UNIVERSITY OF LONDON

GOLDSMITHS COLLEGE

B. Sc. Examination, Spring 2011

COMPUTING AND INFORMATION SYSTEMS

IS52021A / CIS225

Database Systems

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, in which the underlined attributes form primary keys:

Computer(computerId, computerName, make, cpu, ram, hd)
Store(storeId, storeName, address, telephone, website)
Availability(storeId, computerId, price)

The first relation stores details of computers to be sold, consisting of the computer id, the computer name, the make, the name of the CPU, the amount of RAM memory (in GB), and the capacity of the hard disk (in GB). The second relation contains details on computer stores, namely the store id, name, address, telephone and website. In particular the address is a string of characters that ends with the name of the city. The third relation contains information about the store where a computer is sold, and its price in that store.

You are required to write the following queries in SQL:

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

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

c) Display the name, the CPU and the RAM of all Dell computers. **[2]**

d) Display the name of all the computers, the name of the stores where they are sold, and their price. **[3]**

e) Display the computerId and the average price per computer across all stores, if the average price is under £1000. **[3]**

f) Display the name of stores in London that sell Sony computers based on a CPU called Intel Core i7-870. **[4]**

g) Find the name and address of the stores that sell all the computers that are available in the other stores. **[4]**

h) Display the website of the stores that sell the cheapest computer having a RAM memory of at least 16 GB, and a hard disk capacity of at least 1000 GB. **[5]**

Question 2

a) Consider the database formed of the following tables:

Branch (branchno, address, telephone, managerid)
Staff (staffno, first_name, last_name, job, salary, branchno)

The underlined attributes form primary keys, and managerid is a foreign key referencing Staff and is matching Staff's primary key.

Using Branch and Staff as underlying tables:

i. Provide an example of a view in SQL based on the two tables, that can be updated, and briefly explain why the view is updatable. Illustrate with an example showing how an insertion and an update SQL commands are performed on this view. **[8]**

ii. Provide two examples of views in SQL based on the two tables, that cannot be updated, and briefly explain why they are not updatable. **[8]**

b) Explain the two-tier and three-tier client server architectures used in the context of database systems. **[6]**

c) Is it allowed in relational databases for a table to contain a foreign key referencing the same table? Justify your answer (you may use a brief example). **[3]**

Question 3

Consider the database below storing information about properties for rent that are handled by staff working in the branches of an estate agent.

Property (propertyno, no, street, postcode, city, type, rooms, rent, staffno)
Branch (branchno, no, street, postcode, city, managerid)
Staff (staffno, first_name, last_name, job, salary, branchno, telephone)

The underlined attributes form primary keys in the respective tables. The meaning of the attributes is the following: managerid in table Branch corresponds to a staffno and indicates which staff member manages each branch; branchno in table Staff indicates in which branch the member of staff works; staffno in table Property indicates which staff member handles that property; all the other attributes are self explanatory in this context. You are required to:

a) Create the three tables in SQL. In particular, all primary key and foreign key constraints need to be expressed. Moreover, impose the conditions that the type of a property is either flat or house, and that the salary of staff cannot be more than 50k. Provide the necessary SQL commands in such an order that the DBMS does not return an error. **[12]**

b) Write a constraint in SQL that ensures that no staff member handles more than 60 properties. **[4]**

c) Write a constraint in SQL that ensures that all the properties handled by staff working in a branch are in the same city as the branch. **[4]**

d) Write an SQL statement that increases the rent of all houses by 10%, if the number of rooms is 4 or more. **[3]**

e) Write an SQL statement that increases the salary of all staff by 5%. **[2]**

SECTION 2

Question 4

a) You are required to:

i. Draw an ER diagram using precisely three entity types and two binary relationship types, that has a problem of chasm trap. You are not required to include attributes in this diagram, but you have to include the multiplicity constraints. **[4]**

ii. Draw a semantic net corresponding to the ER diagram from (i). Use this semantic net to justify why the diagram has a chasm trap problem. **[3]**

iii. Transform the ER diagram from (i) into a new ER diagram such that the chasm trap problem is solved. **[2]**

iv. Draw a semantic net corresponding to the ER diagram from (iii). Use this semantic net to justify why the diagram does not have a chasm trap problem. **[3]**

v. Name one problem with ER diagrams apart the chasm trap. **[1]**

b) As explained in the quoted text on top of the ER diagram below, the ER diagram represents the relationship type *Advertises* that links the entity types *Newspaper* and *PropertyForRent*.

