UNIVERSITY OF LONDON

Goldsmiths College

Department of Computing

BSc Examination 2005

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.

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**.

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

SECTION 1

QUESTION 1

Part 1

a) What is a view? [2]

b) Discuss the difference between a view and a base relation. [3]

c) Explain what happens when a user accesses a database through a view. [3]

Part 2

a) Explain in detail the conditions under which aggregate functions and the GROUP BY clause can be used in a SQL SELECT statement. [6]

b) Give an example of a query that uses the GROUP BY clause; express the query also in natural language (English). [3]

Part 3 Which are the two major components of SQL, and what function does each of them serve? Do not use more than two statements in your answer. Give three examples of commands indicating the component to which they belong, and explain them in natural language (English). [8]

QUESTION 2

The following tables form part of a database held in a Relational Database Management System:

Employee	(**empID**, fName, lName, address, DOB, sex, position, salary, deptNo)
Department	(**deptNo**, deptName, mgrEmpID)
Project	(**projNo**, projName, budget, deptNo)
WorksOn	(**empID**, **projNo**, hoursWorked)

where	Employee	contains employee details and **empID** is the key.
	Department	contains department details and **deptNo** is the key. mgrEmpID identifies the employee who is the manager of the department. There is only one manager for each department.
	Project	contains details of the projects in each department and the key is **projNo** (no two departments can run the same project).
and	WorksOn	contains details of the hours worked by employees on each project, and **empID/projNo** form the key.

Part 1 Translate the following queries in SQL.

1) List all department names in alphabetical order. [2]
2) List all the details of employees who are male. [2]
3) List the average salary of managers. [2]
4) List the names and addresses of all employees who are programmers. [2]
5) Produce a list of the names and addresses of all employees who work for the 'Finance' department. [2]
6) Produce a complete list of all managers who are due to retire this year, in alphabetical order of last name (note that date_part('year', date) extracts the year from a date; the age for retirement is 65). [3]
7) Produce a report of the total hours worked by each employee, including the employee's id, last name, first name and department name, arranged in order of department name and, within the same department, in the alphabetical order of the last name of the employees. [3]
8) List the departments that have all the project budgets above £1 million (or, in other words, the departments that have no project budgets under 1 million). [3]
9) List the total number of employees in each department for those departments with less than 20 employees. Create an appropriate heading for the columns of the results table. [3]

Part 2 Create a view of employee details that excludes the date of birth and salary for all employees who work in the 'IT' department. [3]

QUESTION 3

Part 1 Draw an ER diagram for the following requirements: [20]

A company called Perfect Pets runs a number of clinics. A clinic has many staff and a member of staff manages at most one clinic (not all staff manage clinics). Each clinic has a unique clinic number (clinicNo) and each member of staff has a unique staff number (staffNo).
When a pet owner contacts a clinic, the owner's pet is registered with the clinic. An owner can own one or more pets, but a pet can only register with one clinic. Also, each pet owner should be registered with one clinic only (usually, with the one they contacted the first time). Each owner has a unique owner number (ownerNo) and each pet has a unique pet number (petNo).
When a pet is brought along to the clinic, it undergoes an examination carried out by a member of the consulting staff. The examination may result in the pet being prescribed one or more treatments. The company has a large number of approved treatments, and each individual treatment (for an individual pet) must correspond to such an approved treatment. Each examination has a unique examination number (examNo) and each type of approved treatment has a unique treatment number (treatNo).

Part 2 Choose a part of your diagram consisting of two entity types linked through a one-to-many relationship type. Consider two or more attributes per entity type and translate this part into the relational model, using SQL commands. [5]

SECTION 2

QUESTION 4

Part 1 Consider the following relation.

project	deadline	phase	contractor	budget	expenditure	reward/penalty

Express the following statements as irreducible functional dependencies:

a) each project has a unique deadline; [1]

b) any phase of any project is carried out by a unique contractor; [1]

c) each contractor negotiates a (unique) budget for each phase of each contracted project (for example, if a contractor has contracted the phase "design" for two different projects, they may have negotiated two different budgets); [1]

d) 'expenditure' represents the total sum of money spent by a contractor for each particular phase of each particular project; [1]

e) there is a system of rewards and penalties with reference to the expenditure; these depend solely on the sum that was under-spent or over-spent (e.g., reward/penalty can be calculated as "90% * (budget-expenditure)"); [1]

Advice: at points c and d you may be tempted to propose a reducible functional dependency; the answer to point b will help you to resolve this possible problem.

Part 2 Consider the table/relation described above. Show the table with a few inserted tuples in order to illustrate some data redundancy; show explicitly where redundancy occurs. [3]

Part 3 Consider the table/relation described above. Using explicit data values, give an example of an update anomaly. Accompany the example by a brief explanation. [5]

Part 4 Consider the following relation. It stores information about exams taken by students.

student	exam-paper	course	date	room	result

Consider the following functional dependencies:

student, exam-paper → result
exam-paper → course, date
exam-paper → room
course, date → exam-paper

Assume they completely express all the functional dependencies existing in the given relation (i.e., the others are either trivial or can be deduced from the given ones).

The given relation is not in BCNF. Decompose/transform it (non-loss) into a set of relations in BCNF. The normalisation/decomposition process should be carried out through the application of Heath's theorem. In this process, **you must consider the functional dependencies in the order in which they are listed above**. For each decomposition state:

a) the relation that is to be decomposed;

b) the functional dependency (or dependencies) on which the decomposition is based;

c) the resulting relations;

d) the candidate keys for each resulting relation;

e) for each resulting relation, whether it is or it isn't in BCNF.

