UNIVERSITY OF LONDON

GOLDSMITHS COLLEGE

B. Sc. Examination, May 2010

COMPUTING AND INFORMATION SYSTEMS

IS52021A / CIS225

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

Consider the following database:

Parts(pid: integer, pname: string, colour: string)
Suppliers(sid: integer, sname: string, address: string)
Catalogue(sid: integer, pid: integer, cost: real)

The Catalogue relation provides the costs charged for parts whose name, colour and id are provided in the first relation above, by suppliers whose id, name and address are included in the second relation above. In particular, assume that the city is placed at the end of the address. Write the following queries in SQL:

a) Display the number of all suppliers. **[2]**

b) Display the price of the most expensive part. **[2]**

c) Display the average price per part across all suppliers. **[3]**

d) Find all suppliers in York that provide curtains. **[3]**

e) Find the names of the blue parts for which there is some supplier. **[3]**

f) Find the names of the suppliers who supply every part. **[4]**

g) For each part, find the name of the supplier who charges the most for that part. **[4]**

h) Find the ids of suppliers who supply only red parts. **[4]**

Question 2

a) Briefly describe how strong and weak entity types differ, using no more than two statements. Provide an example of each entity type. **[4]**

b) Describe what relationship types represent in an ER model and provide examples of unary, binary, ternary, and quaternary relationships. **[5]**

c) Suppose you are involved in the development of a database system for the Sales Department of a company. The operation of the Department can be described as follows.

They have a file of products that they provide to their customers. Each type of product has a unique product number, as well as a description, a cost and a price. For each type of product, the number of items available in stock and the number of allocated items are updated regularly. When the number in stock decreases to the reorder level, the product is reordered in a pre-decided quantity.

They have a file of customers. Each customer is given a unique customer number. This file also contains customer names that consist of their first and last names, and customer addresses

composed of street, city and postcode and the customer telephone number. Each customer has a credit limit, which is used to validate their orders.

A customer may place zero, one or more orders at a time, and an order is always placed by one customer alone. Each order is identified by a unique order number. Other information as to orders includes the date due, the total price, and the status, that is, an order may be outstanding, partially delivered, or fully delivered and invoiced.

An order may involve one or more than one type of products, and a type of products may be involved in more than one order. For each ordered product, its quantity, total price, and status (i.e., outstanding, partially delivered, or fully delivered) are recorded and updated regularly.

Given this information, draw an ER diagram for this Sales Department including entity names, relationship names, multiplicity constraints, and all the aspects regarding the attributes: primary keys, simple attributes, derived attributes, composed attributes, multiple-valued attributes, if they exist. **[16]**

Question 3

The following database formed of three tables Staff, Department and StaffInDept, stores information about the staff members of a company working across various departments. In particular the table StaffInDept shows the percentage of time, provided in the column *timepercentage,* that each staff memeber, identified by *staffno*, works in each department, identified by *deptno*. Each department has one manager identified by *managerno*, which corresponds to a *staffno*.

Staff (staffno: integer, name: string, age: integer, salary: integer)
Department (deptno: integer, budget: integer, managerno: integer)
StaffInDept (staffno: integer, deptno: integer, timepercentage: real)

You are required to:

a) Identify all the foreign keys in the database, indicating for each of them, the parent and child tables. **[3]**

b) Create the table Staff including a constraint that guarantees that every staff member is between16 and 65 years old. **[4]**

c) Create the table Department including a constraint that guarantees that all managers have a salary over £50,000. **[5]**

d) Create the table StaffInDept. **[3]**

e) Create an assertion that guarantees that the total percentage of time that every staff works across departments is no more than 100%. **[4]**

f) Increase the salaries of managers by £10,000. **[2]**

g) Delete the data on staff members which work 0% of time across all departments. **[4]**

SECTION 2

Question 4

a) Assume that the entity type *student* is linked to the entity type *module* via a "many to many" relationship type called *takes* (that is, a student may take many modules, and a module may be taken by many students). Specify the correct number of tables one would get if one converted the ER diagram described to the relational model. **[2]**

b) Convert the following EER diagrams to the relational model:

(i)

