

**UNIVERSITY OF LONDON**

**GOLDSMITHS COLLEGE**

**B. Sc. Examination 2007**

**COMPUTING AND INFORMATION SYSTEMS**

**IS52021A / IS52003A  
(CIS225 / CIS209)**

**Database Systems  
Internal**

**Duration: 3 hours**

---

This paper consists of **2** sections. Each section has **3** questions. Answer **2 questions** from each section. Each question carries **25** marks. Full marks will be awarded for **complete** answers to **4** questions.

There are 100 marks available on this paper.

The mark carried by each part is printed within square brackets. **Gauge the time to be spent on each part by the number of marks awarded.**

No calculators should be used.

**THIS PAPER MUST NOT BE REMOVED FROM THE EXAMINATION ROOM**

# SECTION 1

## Question 1

a) Define the following terms in no more than two sentences per term:

- |       |   |     |
|-------|---|-----|
| i.    | Data  | [1] |
| ii.   | Database  | [1] |
| iii.  | Database Management System                              | [1] |
| iv.   | Security  | [1] |
| v.    | Integrity   | [1] |
| vi.   | Views   | [1] |
| vii.  | Index   | [1] |
| viii. | Relation in the context of the relational data model    | [1] |
| ix.   | Intension in the context of the relational data model   | [1] |
| x.    | Cardinality in the context of the relational data model | [1] |
| xi.   | Foreign key in the context of the relational data model | [1] |

b) Define the terms of Data Manipulation Language and the Data Definition Language in no more than two sentences per term. For each such type of language provide three examples of commands of different types in the context of the relational databases, and explain what their effects are. [14]

## Question 2

The relational schema shown below is part of a hospital database. The names of the tables and their attributes are self-explanatory. The Patient table contains information about the patient names and addresses, and the dates of birth. Patients are admitted in wards and admission dates are recorded in table Contains. The table Ward contains information about the ward names, types, and the number of beds in each ward. Treatments are prescribed to each patient, and treatments details are stored in table Prescribed, namely the patient number, the drug number, the number of unit of the drug per day, and the start and the finish dates of the treatment with that drug. The table Drug keeps information regarding the drug names and the drug unitary costs. The primary keys are highlighted in bold.

*Patient* (**patientNo**, patName, patAddr, DOB)

*Ward* (**wardNo**, wardName, wardType, noOfBeds)

*Contains* (**patientNo**, **wardNo**, **admissionDate**)

*Drug* (**drugNo**, drugName, costPerUnit)

*Prescribed* (**patientNo**, **drugNo**, unitsPerDay, **startDate**, finishDate)

Formulate the following SQL statements:

- a) List all the patients' details, alphabetically by name, and the dates they were admitted in the word. [3]
- b) List all the patients contained in the 'Surgical' ward. [3]
- c) List all the patients admitted today. [3]
- d) Find the names of all the patients being prescribed 'Morphine'. [3]
- e) What is the total cost of Morphine supplied to a patient called 'John Smith'? [4]
- f) What are the maximum, minimum and average number of beds in a ward? Create appropriate column headings for the results table. [2]
- g) For each ward that admitted more than 10 patients today, list the ward number, ward type and number of beds in each ward. [4]
- h) List the numbers and names of all patients and the drugNo and number of units of their medication. The list should also include the details of patients that are not prescribed medication. [3]

### Question 3

A cleaning company specializes in providing cleaning services for both domestic and commercial clients. Each type of client has a set of requirements. For example, the Cardboard Box Company requires cleaning services from Monday to Friday 7am until 9am and 5pm until 7pm each day, but P. Nuttall only requires cleaning services on a Wednesday from 10am until 1pm.

The following information is collected for each client: the name, the address, the telephone and fax numbers. Whenever a new client is taken on, an administrator of the cleaning company assesses how many cleaning staff are required for the premises prior to assigning any staff to the job. Note that this is the ideal number, it may differ in practice. In addition, the administrator also assesses whether any specialist equipment is required and when. For example, three industrial floor cleaners may be needed on two out of five occasions for one commercial client.

The cleaning company keeps the following information about its staff members: the name, the address, the salary and the tax code, as well as a contact telephone number. The cleaning staff work in groups of six, with a supervisor to oversee the work done. The other staff are administrative staff who manage the day-to-day office work including visiting new clients and ensuring the specialist equipment is properly maintained.

