UNIVERSITY OF LONDON

GOLDSMITHS COLLEGE

B. Sc. Examination 2006

COMPUTING AND INFORMATION SYSTEMS

IS52021A (CIS225) / IS52003A (CIS209)

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

a) In order to explain what a fan trap is in ER modelling, do:

i. Define the concept of fan trap. [2]

ii. Provide an example of a diagram containing a fan trap. Show, through a new diagram, how the fan trap can be resolved, and justify your solution. [4]

b) Define the concept of weak entity type. Provide an example to illustrate this concept. [3]

c) What is a ternary relationship in ER modelling? Provide an example of a ternary relationship. [3]

d) In the context of EER modelling:

i. Describe and illustrate, using an example, the process of attribute inheritance (in no more than five sentences). [3]

ii. Describe what a superclass and a subclass represent, and provide an example; describe the relationship existing between a superclass and its subclass (use no more than five sentences). [3]

iii. Describe what a shared subclass represents and how this concept relates to multiple inheritance; illustrate the concept using an example (use no more than four sentences). [2]

e) Define each of the following terms in the context of the relational data model (use no more than one sentence per term):

i. Tuple [1]

ii. Relation [1]

iii. Intension [1]

iv. Cardinality [1]

v. Database [1]

Question 2

Consider the following schema of a database that stores data about employees working in various departments:

Department (did: integer, budget: integer, managerid: integer)
Employee (eid: integer, ename: string, age: integer, salary: integer)
Works (eid: integer, did: integer, pct_time: real)

Each department has one manager (managerid) which is also employee.
An employee can work in more than one department, and the pct_time field of Works table gives the percentage of time (represented as a real number between 0 and 1) an employee works in a given department. Using SQL, you are required to:

a) Define the table Employee including primary and foreign keys constraints as well as a table constraint on Employee that will ensure that every employee's name is known and that all employees are less than 65 years old. [4]

b) Define the table Department including primary and foreign keys constraints as well as a table constraint on Department that will ensure that all managers are more than 30 years old. [4]

c) Define the table Works including primary and foreign keys constraints. Then define an assertion that insures that the sum of the percentages of time an employee works in the departments is not more than 100%, for any employee. [7]

d) Mention an option that may be included in the definition of the foreign key "eid" in the table Works such that no delete command on the table Employee involves any modification of the content of the table Works. Illustrate your answer using an example of a delete command explaining its effect on the two tables. [2]

e) Write the commands for increasing the salaries of all employees by 10%, except for managers, whose salaries have to be increased by 15%. [4]

f) Write one statement that removes the details of the employees whose salaries exceed the salary of the manager of one or more of the departments in which they work. [4]

Question 3

Consider the following relations that keep track of flights information:

Flights(flno: integer, from_city: string, to_city: string, distance: integer, departs: time, arrives: time, price: integer)
Aircraft(aid: integer, aname: string, cruising_range: integer)
Certified(eid: integer, aid: integer)
Employees(eid: integer, ename: string, salary: integer)

Note that Employees relation describes pilots and other kinds of employees as well; every pilot is certified for some aircraft (otherwise, he or she would not qualify as a pilot), and only pilots are certified to fly.

Write each of the following queries in SQL:

a) Find the smallest price of the direct flights from London to Paris. [2]

b) Find the average salary of the employees. [2]

c) Find the number of pilots certified for Airbus aircrafts. [2]

d) Find the details of the pilots certified for Boeing aircrafts. [3]

e) Find the minimum salary of the pilots who can operate planes with a cruising range greater than 5000 miles. [3]

f) Find the aircraft id's (aid) of all aircrafts that can be used on non-stop flights from Tokyo to London. [3]

g) Find the names and eid's of all the employees who earn more than the average salary of all the employees. [3]

h) Find the names of the employees who earn the second highest salary. [3]

i) Find the eid's of the employees who are certified for the largest number of aircrafts. [4]

SECTION 2

Question 4

The director of *EasyDrive School of Motoring* has provided the following description. The school has several offices in most of the main cities of Scotland. Each office has a manager (who tends to also be a senior instructor), several senior instructors, instructors, technical and administrative staff. The manager is responsible for the day-to-day running of the office. Clients must first register at an office, providing their personal details. Before the first lesson, a client is requested to attend an interview with an instructor to assess the needs of the client and to ensure that the client holds a valid provisional driving licence. After the interview, the first lesson is booked. A client may request individual lessons (which last for one hour) or book a block of lessons for a reduced fee. A lesson is with a particular instructor, in a particular car, at a given time. A client is free to ask for a particular instructor or to request that an instructor be changed at any stage throughout the process of learning to drive. After each lesson, the instructor records the progress made by the client and notes the mileage used during the lesson. The school has a pool of cars of various brands and models, which are adapted for the purposes of teaching. Cars are inspected by the technical staff members at regular intervals for faults, and comments are recorded. Once ready, a client applies for a driving test date within a testing centre. The instructor is not responsible for testing the client, but should be available to drop off and pick up the client before and after the test at the testing centre. To obtain a full driving licence the client must pass both the practical and the theoretical parts of the test. The school records the results on the tests. If a client fails to pass, the instructor must record the reason for the failure.

You are required to:

a) Draw an ER diagram for the specification above, including the appropriate attributes. [19]

b) Choose a part of the diagram obtained in your answer to part a) comprising two entities and translate it into a relational model. [6]

Question 5

Consider the relation 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

You are required to:

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) Define the concept of transaction. Illustrate your definition with an example. [3]

b) Explain the ACID properties of transactions including a brief example for each property (use no more than five sentences per property). [12]

c) Give an example of two transactions that illustrates the "lost update problem". Show how this problem may be avoided in the case of this example. [5]

d) Give an example of two transactions that illustrates the "uncommitted dependency problem". Show how this problem may be avoided in the case of this example. [5]