香港公開大學
THE OPEN UNIVERSITY OF HONG KONG

The Open University

Higher Level Course Examination (Specimen Examination Paper)
M358 Relational Databases

Time Allowed: 3 hours

Examination Number								
Student Number								

There are **TWO** parts to this paper. You should answer **BOTH**.

Part A (50%)

You should attempt **ALL** questions. Your answers to Part A questions should be written *in the spaces provided inside this examination paper*.

Part B (50%)

You should attempt not more than **TWO** questions. Each Part B question carries 25 marks.
Your answers to Part B questions should be written in the answer book(s) provided.

Instructions:

1. This examination paper should be answered in ENGLISH.
2. Begin each question on a new page.
3. Write the question number at the top of each page you have worked on.
4. Do all your rough work in the answer book but remember to cross it out afterwards. You may ask the invigilator for supplementary answer books if you need them.
5. You should write clearly. Marks may be deducted where the writing is very difficult to read.
6. At the end of the examination, check that you have written your examination number and student number on the cover of this question paper and each answer book used. Also check that you have written the course code and completed the examination number column on the answer book flap. **Failure to do so will mean that your work cannot be identified**.
7. Fasten together your answer book, and supplementary answer book(s) if any with the treasury tag provided. **Both** the question paper and the answer book(s) must be handed in to the invigilator.
8. **Do NOT open this question paper until you are told to do so, otherwise you may be disqualified.**

Admissible materials in this examination:

1. English language and/or English-Chinese language dictionary without additional hand-written or printed notes is allowed. Electronic dictionary is NOT allowed.

PART I

Answer ALL questions from this part.

Question 1

(a) Give a brief explanation of the difference between the terms *information* and *data*.

[2]

(b) For the following example form:

Department Name	Customer Support Services
Department Code	CSS11
Manager	John Sykes
Staff Employed	17

(i) Give a sentence that represents your understanding of its information content.

[1]

(ii) The value for Staff Employed is *derived data*. Explain what this means in terms of data in a database.

[2]

Question 2

Consider the diagram in Figure 1, which shows occurrences of the Supplies relationship connecting occurrences of the entity types Supplier and Part.

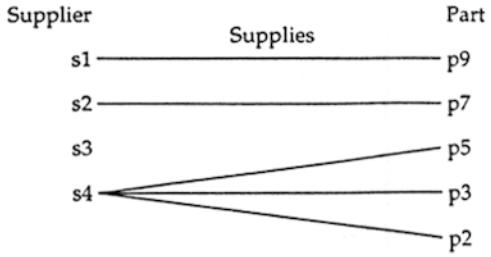

Figure 1

(a) State the number of occurrences of the Supplies relationship. _____ [1]

(b) From the given occurrence diagram, explain what you can deduce about the degree and participation conditions of the Supplies relationship.

[2]

(c) From your answer to (b) and the further requirement that a part must be supplied by a single supplier, draw an E–R diagram showing the Supplies relationship between the two entity types Supplier and Part, including degree and participation conditions.

[2]

Question 3

The relational model for the University example used in the course has the following relations: *Student*, *Enrolment*, *Staff*, *Course* and *Assignment*.

In particular, the participation of a student on a course and their tutor for that course is represented by the *Enrolment* relation, which is defined as:

relation *Enrolment*
StudentId: IdentifiersOfStudents
CourseCode: CodesOfCourses
TutorNo: StaffNumbers

Identify the following keys of the *Enrolment* relation, or state 'None' if there is no such key.

(a) The primary key.

[1]

(b) All foreign keys, together with the relation referenced.

[3]

(c) All alternate keys.

[1]

Question 4

(a) A relation has an extension containing the following tuples:

11	a87	z41
16	a43	x58
38	a67	z23
53	a21	x84
76	a75	x45

What is the cardinality and what is the degree of this relation?

[2]

(b) Relations R and S both have an attribute A which is defined on the same domain. A is the primary key of R and an alternate and foreign key of S, referencing R. The cardinality of R is 300 and of S is 60, and there are no nulls in any tuple of R or S. What is the possible cardinality of the relations produced by the following relational algebra expressions?

(i) **select** R **where** A = 'xyz'

[1]

(ii) **join** R **and** S **where** $A = A$

[1]

(iii) R **times** S

[1]

Question 5

Consider the following two tables: **employee** which contains information about an organization's staff, and **assignment** which contains information about projects that an employee works on, and the percentage of their time committed to each project.

assignment

project_no	*emp_no*	*committed_time*
p1	e1	20
p1	e2	20
p2	e1	30
p3	e2	80
p3	e1	20

employee

emp_no	*name*
e1	Patel
e2	Said
e3	Timman

(a) For each of the following SQL queries give the table produced and a simple sentence describing the request that the query answers.