By stating any assumptions you make when providing your answers, you are required to:

- a) Develop an Enhanced Entity–Relationship model from the above information. [11]
- b) Produce a set of tables from your Enhanced Entity–Relationship model clearly identifying each primary and foreign key. [14]

## **SECTION 2**

#### Question 4

Consider the relation R with the following attributes:

(VIN, Model, Capacity, Date, Garage, Address)

The relation represents information about vehicles, which are sent to various garages.

VIN stands for Vehicle Identification Number, Model represents a vehicle's model (e.g. Beetle), Capacity represents the vehicle's engine capacity (e.g. 2300 cc), Garage stores a garage's name, and Address stores the garage's address.

Consider the following semantic assumptions:

- (1) Each vehicle has a unique VIN, a non-unique model and a non-unique capacity.
- (2) Different vehicles of the same model can have different capacities.
- (3) On a given date, each vehicle can be sent to no more than one garage.
- (4) On different dates, a given vehicle may be sent to any garage. It may also be sent to the same garage more than once.
- (5) Different vehicles may be sent to a given garage at any time.
- (6) Each garage has a single but non-unique address (there may be two garages with the same address).

In each of the expressions in (a) (b) and (c) below, substitute the question mark with an attribute or a set of attributes, such that the resulting expression represents a non-trivial functional dependency that is satisfied by the relation R.

If more than one answer is right for a given expression, choose the *smallest* possible set if the question mark is on the left-hand side of the arrow, and the *largest* possible set if the question mark is on the right-hand side of the arrow.

- a) ?  $\rightarrow$  (Model, Capacity) [3]
- b) Garage  $\rightarrow$  ? [3]
- c) (?, Date)  $\rightarrow$  Garage [3]
- d) Find a candidate key for the given relation R. [2]
- e) Provide at least one explanation why R is not in BCNF. Explain your answer. [2]
- f) Using the functional dependencies in (a) and (b) above, apply Heath's theorem twice to derive an equivalent set of relations, all in BCNF. Each time you apply Heath's theorem, state explicitly which functional dependency you apply it to. For each new relation you obtain at each stage of the process, state if it is in BCNF; if it is not, provide a reason. State explicitly all the relations in the final set you obtain, and provide a candidate key for each. [12]

**Question 5**

a) Give an example of a transaction consisting of at least three database operations. Using either pseudo-code or English, demonstrate or explain how an application which invokes the transaction would implement the “all or nothing” approach to its completion. Which of the ACID properties is reflected in this approach? Note: if using pseudo code, you may use the additional primitives “BEGIN TRANSACTION”, “COMMIT TRANSACTION”, “ROLLBACK TRANSACTION”, and “ERROR”. [10]

b) Explain the term *checkpoint*. Provide a brief explanation of its role in database recovery and mention the two types of memory that are involved. [5]

c) There are five types of transactions that can be identified when recovering from a system failure, according to their position in relation to the most recent checkpoint and to the failure itself. Describe each of them, stating, in each case, the corresponding recovery action that a DBMS must take (a diagram may help your explanation). [10]



## Question 6

Assume that **R** is a relation with attributes **a<sub>1</sub>** through **a<sub>5</sub>**, and with **a<sub>1</sub>** as a primary key:  
**(a<sub>1</sub>, a<sub>2</sub>, a<sub>3</sub>, a<sub>4</sub>, a<sub>5</sub>)**

Assume that the following functional dependencies hold for **R**:

**fd1:**  $a_1 \rightarrow (a_2, a_3, a_4, a_5)$

**fd2:**  $a_2 \rightarrow a_3$

**fd3:**  $a_3 \rightarrow a_4$

**fd4:**  $a_4 \rightarrow a_5$

Assume that no other non-trivial functional dependencies hold, except for these four and for transitive dependencies, which may be derived from them.

- a) Explain the term *functional dependency*. Choose one of the functional dependencies fd1 through fd4 above and use it to demonstrate your explanation [3]
- b) Explain why the first of the above dependencies (fd1) is already given in the above description of the relation R [3]
- c) Based on fd1 through fd4, determine if the following functional dependencies hold for R (explain your answers):
1.  $(a_2, a_3) \rightarrow a_4$  [3]
  2.  $a_2 \rightarrow a_5$  [3]
  3.  $a_2 \rightarrow (a_3, a_5)$  [3]
  4.  $a_2 \rightarrow a_1$  [3]
- d) Define the concept of Boyce-Codd normal form (BCNF). [3]
- e) Explain the relevance of the assumption that no other non-trivial and non-transitive dependencies hold, to the claim that R is not in BCNF. [4]