NAPIER UNIVERSITY

SCHOOL OF COMPUTING

FIRST DIET MALASIA (SEMESTER ONE) EXAMINATION

SESSION 2001-2002

MODULE: CO22001

DATABASE SYSTEMS

DATE: **DURATION: 2 HOURS**

                                                                   START TIME:

EXAMINER(S)

G.RUSSELL
C.HASTIE

QUESTION PAPER DATA

Number of pages - FOURTEEN
Number of questions - FORTY
Number of sections - ONE

INSTRUCTION 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 THE FULL INSTRUCTIONS BEFORE COMMENCING WRITING

1. Transactions are described as supporting the ACID model. What does the "C" stand for in ACID?
   A. Concurrency control
   B. Correctness
   C. Completeness
   D. Computation
   E. Consistency preservation

   (1)

2. In the transaction schedule shown below for a system without concurrency control, what is the name of the error introduced?
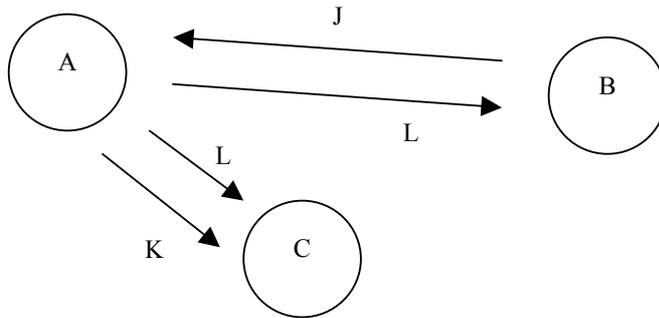
time	Transaction A	Transaction B
1		WRITE(A)
2	READ(A)	
3		ABORT
4	A = A + 3	
5	WRITE(A)	
6	COMMIT	

   A. Lost Update
   B. Inconsistency Analysis
   C. Inconsistent Update
   D. Lost Dependency
   E. Uncommitted Dependency

   (1)

Transaction Scenario 1

From a transaction schedule the following Precedence Graph was produced.

3.   From Transaction Scenario 1, the precedence graph shows
   A. that the transaction schedule is serialisable
   B. that the transaction schedule is unserialisable
   C. that the transaction schedule could be both serialisable and unserialisable
   D. nothing as there is not enough information in the graph
   E. nothing, as precedence graphs do not work for more than two transactions

(1)

4.   The graph in Transaction Scenario 1 was produced from the following schedule.

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	COMMIT	COMMIT

Compare the precedence graph to the transaction schedule, and select the TRUE statement from the following:
   A. The precedence graph is accurately drawn
   B. One or more of the loops are missing
   C. One or more of the arrows are pointing the wrong way
   D. Errors in BOTH loops and arrow direction
   E. The circles should contain the attributes, and the loops labeled with the transaction names.
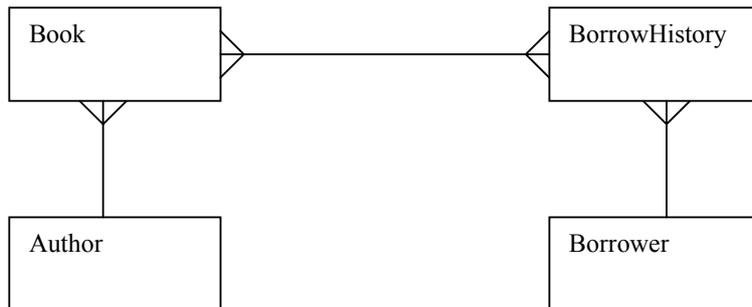
(1)

5.   Select the true statement from the following statements comparing fine and coarse grain locking granularity.
      A. fine grained locking required less processing by the DBMS
      B. coarse grained locking has less concurrency
      C. the grain size is best selected using Dykras' Algorithm
      D. the level of concurrent access is unaffected by grain size
      E. None of the above

                                                                                                    (1)

6.   In a database with a strict locking mechanism deadlock can occur. Select the TRUE statement.
      A. Deadlock can be avoided using Two-Phase Locking
      B. Deadlock can be avoided using better primary keys
      C. Deadlock can be avoided using a Precedence Graph
      D. Deadlock can be avoided using Timestamping
      E. None of the above

                                                                                                    (1)

7.   With respect to Two-Phase Locking, select the TRUE statement.
      A. Locks are only required when accessing keys
      B. Locks can be acquired at any point in a transaction
      C. If a needed lock cannot be acquired then the transactions are deadlocked
      D. Before accessing an item a lock must first be acquired
      E. None of the above

                                                                                                    (1)

8.   Using DEFERRED UPDATE, modifications made by a transaction which has not yet aborted or committed
      A. are applied to the log file and then the disk
      B. are applied only to the log file
      C. are applied only to the disk
      D. are applied in the same way as in IMMEDIATE UPDATE.
      E. None of the above

                                                                                                    (1)

9.   In IMMEDIATE UPDATE, if the DBMS fails and is then restarted then the log file is parsed. Which one of the following statements is true?
      A. Only apply the old data if the log entry ends in ABORT.
      B. Never re-apply the old data.
      C. Always re-apply the new data.
      D. If a log entry DOES NOT end in COMMIT, apply the old data.
      E. Re-apply the new data ONLY if the entry ends in COMMIT.

