

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

**FIRST DIET (SEMESTER TWO) EXAMINATION**

**SESSION 2000-2001**

**CO72010: DATABASE SYSTEMS**

**DATE:**

**DURATION: 2 HOURS**

**START TIME:**

**EXAMINER:**

**MR. A. CUMMING**

**QUESTION PAPER DATA**

**Number of Pages - ELEVEN**  
**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 entered on the answer sheet provided using a horizontal line.**  
**A pencil must be used.**  
**Append Answer Sheet to Question Paper.**

**PLEASE READ THE FULL INSTRUCTIONS BEFORE COMMENCING WRITING**

Question 1- 5 concern the following database for a recruitment agency:

job		
reference	employer	salary
01	Napier	£20000
02	GCHQ	£22000
03	Napier	£24000

requirement	
job	skill
01	Unix Admin
01	Oracle Admin
02	Unix Admin
02	Number Theory

The database currently holds three job vacancies. Two of the vacancies have specific skills requirements.

1. Choose the SQL statement which will return details of the jobs of interest to a candidate with experience of Unix Administration.
  - (a) `SELECT * FROM job  
WHERE reference='Unix Admin'  
AND skill = 'Unix Admin';`
  - (b) `SELECT * FROM job, requirement  
WHERE reference=job AND skill='Unix Admin';`
  - (c) `SELECT * FROM job  
WHERE skill = 'Unix Admin';`
  - (d) `SELECT * FROM 'Unix Admin';`
  - (e) `SELECT 'Unix Admin' FROM job, requirement;` (1)
  
2. A list of jobs together with a count of the required skills is needed. It is important that jobs such as 03, which has no skills specified, are included. Which of the following operations is most appropriate?
  - (a) CARTESIAN PRODUCT
  - (b) INNER JOIN
  - (c) INTERSECTION
  - (d) LEFT or RIGHT OUTER JOIN
  - (e) UNION (1)
  
3. Select the term which best describes the cardinality of the relationship: job to the requirement.
  - (a) many to many
  - (b) many to one
  - (c) one to many
  - (d) one to one
  - (e) some to many (1)
  
4. Which of the following show appropriate primary keys for the tables?
  - (a) job(reference, employer, salary) requirement(job, skill)
  - (b) job(reference, employer, salary) requirement(job, skill)
  - (c) job(reference, employer, salary) requirement(job, skill)
  - (d) job(reference, employer, salary) requirement(job, skill)
  - (e) job(reference, employer, salary) requirement(job, skill) (1)

5. The table requirement was created using the following SQL statement:

```
CREATE TABLE requirement (  
    job INTEGER,  
    skill VARCHAR(50),  
    FOREIGN KEY job REFERENCES job(reference));
```

Select the true statement concerning the requirements table.

- (a) it cannot store a field with a NULL value for job
  - (b) it must be created BEFORE the job table
  - (c) the value 'Windows 2000 Admin' is NOT permitted for the skill attribute
  - (d) referential integrity checks will prevent invalid values for job
  - (e) referential integrity checks will prevent identical rows being inserted (1)
6.  $\Pi$  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)  $\sigma_{a,c} R$
  - (b)  $\Pi_{a < 2} R$
  - (c)  $\Pi_{a,c}(\sigma_{a=c} R)$
  - (d)  $\sigma_{a < c}(\Pi_{a,c} R)$
  - (e)  $\Pi_{a < c}(\Pi_{a,c} R)$  (1)
7. A many to many relation is required between two tables EMPLOYEE and PROJECT to indicate that a project may have many staff and an employee may be working on more than one project. It would be unusual for an employee to be involved in more than three projects. The relationship between EMPLOYEE and PROJECT would be best represented by:
- (a) an attribute **employee** in the PROJECT table
  - (b) an attribute **project** in the EMPLOYEE table
  - (c) attributes **employee1, employee2, employee3** in the PROJECT table
  - (d) attributes **project1, project2, project3** in the employee table
  - (e) a new table with two attributes: **employee** and **project** (1)

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.

8. Select the true statement:
- (a) A call to mysql\_fetch\_row "advances" the cursor.
  - (b) A single call to mysql\_fetch\_row returns a column
  - (c) If the result identifier refers to an empty result then mysql\_fetch\_row will cause an error
  - (d) mysql\_fetch\_row can be used to create a new table
  - (e) The number of remaining rows may always be determined by a call to mysql\_fetch\_row. **(1)**
9. 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 **(1)**
10. Which of the following is NOT an example of data redundancy?
- (a) a value in the database can be derived directly from another value.
  - (b) a value in the database can be derived by performing a calculation on other values.
  - (c) a relationship between two entity types can be derived from another relationship
  - (d) two attributes in the database have the same value
  - (e) none of the above **(1)**