(i)
```
SELECT AVG (committed_time) AS av_time
FROM assignment
WHERE project_no = 'p3'
```

[1]

(ii)
```
SELECT project_no, committed_time
FROM assignment a, employee e
WHERE name = 'Said'
   AND a.emp_no = e.emp_no
```

[2]

(b) Give a list of steps describing the logical processing of query (a)(ii) according to the model given in the course text.

[2]

Question 6

The **course** table from the Study database contains the following data.

course_code	*title*	*credit*	*quota*
c1	Logic	60	null
c2	Syntax	30	null
c3	Functions	60	null
c4	Semantics	60	24
c7	Pragmatics	30	20

Suppose that a view is defined over the **course** table as follows.

```
CREATE VIEW long_course (cx,cy,cz) AS
  SELECT credit, course_code, quota
  FROM course
  WHERE credit >= 60
```

(a) Give the values that appear in column **cy** when the view is displayed in full.

[1]

(b) Give both the name of the column in the base table and the course code from each row modified by the following update.

```
UPDATE long_course
    SET cx = cx + 15
    WHERE cz > 20
```

[2]

(c) Assuming that the SQL statement given in part (b) has been executed, give the first two rows displayed by the following query.

```
SELECT *
FROM long_course
ORDER BY cy DESC;
```

[2]

Question 7

(a) Give an SQL statement that allows any user to retrieve all the rows from the table **country**, but not add to, change or delete any of its data. Completely accurate SQL syntax is not required.

[1]

(b) Briefly describe the effect of executing the following SQL statement.

```
GRANT UPDATE ON country TO secretary WITH GRANT OPTION
```

[2]

(c) Describe the effect of user **secretary** executing the following SQL statement assuming that the statements in parts (a) and (b), and no other GRANT statements, have previously been executed. What privileges does user **admin** have with respect to the **country** table?

```
GRANT INSERT ON country TO admin
```

[2]

Question 8

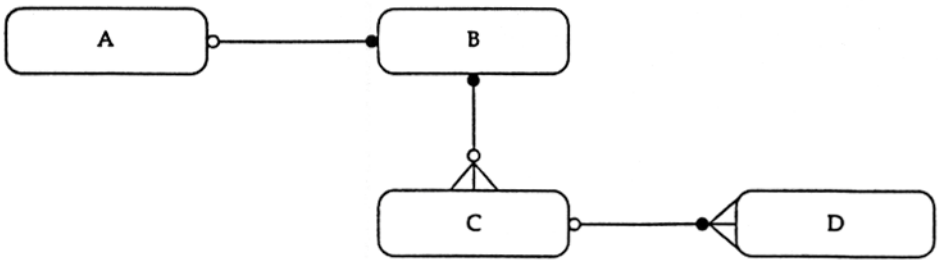

Figure 2

In the E–R diagram in Figure 2, the relationships between the four entity types, A, B, C and D, express constraints on occurrences of the entity types. (We have omitted names for relationships because they are not required in your answer.) Complete the following five sentences to give these constraints, where the first and last space in each sentence is one of A, B, C or D, and the middle space is one of *at least*, *no more than* or *exactly*.

An occurrence of ___ must be related to _____ one occurrence of ___ [1]

An occurrence of ___ must be related to _____ one occurrence of ___ [1]

An occurrence of ___ must be related to _____ one occurrence of ___ [1]

An occurrence of ___ must be related to _____ one occurrence of ___ [1]

An occurrence of ___ must be related to _____ one occurrence of ___ [1]

Question 9

The relation X has the following heading:

$X(\underline{A}, \underline{B}, C, D, E)$

The following functional dependencies (FDs) hold in X:

$A,B \mapsto C$
$A,B \mapsto D$
$D \mapsto E$

(a) Give one further FD that can be obtained by derivation.

[1]

(b) State the highest normal form of X, and justify it.

[2]

(c) Give the headings of the relations in BCNF that result from normalizing X.

[2]

Question 10

During database design, a data type has to be specified for each column.

(a) For values represented by numeric digits, give three aspects of these values that should be taken into account in choosing an appropriate data type.

[3]

(b) What does the SQL CAST function do that allows the choice of data type for a column to be unconstrained by the format of data for application processes?

[2]

PART II

Answer TWO questions from this part.

Question 11

Figure 3 gives an E–R model that describes how members of a walking club attend walks. On any one day when a walk is being held one member leads a walk over a given route. Other members attending on that date are called attendees. The walk leader is not counted as an attendee.