                                                                                                    (1)

10.   When comparing Main Memory of a computer to the data stored on a Hard Drive, select
        the TRUE statement.
            A. Main memory's transfer rate is higher than a hard drive
            B. Main memory's transfer rate is lower than a hard drive
            C. Main memory's transfer rate is about the same as a hard drive
            D. Transfer rate is dependent on the seek time
            E. Transfer rates cannot be compared.

                                                                                        (1)

11.   Using SERIAL data organisation -
            A. new records are written in primary key order
            B. new records are written in candidate key order
            C. new records are written randomly
            D. new records can only overwrite old records
            E. new records are written at the end of the PREVIOUS record

                                                                                        (1)

12.   Consider the problem of selecting multiple rows using Embedded SQL. Select the TRUE
        statement.
            A. You must use a cursor
            B. You must use a view
            C. You must use a SUBSELECT
            D. Can only be done in C++
            E. This cannot be done.

                                                                                        (1)

13.   The Data Dictionary of a DBMS can be used for a variety of tasks. Which of the
        following is NOT ONE OF THESE?
            A. Data analysis
            B. Invoicing
            C. Enforcement of standards
            D. Costing Change
            E. Improved Documentation

                                                                                        (1)

ER Scenario 1

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.

```
┌──────────────┐                              ┌──────────────────┐
│ Book         ├──────────────────────────────┤ BorrowHistory    │
│              │                              │                  │
└──────┬───────┘                              └─────────┬────────┘
       │                                                │
       │                                                │
┌──────┴───────┐                              ┌─────────┴────────┐
│ Author       │                              │ Borrower         │
│              │                              │                  │
└──────────────┘                              └──────────────────┘
```

