

No: _____

NAPIER UNIVERSITY
SCHOOL OF COMPUTING

CO22001

DATABASE SYSTEMS (LAUDER COLLEGE)

ACADEMIC SESSION: 2002-2003

EXAMINATION DIET: MAY/JUNE

TRIMESTER: TWO

EXAMINATION DURATION: 2 HOURS

READING TIME: NONE

EXAM PAPER INFORMATION

Number of pages – TWENTY-ONE

Number of questions – FORTY

Select ONE from (a) to (e)

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

For full instructions see next page.

EXAMINER: DR. G. RUSSELL

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. [C] [H] [E] [S] [N] [E] [Y] [F]

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.

Relation P		Relation Q	
ColW	ColX	ColY	ColZ
A	4	B	7
B	5	D	4
C	6	C	6

Consider the relations P and Q above. Select the number of rows in the table resulting from the following join.

P RIGHT OUTER JOIN_{ColX = ColZ} Q

- a. 2
- b. 4
- c. 7
- d. 3
- e. None of the above.

Mark: (1)

2. In Relational Algebra, a "tuple"

- a. is a collection of attributes describing some real-world entity.
- b. is an index.
- c. is a set of atomic values.
- d. is effectively a column of a relation.
- e. is a collection of relations describing a mini-world view.

Mark: (1)

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

Mark: (1)

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

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

Mark: (1)

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

Mark: (1)

6. Consider the following functional dependencies

$a, b \Rightarrow c, d$ $e, g, h \Rightarrow f, j$
 $a, c \Rightarrow b, d$ $p, q \Rightarrow r, s$
 $e, f, g \Rightarrow h, i$ $s \Rightarrow t$
 $f, g \Rightarrow j$ $q \Rightarrow u$
 $g, h \Rightarrow i$

Which of the following relational schemas might be the result of normalising $R(\underline{s}, \underline{q}, t, u)$?

- a. The schema $R(\underline{s}, \underline{q}, t, u)$
- b. The schema $R1(\underline{s}, \underline{q})$ $R2(\underline{s}, t)$ $R3(\underline{q}, u)$
- c. The schema $R1(\underline{s}, \underline{q})$ $R2(\underline{s}, t)$ $R3(\underline{q}, u)$
- d. The schema $R1(\underline{s}, \underline{q}, t)$ $R2(\underline{s}, \underline{q}, u)$
- e. The schema $R1(\underline{s}, \underline{q})$ $R2(\underline{q}, t)$ $R3(\underline{t}, u)$

Mark: (1)

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

$R(\underline{p}, q, r, s, t)$
 $p, q \rightarrow r, s, t$
 $r, s \rightarrow p, q, t$
 $t \rightarrow s$

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

Mark: (1)

8. Consider the following functional dependencies:

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

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

Mark: (1)


9. Sometimes it is best NOT to fully normalise a database. Select the strongest argument which would support this statement.
- a. A fully normalised database may perform too slowly
 - b. Normalisation will increase the number of attributes in each table
 - c. Normalisation will reduce the number of tables
 - d. An un-normalised database may make some queries too complicated
 - e. A fully normalised database may result in tables which are too large

Mark: (1)

10. When implementing security in a DBMS, which of the following is NOT supported by the GRANT command?
- a. Providing DELETE privileges.
 - b. Removing privileges of other people.
 - c. Supporting the devolution of access control to non-DBAs.
 - d. Changing passwords.
 - e. Providing SELECT privileges.

Mark: (1)

11. When producing relations for the diagram shown, where almost all P entities relate to a Q entity:



Which of the following is best?

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

Mark: (1)

12. 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

Mark: (1)

13. The DOMAIN of an attribute is which of the following?

- a. The range of values that the attribute is allowed to take.
- b. The type of the attribute.
- c. Which tables the attribute is allowed to be in.
- d. More than one of the above.
- e. None of the above.

Mark: (1)

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

Mark: (1)

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

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

Mark: (1)

16. Which of the following best describes the costs of the operations insert, delete and seek on a table where records are stored in primary key order? Deleted records may be "flagged".

- a. insert, delete and seek are all cheap
- b. insert is expensive, delete and seek are cheap.
- c. insert and delete are cheap, seek is expensive
- d. insert and delete are expensive, seek is cheap
- e. insert, delete and seek are all expensive

Mark: (1)

17.

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

Select the term which best describes the cardinality of the relationship between the table job to the table requirement.