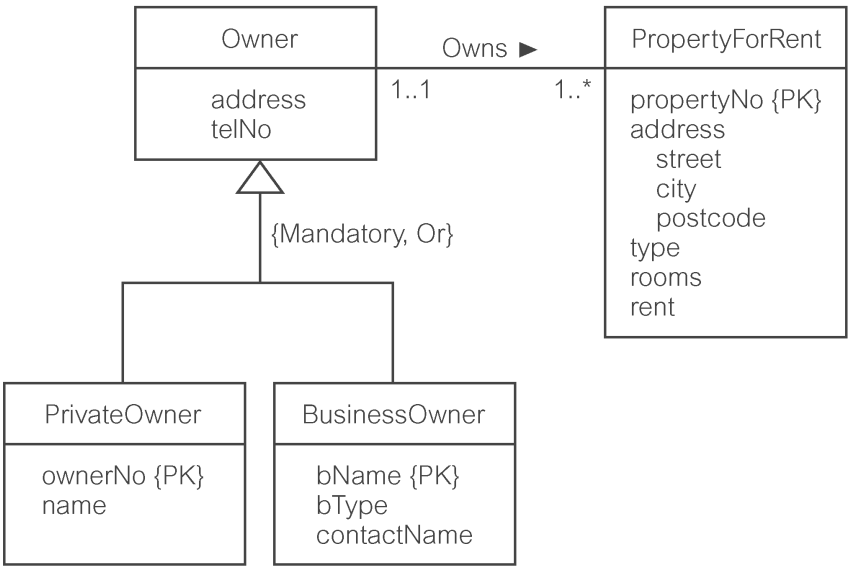

[9]

(ii)

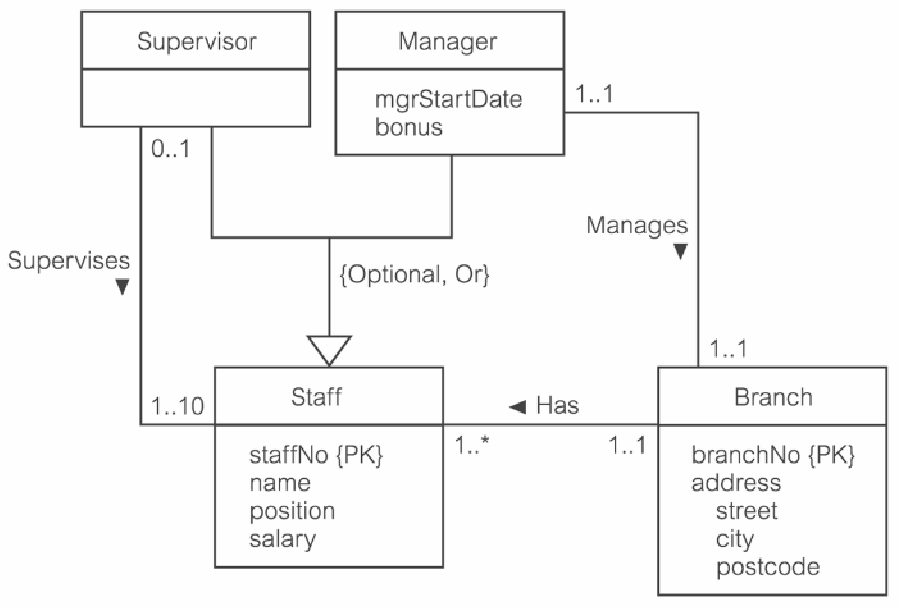

[14]

Question 5

a) The relations below, storing atomic values, satisfy the specified functional dependencies (FDs). Using these FDs, identify the candidate keys for each relation, and mention whether or not the relations are in each of the normal forms 2NF or 3NF. Justify your answer for each relation and for each normal form.

(i) **Staff**(staffno, staffname, job)
staffno -> staffname
staffno -> job **[5]**

(ii) **Staff**(staffno, staffname, jobcode, yearsinjob)
staffno -> staffname
staffno,jobcode -> yearsinjob **[5]**

(iii) **Staff**(staffno, staffname, jobcode, jobdescription)
staffno -> staffname,jobcode
jobcode -> jobdescription **[5]**

(iv) **Staff**(staffno, staffname, project, hoursworked)
staffno -> staffname
staffno, project -> hoursworked **[5]**

b) Decide whether the relation below, whose primary key is underlined, is in BCNF, and if not, illustrate the process of normalizing it. Clearly state each relation from the final decomposition, with its candidate key.

Student(studentno, studentname, course, mark)
studentno, course -> mark
studentno -> studentname **[5]**

Question 6

a) The following database stores data about the staff members from an estate agent, and about the properties to rent that they handle.

Staff(staffno, fname, lname, DOB, position, salary, phone, email, homeaddress)

Property(propertyno, type, rooms, address, rent, description, staffno)
FK staffno references Staff(staffno) {*indicates who handles the property}*

In order to tailor the access to data for different types of users, you are required to:

(i) Create a view *manager_view* that contains all details of staff members and the number of properties each of them handles. Grant read access on this view to the manager. **[4]**

(ii) Create a view *staffinfo* that contains all details of staff except the confidential information DOB, homeaddress and salary. Grant read access on this view to all staff members. **[3]**

(iii) Create a view *unhandled_properties* that contains all the properties that are not yet allocated to some staff member. Grant read and update access on this view to the manager and the secretary. **[3]**

b) Illustrate the lost update problem using a concrete detailed example, and provide a solution to prevent this problem. **[7]**

c) Enumerate and briefly define the ACID properties of transactions. **[8]**