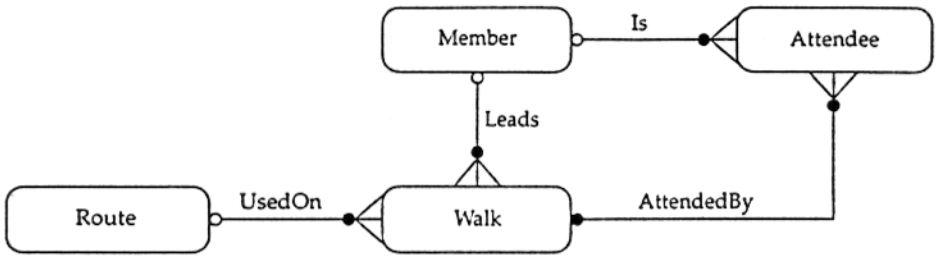

Entity types

Member (MemberNo, Name)
Walk (WalkDate)
Attendee (MemberNo, WalkDate)
Route (RouteNo, Description)

Constraint

A member who leads a walk is not also included as an attendee for that walk.

Assumption

There is only one walk on any given date.

Figure 3 Walks E–R model

(a) Given the following domains, produce a relational model corresponding to this E–R model, which should include all primary and foreign keys but not explicit representation of the participation constraints. We have given you the *Member* relation to remind you of the format required but you should reproduce it in your answer and add to it, if necessary. [10]

domains
MemberNumbers = integer
MemberNames = string
WalkDates = date
Routes = integer
Descriptions = string

relation *Member*
MemberNo: MemberNumbers
Name: MemberNames
primary key *MemberNo*

(b) For each participation condition shown in the E–R model, *either* explain how it is already represented in the model you gave as your answer to part (a), *or* give an addition to your model to include it. [10]

(c) Use the relational algebra described in the course to augment the model you gave as your answer to part (a) with the constraint specified in the Constraint part of Figure 3. [5]

Question 12

In the Study database used in the course, the **city** table includes the columns **name** and **country**, and the **country** table includes the columns **name** and **capital**. Suppose the following two embedded SQL statements are included in an application program which, amongst other statements concerned with details such as checking for errors, initially requests a user to input a value for the variable *y* and then finally prints out a value for the variable z.

```
EXEC SQL
SELECT country
INTO :x
FROM city
WHERE name = :y;

EXEC SQL
SELECT capital
INTO :z
FROM country
WHERE name = :x;
```

(a) Describe (do not give specific syntax) what else must be included in the program for these two statements to be acceptable as valid embedded SQL. [3]

(b) Assuming that the values for the **name** column in the **city** table are unique, give the request which is answered by these two statements together. [4]

(c) Give a single embedded SQL statement, which includes a subquery, that answers the same request. [6]

(d) Describe the logical processing model for the single statement you gave as your answer to part (c), and compare it with the combined logical processing of the two statements given in the question. [6]

(e) Explain why the assumption in part (b) is necessary for the proper execution of the given statements. If this assumption is not valid, describe an alternative set of statements which could be used in a program to answer the same request for each city that satisfies the condition; do not give detailed syntax but specify the intent of each statement and the use of the host variables. [6]

Question 13

A company, LearningTV, makes television and video programmes for a distance-learning university. It has established the following requirements for the data it wishes to record about the production of these programmes.

> Each programme is produced for use with a specific university course, and each course may have many programmes associated with it. For a programme it is required that its programme code, title and running time is recorded, and similarly the course code and the title of the course.
>
> A programme is the responsibility of a producer, one of the staff of LearningTV who has a staff number, name and telephone number recorded. A programme also has an advisor, one of the academic staff of the university who is identified just by name; their telephone number is also recorded, together with an estimate of the time they contribute to each programme. Each academic may be an advisor for many programmes associated with various courses, and a course may have many academics acting as advisors for its programmes.

(a) Give an E–R model to represent these requirements. This should include an E–R diagram, showing the degree and participation conditions of all relationships, and entity types, constraints and assumptions (you do not need to include assumptions of uniqueness for codes, and so on). [18]

(b) Suppose a programme also has a director, one of the staff of LearningTV. For each of the following distinct situations, explain whether the use of entity subtypes may be beneficial.

(i) A member of staff of LearningTV may be both a producer and a director, but not both for any one programme. [3]

(ii) Staff of LearningTV have a position that specifies whether they are either a producer or a director, which determines their role for all programmes. [4]

[END OF QUESTION PAPER]