i. Redraw the diagram by including at least one more attribute per entity type, indicating the primary keys and the multiplicity constraints, knowing that each property for rent is advertised by at least one and at most five newspapers, and that each newspaper advertises no more than thirty properties. **[4]**

ii. Do you get the same meaning of the ER diagram if the two attributes of the relationship type Advertises, standing for the date and the cost of an advertisement, were moved into the entity type Newspaper? Briefly justify your answer. **[3]**

iii. Convert the diagram from (b.i) to the relational model. For all the extracted relations, clearly mention their names, attributes, primary keys, and foreign keys. **[5]**

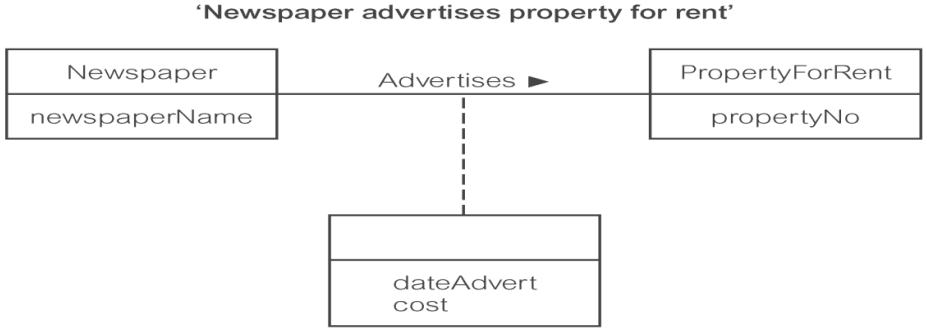

Question 5

a) Draw an EER diagram for the following specification, in which you are required to make use of the specialisation/generalisation relationship type also.

Each branch of a company has a branch number, and an address composed of street, city and postcode. The information about all the staff members working in the company's branches consists of the staff number, the name, position and salary. Any staff member is either full time permanent or part time temporary. The information available on full time permanent staff consists of a salary scale and in a holyday allowance, while the information on part time temporary staff consists of an hourly rate. On the other hand there are different subcategories of staff including the branch managers, with the manager start date and bonus as specific information, the sales personnel, with sales area and car allowance as specific information, the secretaries, with the typing speed as specific information, and other subcategories. It is also known that any staff member can be in one or more subcategories from those mentioned above. **[19]**

b) In the context of transaction concurrency control:

i. Define the concept of deadlock. **[2]**

ii. Provide an example of two detailed transactions that lead to a deadlock, and briefly explain why the deadlock occurs in this case. **[4]**

Question 6

a) Consider the relations with their functional dependencies given below. The relations are used to store information about a large number of students in a university, such that it may be possible two students to have the same name. You are required to use the functional dependencies in order to find the candidate keys of these relations. Then state if the relations are in each of the normal forms 2NF or 3NF. Justify your answer for each relation and for each normal form.

i. **Student**(sid, studentName, studyLevel)
sid -> studentName
sid -> studyLevel **[5]**

ii. **Student**(sid, studentName, courseCode, finalMark)
sid -> studentName
sid, courseCode -> finalMark **[5]**

iii. **Student**(sid, studentName, programme, programmeDescription)
sid -> studentName, programme
programme -> programmeDescription **[5]**

b) The following relation R satisfies the functional dependences Fd1, Fd2 and Fd3 below. You are required to normalise R into the Boyce-Codd normal form (BCNF), showing the process of normalisation in detail, including the candidate keys used in the process. Write the final decomposition of R explicitly.

R (sid, studentName, address, courseCode, term, courseworkMark, examMark, finalResult)

Fd1: sid -> studentName, address
Fd2: sid, courseCode -> examMark, finalResult
Fd3: sid, courseCode, term -> courseworkMark **[10]**