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)
```

14. Considering the information in ER Scenario 1, which of the following statements is TRUE?
    A. A book can only be borrowed once
    B. A book can only have a single author
    C. A borrower can only borrow the same book once
    D. Many-to-many relationships should never appear in an ER diagram
    E. A borrower must borrow multiple books.

(1)

15.   Considering the information in ER Scenario 1, which of the following statements contains a good set of relations which would result from mapping the scenario into relations?

    A.   `Author(`<u>`authorID`</u>`,name,country)`
                `Book(`<u>`ISBN`</u>`,title,publisher,`*`authorID,BID`*`)`
                `BorrowHistory(`<u>`BID`</u>`,when-borrowed,when-due-back,`*`ID-NUM,ISBN`*`)`
                `Borrower(`<u>`ID-NUM`</u>`,name,address,date-of-birth)`

    B.   `Author(`<u>`name`</u>`,country)`
                `Book(`<u>`title`</u>`,publisher,`*`name`*`)`
                `BookBorrow(`<u>*`title,BID`*</u>`)`
                `BorrowHistory(`<u>`BID`</u>`,when-borrowed,when-due-back,`*`name`*`)`
                `Borrower(`<u>`name`</u>`,address,date-of-birth)`

    C.   `Author(`<u>`authorID`</u>`,name,country)`
                `Book(`<u>`ISBN`</u>`,title,publisher,`*`authorID`*`)`
                `BookBorrow(`<u>*`ISBN,BID`*</u>`)`
                `BorrowHistory(`<u>`BID`</u>`,when-borrowed,when-due-back,`*`ID-NUM`*`)`
                `Borrower(`<u>`ID-NUM`</u>`,name,address,date-of-birth)`

    D.   `Author(`<u>`name`</u>`,country)`
                `Book(`<u>`title`</u>`,publisher,`*`name,BID`*`)`
                `BorrowHistory(`<u>`BID`</u>`,when-borrowed,when-due-back,`*`name,title`*`)`
                `Borrower(`<u>`name`</u>`,address,date-of-birth)`

    E.   `Author(`<u>`authorID`</u>`,name,country,`*`ISBN`*`)`
                `Book(`<u>`ISBN`</u>`,title,publisher)`
                `BookBorrow(`<u>*`ISBN,BID`*</u>`)`
                `BorrowHistory(`<u>`BID`</u>`,when-borrowed,when-due-back)`
                `Borrower(`<u>`ID-NUM`</u>`,name,address,date-of-birth,`*`BID`*`)`

                                                      (1)

16.   ER Scenario 1 is a good example of:
    A. A Chasm Trap
    B. A Cliff Trap
    C. A Fan Trap
    D. A Tsunami
    E. None of the above

                                                      (1)

17.   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 the foreign key in the A relation.
    B. A and B should be kept separate with the foreign key in the B relation.
    C. A should be subsumed by B
    D. B should be subsumed by A
    E. A and B should be kept separate with a foreign key in both A and B.

                                                      (1)

18. Enhanced ER Models offer a few advantages over normal ER diagrams. Which of the following is one of these?
    A. Allegation
    B. Aggregation
    C. Cataloguing
    D. Capitalisation
    E. Mitigation

                                                                                          (1)

19. For a relation to be in third normal form, which of the following is NOT a requirement?
    A. It must be in second normal form
    B. There must be no transitive functional dependencies
    C. All non-key attribute is fully functionally dependent on the whole key
    D. Every determinant is a candidate key
    E. more than one of the above.

                                                                                          (1)

20. ROLLBACK may occur under a number of circumstances. Which of the following are NOT one of these?
    A. A transaction aborted.
    B. A cascade abort is required.
    C. The machine on which the DBMS was executing failed suddenly.
    D. The DBMS application was corrupted by a power spike.
    E. A committed transaction needs to be undone.

                                                                                          (1)

21. Indexes speed up data access. Which of the following are TRUE?
    A. Primary indexes can have duplicate keys.
    B. An attribute which only has a limited number of possible values will still have access performance improved using an index.
    C. Secondary indexes must have unique keys.
    D. Columns which are frequently modified are good candidates for indexing.
    E. None of the above.

                                                                                          (1)

22. In ER Modelling, "Chasm traps" can occur when entities are related via a relationship with
    A. Partial Overhangs
    B. Partial Differentiation
    C. No optionality
    D. No Partial Precipitation
    E. Partial Participation

                                                                                          (1)

23. Consider the following set of functional dependencies:

alpha, beta => gamma, delta
alpha, gamma => beta, delta
fee, fie, fo => ho, iota
fie, gamma => jota
fo, fum => iota

fee, fo, fum => fie, fot
pi, quo => rho, sigma
sigma => tau
sigma => zeta

Which of the following best describes the relation R(fee, fie, fo, fum, fiz, fot)

A. First Normal Form
B. Second Normal Form
C. Third Normal Form
D. Forth Normal Form
E. Boyce-Codd Normal Form

(1)

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

A. Lost Updates
B. Deletion of data
C. Insertion problems
D. Deferred updates
E. Deadlock

(1)

25. 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:

Name, DOB   => Handicap, HomeClub
DOB   => Category   (i.e. Junior, Ordinary, Senior or Veteran etc.)
Handicap   => BufferValue, Increment, Decrement   (The SGU Handicap system)
Yardage   => SSS   (The Standard Scratch Score – simply an
                                               indication of the difficulty of the course based on its
                                               total length in yards.)

The result of normalising the above relation in 1NF to 3NF should produce which of the following relational schemas?