11. Referential Integrity is where
- (a) references in the database are identical to references in an ER diagram.
  - (b) foreign keys cannot be changed by queries.
  - (c) foreign keys in one table always refer to foreign keys in another table.
  - (d) primary keys in one table always refer to primary keys in another table.
  - (e) foreign keys always refer to a candidate key or are NULL. **(1)**
12. A primary key in a table
- (a) must be different from a primary key in another table
  - (b) must be different from foreign keys in another table
  - (c) can be the same value as other primary key values in the same table
  - (d) must not contain NULL
  - (e) must be a number **(1)**
13. When following the Database Analysis Life Cycle
- (a) You first design the database and then write the specification
  - (b) Evaluation is directly after Loading.
  - (c) You test the database before loading in the user data
  - (d) User training is performed as part of evolution.
  - (e) Operation follows directly after Implementation. **(1)**
14. 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)**
15. A fan trap in an ER diagram can occur where
- (a) an entity type has optional relationships.
  - (b) an entity type is related to other entity types only via N:1 relationships.
  - (c) an entity type is related to other entity types only via 1:N relationships.
  - (d) an entity type is related to other entity types only via M:N relationships.
  - (e) a relationship hides an attribute. **(1)**
16. In Enhanced ER diagrams, a subclass
- (a) may contain many superclasses.
  - (b) can only exist in Chen's notation.
  - (c) Is part of Chun's notation.
  - (d) may contain only one superclass.
  - (e) is contained by one superclass. **(1)**

17. Which of the following could be the result of Generalisation:
- (a) superclass - card(cardnumber, issuer, cardholder)  
     subclass - visa(expiryDate, creditLimit)  
     subclass - switch(issueDate, colour)
  - (b) card(cardnumber, issuer, cardholder, expiryDate, creditLimit, issueDate, colour)
  - (c) superclass - card(cardnumber, issuer, cardholder)  
     subclass - visa(expiryDate, creditLimit)  
     subclass - switch(expiryDate, issueDate, colour)
  - (d) subclass - card(cardnumber, issuer, cardholder)  
     superclass - visa(expiryDate, creditLimit)  
     superclass - switch(issueDate, colour)
  - (e) subclass - card(cardnumber, issuer, cardholder)  
     subclass - visa(expiryDate, creditLimit)  
     subclass - switch(issueDate, colour) (1)
18. 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 relation.
  - (b) Combine A and B together into a single relationship.
  - (c) Keep them separate and put a foreign key in both A and B.
  - (d) Keep them separate and put a foreign key in one of A or B.
  - (e) Use a primary key which is a composition of the primary keys of A and B. (1)
19. The process of combining entity types from an ER diagram into a single relation in a relational database is to
- (a) consume
  - (b) subsume
  - (c) assume
  - (d) generalise
  - (e) specialise (1)
20. Consider the following functional dependencies:
- $$\begin{array}{lcl} a,b & \rightarrow & c,d \\ e & \rightarrow & c \\ b & \rightarrow & e,f \end{array}$$
- Given the functional dependencies shown above, what normal form is the following relation in: R(a, b, c, d, e, f)
- (a) unnormalised
  - (b) first normal form
  - (c) second normal form
  - (d) third normal form
  - (e) BCNF (1)

21. Given the same functional dependencies as shown above, which option shows the relations normalised to third normal form of:  $R(\underline{a}, \underline{b}, c, d, e, f)$
- (a)  $R(\underline{a}, \underline{b}, c, d, e, f)$
  - (b)  $R(\underline{a}, \underline{b}, d)$   
 $R(\underline{e}, c)$   
 $R(\underline{b}, e, f)$
  - (c)  $R(\underline{a}, \underline{b}, c, d)$   
 $R(\underline{c}, e)$   
 $R(\underline{b}, e, f)$
  - (d)  $R(\underline{a}, \underline{b}, c, d)$   
 $R(\underline{c}, e)$   
 $R(\underline{e}, f, b)$
  - (e)  $R(\underline{a}, \underline{b}, c, d, e, f)$   
 $R(\underline{e}, c)$   
 $R(\underline{b}, e, f)$
- (1)**

22. Consider the following set of functional dependencies:
- |              |   |              |   |              |   |            |
|--------------|---|--------------|---|--------------|---|------------|
| alpha, beta  | → | gamma, delta | → | fee, fo, fum | → | fie, fot   |
| alpha, gamma | → | beta, delta  | → | pi, quo      | → | rho, sigma |
| fee, fie, fo | → | ho, iota     | → | sigma        | → | tau        |
| fie, gamma   | → | jota         | → | sigma        | → | zeta       |
| fo, fum      | → | iota         |   |              |   |            |

Given the above functional dependencies, which of the following best describes the relation  $R(\underline{fee}, \underline{fie}, \underline{fo}, \text{fum}, \text{fiz}, \text{fot})$ ?