When the normalisation process is finished, state the end result clearly. [12]

QUESTION 5

Consider the following two tables. They are used in all of the following parts.

Programmes

Code	Name	Running
CIS	Computing and IS	yes
CS	Computer Science	yes
IC	Internet Computing	yes
MC	Maths Computing	no

Registrations

StudentId	Username	Name	Qualification	*Programme*	Degree
DOC001	co1ac	Ahmed Cook	A-Level	CIS	null
DOC002	c01aj	Anna Jones	Access	CIS	null
DOC003	c01rp	Rachel Patel	Access	CS	2.1 hons
DOC004	c01sf	Steve Fuller	A-Level	IC	1 hons

The primary keys are printed in italics ('Registrations' has a composite primary key). The column 'Programme' in 'Registrations' is a foreign key referencing 'Code' in 'Programmes'. 'Username' is unique per 'StudentId'.

Part 1 The following insert operation is attempted:

```
INSERT INTO Registrations
        VALUES ('DOC010', 'c01af', 'Anna Fuller', 'A-Level', 'CDWWW', null);
```

The database management system generates an error.

a) Explain the reason why an error was generated. [2]

b) Explain what should be done in order for the insert operation to succeed. [2]

Part 2 Express the following constraints in SQL.

a) The qualification 'Short Course' is not acceptable for any programme. [3]

b) The qualification 'Foundation' is not acceptable for the 'CS' programme. [3]

c) No student should be allowed to register for programmes that are not running. [4]

d) No student can register for two or more programmes. [4]

Part 3 Express the following security rules in SQL.

a) The list of graduates (i.e., 'Degree' is not null) can be viewed by anyone. Allow access only to the names. [3]

b) Each student should be allowed to see all his/her records from 'Registrations', but should not be allowed to see any other students' records. [4]

QUESTION 6

Part 1 Consider the table 'Students' and its associated index 'Index-Students-Name' illustrated below.

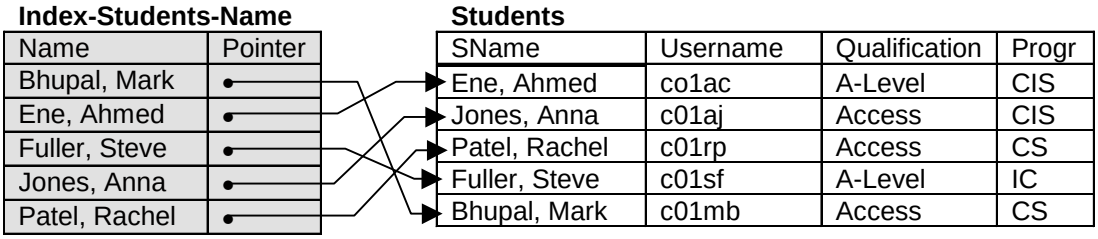

Index-Students-Name

Name	Pointer
Bhupal, Mark	•
Ene, Ahmed	•
Fuller, Steve	•
Jones, Anna	•
Patel, Rachel	•

Students

SName	Username	Qualification	Progr
Ene, Ahmed	co1ac	A-Level	CIS
Jones, Anna	c01aj	Access	CIS
Patel, Rachel	c01rp	Access	CS
Fuller, Steve	c01sf	A-Level	IC
Bhupal, Mark	c01mb	Access	CS

a) Explain how the index is used by the DBMS in the execution of the following query, by comparing with the execution of the same query in the absence of the index. [3]

SELECT * FROM Students WHERE SName = 'Bhupal, Mark';

b) Show the 'Index-Students-Name' and 'Students' tables after the following insertion is executed. [2]

INSERT INTO Students VALUES ('Dunn, Veronica', c01dv', 'Access', 'CIS');

Explain the operations carried out by the DBMS with the disk (in the presence of the index) upon the execution of the above insert statement. [3]

Part 2 Consider the following diagram regarding a popular Internet Bank. It includes information about Customers (name, contact details, security information, etc.), their Accounts (account number, type, balance, etc.) and all the Transactions performed on these accounts (date and type of transaction, amount handled, etc.).

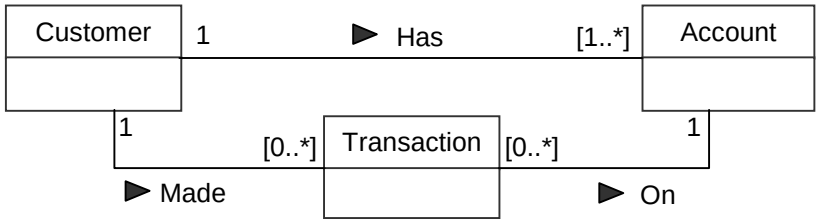

a) Amend the diagram in order to describe the average/expected number of instances per entity; for 'Transaction' consider a period of a month. Give brief explanations as of how you arrived at these numbers. [3]

b) State the average number of accesses per hour for the following queries/transactions. *Give brief explanations as of how you arrived at these numbers*.

1. Customer checks the balance of a particular account. [3]
2. Customer makes a transaction. [3]

c) Draw the corresponding transaction usage map. [2]

Part 3 Using the compatibility matrix, describe the data access protocol for update operations, in the context of concurrent access. You may use simple examples in order to illustrate the points you make. [6]