- a. many to many
- b. one to one
- c. one to many
- d. many to one
- e. some to many

Mark: (1)

18. Continuing from the previous question

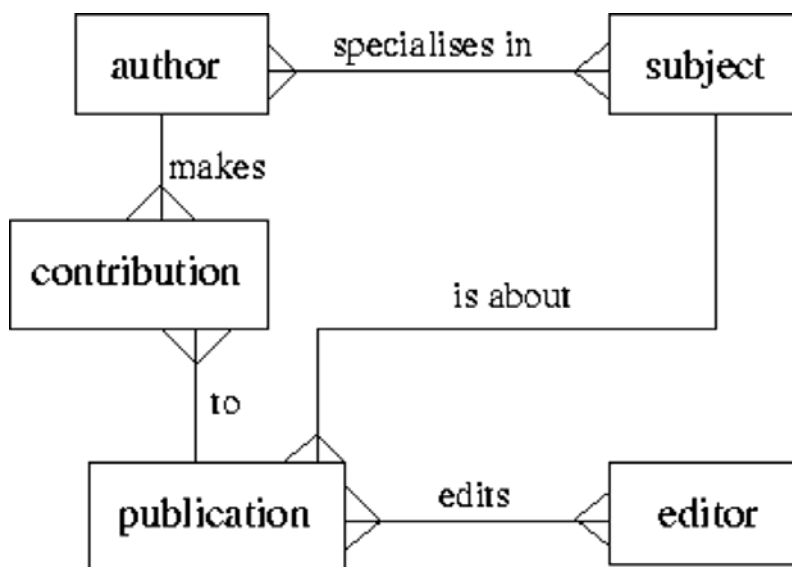
Which of the following show appropriate primary keys for the tables?

- 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)

Mark: (1)

19. 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:

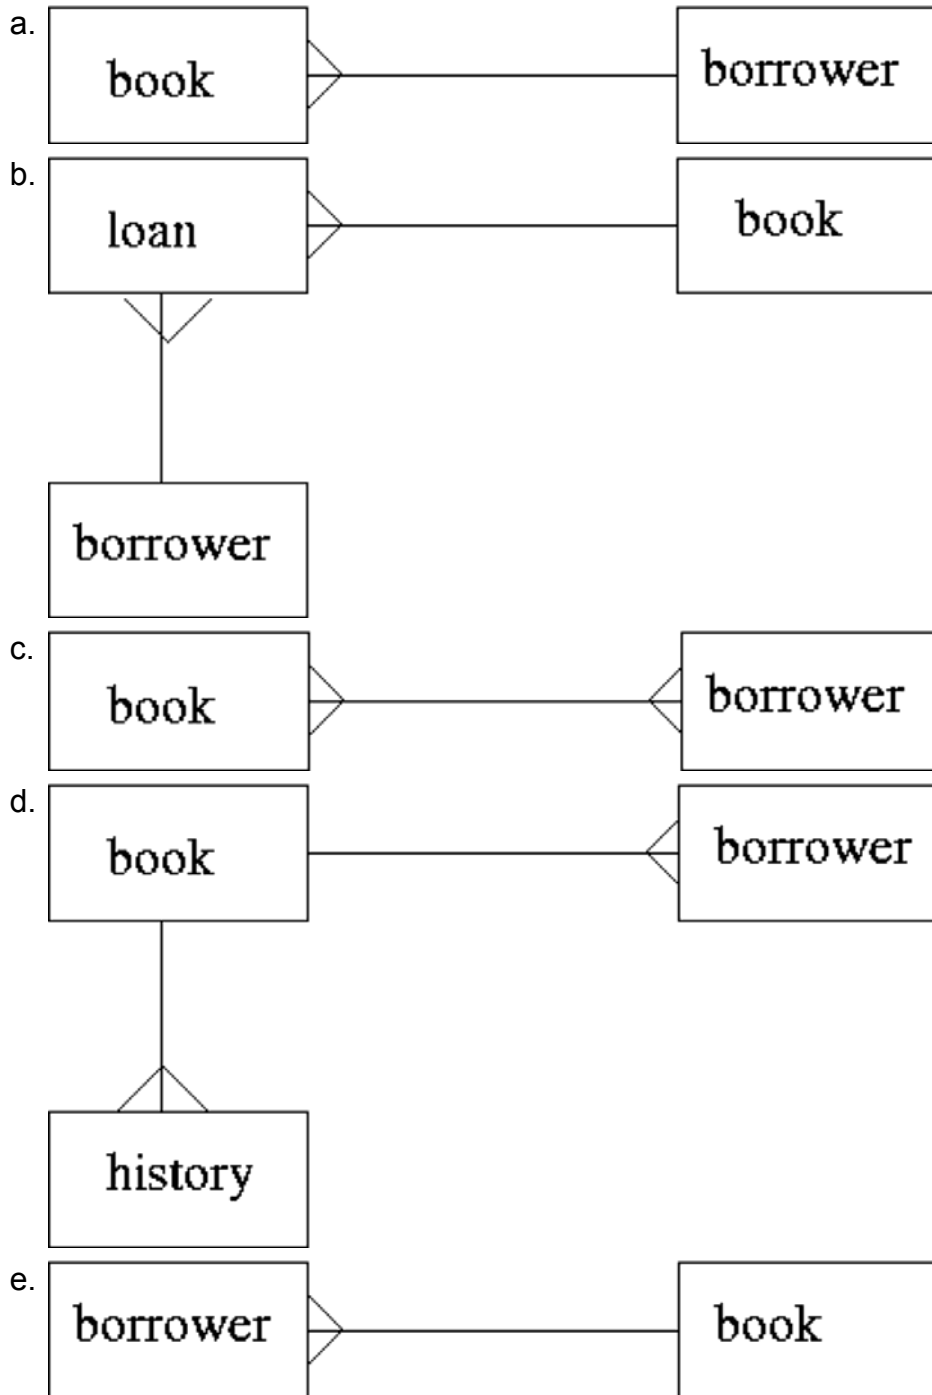


Indicate the relation which has an incorrect cardinality shown:

- a. is about
- b. makes
- c. specialises in
- d. to
- e. None of the above

Mark: (1)

20. A library includes books and borrowers. At any one time a book may be borrowed by a single borrower. A record is kept of the current location of each book and its borrowing history. Select the ER diagram which **best** represents this scheme.



Mark: (1)


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

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

Mark: (1)

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

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

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

Mark: (1)

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

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

Mark: (1)



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

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

Mark: (1)

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

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

Mark: (1)

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

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

Mark: (1)

30. An athletics meeting involves several competitors who participate in a number of events. The database is intended to record who is to take part in which event and to record the outcome of each event. As results become available the winner attribute will be updated with the cid of the appropriate competitor.

Competitor(cid, name, nationality)

Event(eid, description, winner)

Competes(cid, eid)

Competitor			Event			Competes	
cid	name	nationality	eid	description	winner	cid	eid
01	Pat	British	01	running		01	01
02	Hilary	British	02	jumping		02	01
03	Sven	Swedish	03	throwing		03	02
04	Pierre	French				04	02
						04	03

The actual database is to contain several million competitors and nearly one hundred thousand events. The database must be indexed for producing fast results for two particular queries:

- (X) A list of the names and nationalities of the competitors for a given event where the event description is given.
- (Y) A list of event descriptions for a given competitor where the cid is given.

Which of the following is not required?

- An index on name in Competitors
- An index on cid in Competes
- An index on eid in Competes
- An index on description in Events
- An index on eid in Event

Mark: (1)

31. With respect to a DBMS using a hard drive as secondary storage. Select the TRUE statement.

- a. Data should be stored in contiguous blocks to maintain consistency.
- b. A block is the smallest unit which can be read from the disk by the operating system.
- c. Data can be read from a file in "attribute-size" amounts.
- d. The DBMS manages the disk blocks.
- e. The database should be backed up onto primary memory at regular intervals.

Mark: (1)

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

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

Mark: (1)

33. With respect to Two-Phase Locking. Select the FALSE statement.

- a. Locks may be made at any point in the transaction schedule.
- b. A transaction may start to release locks at any time.
- c. The phases are "lock-acquisition" and "lock-release".
- d. Locks are released on an ABORT.
- e. Two-Phase Locking requires key pairs.

Mark: (1)

34. With respect to Two-Phase Locking, select the TRUE statement.

- a. Locks can be acquired at any point in a transaction
- b. Before accessing an item a lock must first be acquired
- c. If a needed lock cannot be acquired then the transactions are deadlocked
- d. Locks are only required when accessing keys
- e. None of the above

Mark: (1)

35. Which of the following is a FALSE statement concerning SQL?

- a. SQL is a standard defined by Oracle.
- b. SQL is an ISO language
- c. SQL is based on a mathematical formulation
- d. SQL allows the user to perform queries on the database.
- e. SQL is based on relational algebra

Mark: (1)

36. Which of the following is TRUE about a foreign key?

- a. It is only used in multi-language database implementations.
- b. It can relate to columns which are not primary keys.
- c. It can relate to multiple rows in another table.
- d. It maintains a relationship between tables.
- e. It can have a value which does not relate to a primary key.

Mark: (1)

37. The standard language SQL contains features to perform which of the following functions:

- a. specifying user passwords
- b. specifying user access rights
- c. specifying frequency of backups
- d. detecting redundant data
- e. specifying disk geometry

Mark: (1)

38.

Departments		Employees		WorkFor	
DeptNo	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?

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

Mark: (1)

39. What is the main advantage of immediate update.

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

Mark: (1)

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

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

Mark: (1)

Total Marks: 40

END OF PAPER