- (a) First Normal Form
  - (b) Second Normal Form
  - (c) Third Normal Form
  - (d) Forth Normal Form
  - (e) Boyce-Codd Normal Form
- (1)**
23. 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)**

A golf club proposes to hold a database about members, instead of the current paper-based card system.

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 – simply an indication of the difficulty of the course based on its total length in yards.)

Questions 24 to 27 relate to the 3<sup>rd</sup> normal form schema 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

24. Select the appropriate description of the Member relation:

- (a) Member: Name, DOB, Handicap, HomeClub, Yardage
- (b) Member: Name, DOB, Handicap, HomeClub, Yardage
- (c) Member: Name, DOB, Handicap, HomeClub
- (d) Member: Name, DOB, Handicap, HomeClub
- (e) Member: Name, DOB, Handicap, HomeClub, Yardage (1)

25. Select the appropriate description of the Categories relation

- (a) Categories: DOB, Category
- (b) Categories: Category, DOB
- (c) Categories: DOB, Category
- (d) Categories: DOB, Category
- (e) Categories: Category, DOB1, DOB2 (1)

26. Select the appropriate description of the Handicap relation

- (a) Handicap: Handicap, BufferValue
- (b) Handicap: Handicap, BufferValue, Increment, Decrement
- (c) Handicap: BufferValue, Increment, Decrement, Handicap
- (d) Handicap: Name, DOB, Handicap
- (e) Handicap: Handicap, Name, DOB (1)



27. What type of functional dependency describes FD4?  
(a) A P.K.D.  
(b) A transitive dependency  
(c) A whole-key dependency  
(d) A virtual dependency  
(e) None of the above (1)
28. 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 (1)
29. In the shadow-paging method of implementing transactions known by the acronym SVM; what does the letter S usually stand for?  
(a) Safe  
(b) Sending  
(c) Shadow  
(d) Simple  
(e) System (1)
30. 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 1-m relationship is mapped into a new relation.  
(e) Each m-n relationship is mapped into a new relation. (1)
31. 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)
32. Assume the relation  $R(\underline{A}, \underline{B}, C, D, E)$  is in 2NF. Which of the following functional dependency must be **TRUE**?  
(a)  $A, C \rightarrow D$   
(b)  $D \rightarrow C$   
(c)  $D \rightarrow E$   
(d)  $A, B \rightarrow C$   
(e) None of the above. (1)

33. Assume the relation  $R(\underline{A}, \underline{B}, C, D, E)$  is in at least 3NF. Which of the following functional dependency must be **FALSE**?
- (a)  $A, B \rightarrow C$
  - (b)  $A, B \rightarrow D$
  - (c)  $A, C \rightarrow E$
  - (d)  $C, D \rightarrow E$
  - (e) None of the above
- (1)**
34. Select the statement which best characterises the performance of the insert, seek and sort operations when comparing a hash index against a B+ tree index.
- (a) insert and seek are similar, but sort will be slower with the hash index
  - (b) seek and sort are similar, but insert will be slower with the hash index
  - (c) sort and insert are similar, but seek will be slower with the hash index
  - (d) insert, seek and sort will all be similar
  - (e) seek will be similar, insert and sort will be slower
- (1)**
35. Consider each of the following properties of an index which uses ordered trees. Select the feature which characterises **POOR** performance.
- (a) The attributes being indexed are short
  - (b) The branching factor is high
  - (c) The index tree is balanced
  - (d) The maximum depth of the index tree is small
  - (e) The tree contains long branches
- (1)**
36. The external view of the ANSI-SPARC architecture chiefly concerns:
- (a) the way individual users see the data
  - (b) the formal description of the data
  - (c) the way the data is actually stored
  - (d) the interface to other applications
  - (e) the data that users outside the company are permitted to view
- (1)**
37. A back-up and recovery regime should protect an organisation against:
- (a) data corruption
  - (b) data validation
  - (c) inconsistent data
  - (d) incorrect data
  - (e) insecure data
- (1)**
38. A well designed relational schema will prevent:
- (a) data corruption
  - (b) data validation
  - (c) inconsistent data
  - (d) incorrect data
  - (e) insecure data
- (1)**

39. The role of the database administrator includes approval of changes to:
- (a) concurrency locks
  - (b) field values
  - (c) records
  - (d) tuples
  - (e) user accounts
- (1)**
40. The standard language SQL contains features to perform which of the following functions:
- (a) specifying repeated field names
  - (b) specifying disk geometry
  - (c) specifying user passwords
  - (d) specifying frequency of backups
  - (e) specifying user access rights
- (1)**

**Total marks [40]**

**END OF PAPER**