A.
Member Details: Name, DOB, Handicap, HomeClub, Yardage
Categories: DOB, Category
Score System: Yardage, SSS
Handicap: Handicap, BufferValue, Increment, Decrement

B.
Member Details: Name, DOB, Handicap, HomeClub
Categories: DOB, Category, HomeClub
Score System: Yardage, SSS, Decrement
Handicap: Handicap, BufferValue, Increment, Decrement

C.
Member Details: Name, Handicap, HomeClub, Yardage, Increment, Decrement
Categories: DOB, Category
Score System: Yardage, SSS
Handicap: Handicap, BufferValue

D.
Member Details: Name, Handicap, HomeClub, Yardage, Handicap, Increment
Categories: DOB, Category
Score System: Yardage, SSS, Decrement
Handicap: Handicap, BufferValue

E. None of the above

26. In the ANSI/SPARC three level database model, the external view:
    A. Is the place where the storage structures link to the database.
    B. Is the place where the users interface to the DBMS.
    C. Is the link between users and the storage structures.
    D. is not part of the model.
    E. is dependent on the underlying DBMS product used (e.g. Oracle, DBASE).

    (1)

27. Data in a database is considered redundant if:
    A. it can be derived from other data in the database.
    B. secondary keys are not unique
    C. it is unique in the database
    D. the data has not yet been COMMITTED to the database
    E. a VIEW has the same data as a TABLE.

    (1)

28. Which of the following is a FALSE statement concerning SQL?
    A. SQL is an ISO language
    B. SQL is based on relational algebra
    C. SQL is based on a mathematical formulation
    D. SQL allows the user to perform queries on the database.
    E. SQL is a standard defined by Oracle.

    (1)

29. A foreign key:
    A. can relate to multiple rows in another table.
    B. can have a value which does not relate to a primary key.
    C. maintains a relationship between tables.
    D. can relate to columns which are not primary keys.
    E. is only used in multi-language database implementations.

    (1)

SQL Scenario 1

Departments	
DeptNo	**Depname**
1	Computing
2	Electrical
3	Geography
4	History
5	Business

Employees	
empno	empname
1	Gordon
2	Ken
3	Brian
4	Colin
5	George

WorkFor	
empno	depno
1	1
3	2
4	1
3	3
1	2
2	5

30.   Using SQL Scenario 1, what is the CARDINALITY of the table "WorkFor"?
      A. 2
      B. 3
      C. 6
      D. 12
      E. none of the above.

(1)

31.   Using SQL Scenario 1, which of the following SQL queries gives a list of departments and their employees?

A.
```
SELECT depname,empname
  FROM   departments, employees
  WHERE  departments.depno = employees.empno
  ;
```

B.
```
SELECT depname,empname
  FROM   departments, employees, workfor
  WHERE  departments.depno = workfor.depno
    AND  workfor.empno     = employees.empno
  ;
```

C.
```
SELECT depname,empname
  FROM   departments, workfor
  WHERE  departments.depno = workfor.depno
  ;
```

D.
```
SELECT depno,empno
  FROM   workfor
  ;
```

E.   None of the above.   (1)

32. Using SQL Scenario 1, which of the following SQL queries gives the number of employees in each department?

A.
```
SELECT depname,COUNT(depno)
FROM    departments, employees, workfor
WHERE   departments.depno = workfor.depno
   AND  workfor.empno    = employees.empno
GROUP BY workfor.depno
;
```

B.
```
SELECT depname,COUNT(*)
FROM    departments, workfor
WHERE   departments.depno = workfor.depno
GROUP BY workfor.empno
;
```

C.
```
SELECT depname,COUNT(workfor.empno)
FROM    departments, workfor
WHERE   departments.depno = workfor.depno
GROUP BY workfor.depno
;
```

D.
```
SELECT depname,COUNT(*)
FROM    departments, employees, workfor
WHERE   departments.depno = workfor.depno
   AND  workfor.empno    = employees.depno
GROUP BY workfor.depno
```

E. None of the above.

(1)

33. Using SQL Scenario 1, what is the CARDINALITY of the relationship between the entity type departments and the entity type workfor?
   A. 2
   B. 5
   C. 1:N
   D. N:1
   E. M:N

(1)

34. Using SQL Scenario 1, what is the primary key of the table workfor?
   A. empno
   B. depno
   C. a composite key of depno and empno
   D. it does not have a primary key
   E. the many-to-many relationship must be eliminated before it can be calculated.

(1)

35. In the Database Analysis Life Cycle, what are the last two steps in the cycle before it begins to repeat?
    A. "Operation" and then "Training"
    B. "Operation" and then "Maintenance and Evolution"
    C. "Operation" and then "Payment"
    D. "Testing and Evaluation" and then "Operation"
    E. "Implementation" and then "Training"

(1)

36. In relational database evolution, "Conceptual Design" is the stage where we map
    A. Specification into relations
    B. Specification into ER diagrams
    C. Specification into Marketing Ideas
    D. ER diagrams into relations
    E. ER diagrams into tables

(1)

37. When mapping an ER diagram into relations where two entity sets A and B exists which are linked by a relationship 1:N where the A side of the relationship is optional
    A. you should subsume A into B.
    B. you have to use create another relation first.
    C. the foreign key is placed in relation A.
    D. the foreign key is placed in relation B.
    E. the foreign key is placed in both relation A and relation B

(1)

Relation Algebra Scenario 1

R

ColA	ColB
A	1
C	4
D	6
E	6

S

ColC	ColD
C	2
D	1
G	6
J	7

38. Using Relation Algebra Scenario 1, what is the cardinality of the resulting relation of:
    R JOIN$_{R.ColA = P.ColC}$ S
    A. 0
    B. 2
    C. 6
    D. 8
    E. None of the above

(1)

39. Using Relation Algebra Scenario 1, what is the cardinality of the resulting relation of:
R FULL OUTER JOIN$_{R.ColA = P.ColC}$ S
    A. 0
    B. 2
    C. 6
    D. 8
    E. None of the above

(1)

40. In relational algebra, the ρ operator represents
    A. RENAME
    B. SELECT
    C. PROJECT
    D. JOIN
    E. UNION

(1)

Total marks [40]
END OF PAPER