIER UNIVERSITY
SCHOOL OF COMPUTING**

**CO22001 IK
DATABASE SYSTEMS
(IKIP COLLEGE)**

ACADEMIC SESSION: 2004-2005

EXAMINATION DIET: MAY 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-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 can never contain a null value for a foreign key attribute.
 - b. a foreign key attribute in a record always refers to another record which does not contain nulls
 - c. a foreign key attribute in a record always refers to another record which contains nulls
 - d. a non-null foreign key attribute always refers to another record
 - e. a record is always referred to from another record

Mark: (1)

2. 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. A virtual dependency
- b. A transitive dependency
- c. A whole-key dependency
- d. A P.K.D.
- e. None of the above

Mark: (1)

3. A lack of normalisation can lead to which one of the following problems :

- a. Deadlock
- b. Insertion problems
- c. Deletion of data
- d. Lost Updates
- e. Deferred updates

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

Mark: (1)

5. Consider the following functional dependencies:

a, b -> c, d
e -> c
b -> e, f

Given the same functional dependencies as shown above, which option shows the relations normalised to third normal form of: R(a, b, c, d, e, f)

- a. R(a, b, c, d)
R(c, e)
R(e, f, b)
- b. R(a, b, d)
R(e, c)
R(b, e, f)
- c. R(a, b, c, d)
R(c, e)
R(b, e, f)
- d. R(a, b, c, d, e, f)
- e. R(a, b, c, d, e, f)
R(e, c)
R(b, e, f)

Mark: (1)

6. 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)

7. Given the following relational database schema:

```
Customer (custno, name, address, telno, credit_rating)
Order (orderno, custno, data, delivery_date, total)
```

The following query is not well-formed SQL. Which of the following describes the main problem with the query.

```
select orderno, custno, date
from order
where total > 100
union
select * from order
order by delivery_date
```

- a. The join condition is missing
- b. Incompatible condition clauses
- c. The selects cannot be unioned
- d. Keywords are in lower case
- e. The layout is faulty

Mark: (1)

8. Continuing from the previous question

The following query is not well-formed SQL. Which of the following describes the main problem with the query.

```
select custno, name
from order, customer
where customer.custno = order.orderno and
delivery_date = '01-Jan-2001';
```

- a. The join condition is wrong
- b. The tables are in the wrong order
- c. The keywords are in lower case
- d. The date is incorrectly formatted
- e. 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. 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)

10. 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 name = 'Marilyn Monroe'
;
```

b. 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
;
```

c. The following SQL...

```
SELECT title FROM movie,casting,actor
WHERE movieid = movie.id
AND actor.id = actorid
AND name = 'Marilyn Monroe'
AND yr = 1959
;
```

d. The following SQL...

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

e. None of the above

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

16. Given the following relation select which of the ER diagrams could describe the relation.

R1 (a, b, c, d) d is a foreign key
 R2 (d, x, y)

a.



b.



c.



- a. Diagram a.
- b. Diagram b.
- c. Diagram c.
- d. two of the above
- e. All three of a,b, and c.

Mark: (1)

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

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

Mark: (1)


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

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

Mark: (1)

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

Mark: (1)

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

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

Mark: (1)

21. 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. Each event has two attributes "first" and "last" each of which refer to Time-Slots
- b. Each event has an attribute "start" which refers to Time-Slots and "duration" which gives the number of slots spanned
- c. A one-to-many relation between Events and Time-Slots is established
- d. Each event has an attribute "start" which refers to Time-Slots and "duration" which gives the length of the event in minutes
- e. A many-to-many relation between Events and Time-Slots is established

Mark: (1)

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

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

Mark: (1)



From Transaction Scenario 1, given the precedence graph, which of the following is TRUE?

- a. That the transaction schedule is serialisable
- b. Nothing as there is not enough information in the graph
- c. That the transaction schedule is unserialisable
- d. Nothing, as precedence graphs do not work for more than two transactions
- e. That the transaction schedule could be both serialisable and unserialisable

Mark: (1)

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

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

Mark: (1)

25. In the acronym ACID, used to describe the important properties of transactions, what does the letter A stand for?

- a. Atomic
- b. Action
- c. After
- d. Also
- e. None of the above

Mark: (1)

26. In symbolic notation, the SELECT operator is also known as

- a. CUP
- b. PI
- c. CAP
- d. SIGMA
- e. OMEGA

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 natural join 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. The option of dropping a secondary index to a table is being considered. Which of the following is most likely to be a consequence of dropping the index?

- a. More disk space may be required
- b. Certain insertions may be faster
- c. Certain seek operations may be faster.
- d. Certain foreign key relations may not be maintained
- e. Certain updates may be slower

Mark: (1)

29. With respect to the B+ tree index method, select the TRUE statement:

- a. Records are physically stored in primary key order.
- b. The index tree may become unbalanced as a result of updates.
- c. Only the primary key field may have a B+ tree index.
- d. B+ trees use a hashing algorithm.
- e. None of the above.

Mark: (1)

30. Indexes speed up data access. Which of the following are TRUE?

- a. Secondary indexes must have unique keys.
- b. Columns which are frequently modified are good candidates for indexing.
- c. An attribute which only has a limited number of possible values will still have access performance improved using an index.
- d. Primary indexes can have duplicate keys.
- e. None of the above.

Mark: (1)

31. The purpose of Embedded SQL is to allow

- a. Programs to be embedded in a database.
- b. SQL queries to be executed as part of a programming language.
- c. Programming language to be embedded in SQL
- d. Databases to be embedded in SQL
- e. None of the above

Mark: (1)

32. Which of the following **best** describes the relationship between C and SQL?

- a. SQL is incompatible with programming language C
- b. There exist mechanisms by which SQL statements can be embedded in C
- c. SQL can be executed from within C programs by means of JDBC
- d. There exist mechanisms by which C statements can be embedded in SQL programs
- e. C is compatible with the specialist database language SQL

Mark: (1)

33. Which of the following is TRUE for two-phase locking.

- a. keys are aquired in the first phase.
- b. locks can only be acquired on primary keys.
- c. locks can be acquired at any time.
- d. lock acquisition is the second phase.
- e. None of the above.

Mark: (1)

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

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

Mark: (1)

35. Two entity types A and B are related by a 1:1 relationship which is optional at both ends. In the process of implementing A and B as relations you would

- a. Combine A and B together into a single relationship.
- b. Keep them separate and put a foreign key in both A and B.
- c. Combine A and B together into a single relation.
- d. Use a primary key which is a composition of the primary keys of A and B.
- e. Keep them separate and put a foreign key in one of A or B.

Mark: (1)

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

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

Mark: (1)

37.

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 Scenario above, what is the CARDINALITY of the table "WorkFor"?

- a. 2
- b. 12
- c. 6
- d. 3
- e. none of the above.

Mark: (1)

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

Mark: (1)

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

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

Mark: (1)

40. With Immediate Update, writing to an attribute results in the DBMS

- a. Immediately performing a COMMIT.
- b. Writing the old and new attribute value to the log.
- c. Immediately writing the change to the database stored on disk.
- d. Releasing concomitant locks.
- e. Writing the new value to the log.

Mark: (1)

Total Marks [40]

End of Paper