UNIVERSITY OF LONDON

GOLDSMITHS COLLEGE

B. Sc. Examination, May 2012

COMPUTING AND INFORMATION SYSTEMS

IS52021A / CIS225

Database Systems

Duration: 3 hours

This paper is in two parts, Part A and Part B. There are a total of three questions in each part. You should answer TWO questions from Part A and TWO questions from Part B.

Full marks will be awarded for complete answers to a total of four questions, two from Part A and two from Part B. Each question carries 25 marks. The marks for each part of a question are indicated at the end of the part in [.] brackets.

There are 100 marks available on this paper.

No calculators should be used.

THIS PAPER MUST NOT BE REMOVED FROM THE EXAMINATION ROOM

PART A

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

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 member, 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 between 16 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]**

Question 3

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)

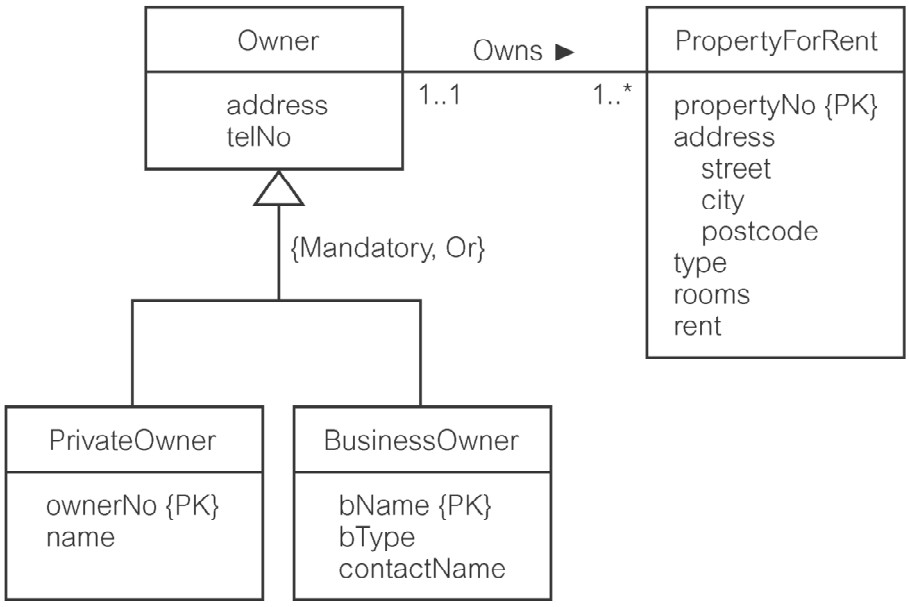

[9]

(ii)

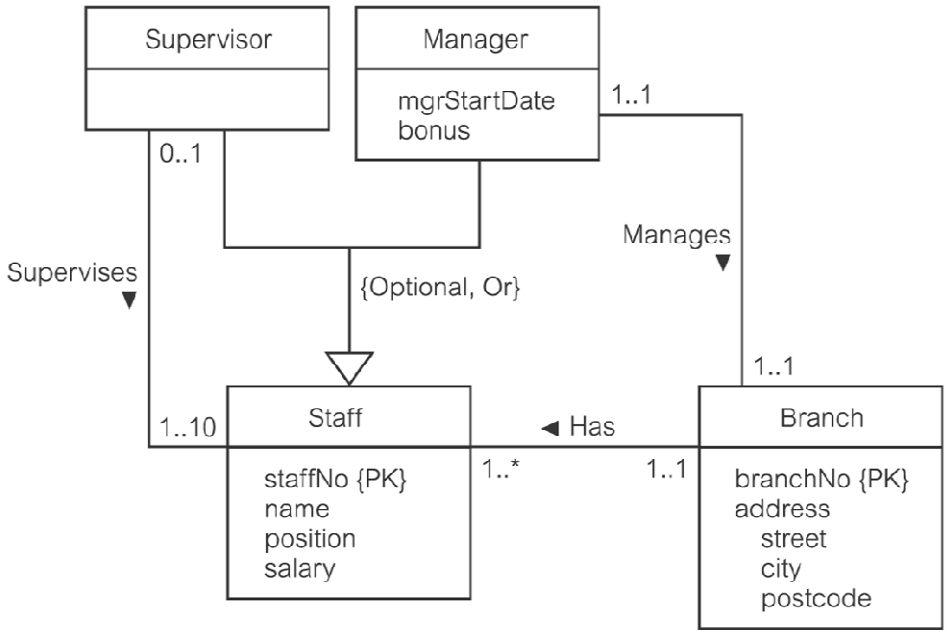

[14]

PART B

Question 4

(a) Define the first, the second and the third normal forms concepts (1NF, 2NF, 3NF). Additionally to each definition of a normal form provide an example of one relation that is in that normal form and one relation that is not in that normal form. Justify your answers in each case. **[15]**

(b) Define the inconsistent analysis problem in the context of database transaction management. Illustrate your definition with a detailed example when this problem occurs and show how it can be solved. **[7]**

(c) Define what concurrency control is, providing very briefly a justification for the need for the concurrency control (do not use more than three statement for your answer). **[3]**

Question 5

Consider the table R below, that stores information about students enrolled in modules. Assume that each module has one assignment per term and no written exam, and that a student gets a result for each module, which is a pass or a fail, based on the overall module mark.

student	**Module**	**term**	**assignmentMark**	**moduleMark**	**result**
Tim Cook	M105	1	70	60	pass
Tim Cook	M105	2	50	60	pass
Anne Blanc	M105	1	75	70	pass
Anne Blanc	M105	2	65	70	pass
Tim Cook	M108	1	36	33	fail
Tim Cook	M108	2	30	33	fail
Tim Cook	M112	1	69	70	pass
Tim Cook	M112	2	71	70	pass

Assume, also, the following functional dependencies:

fd1: (student, module, term) → assignmentMark
fd2: (student, module) → moduleMark
fd3: moduleMark → result

a) Check that the three functional dependencies above are satisfied by the relation R, showing the process of checking. **[6]**

b) Show that the three functional dependencies are irreducible. (*Hint: you can use the relation R, checking if there exists some reduced functional dependency obtained from the three functional dependences that is still satisfied by R*). **[7]**

c) Build a non-loss decomposition of the relation R into a set of relations in BCNF, showing how you apply Heath's theorem for each decomposition. Provide only the schemas of the relations (do not include tuples). Find candidate keys for each resulting BCNF relation. **[12]**

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]**