THE BRITISH COMPUTER SOCIETY

THE BCS PROFESSIONAL EXAMINATION
Professional Graduate Diploma

ADVANCED DATABASE MANAGEMENT SYSTEMS

25th April 2005, 10.00 a.m.-1.00 p.m.
Answer THREE questions out of FIVE. All questions carry equal marks.
Time: THREE hours.

*The marks given in brackets are **indicative** of the weight given to each part of the question.*

1. In early papers on the relational model of data, Dr. E. F. Codd proposed two mechanisms for manipulating data stored in a relational database. These were *relational algebra* and *relational calculus*. Compare and contrast these approaches and discuss their relevance to modern relational database products. **(25 marks)**

2. Explain why a query optimiser within a relational database management system is capable of generating query plans which execute faster than those generated by a human programmer. **(12 marks)**

a) Given two tables:

OrderHeader(OrderNo, CustomerNo, OrderDate)
OrderLine(OrderNo,PartNo,Qty)

Consider the query "Get the CustomerNo of customers who order part P1".

Assume that there are 100 OrderHeader rows and 1000 OrderLine rows and that P1 appears in 50 OrderLine rows.

b) Write down two different relational algebra sequences that satisfy this query and demonstrate that one is more efficient than the other. State one general optimisation principle that is evident from this example. **(13 marks)**

3. *a)* "Now we can state what might be regarded as the fundamental principle of distributed database: to the user, a distributed system should look exactly like a nondistributed system". [C. J. Date]

Discuss the subsidiary objectives that follow from this principle. **(18 marks)**

b) Outline three problems that are associated with distributed databases and give a brief description of the solution to each of them. **(7 marks)**

4. There is increasing interest in using databases to store and process 'geo-spatial' data (also known as 'location-aware' data) that contains data about the geographic location of people and things. An application is given below that uses geo-spatial data.

a) Describe two different ways of representing geo-spatial data such as that implied in the application described below. Include in your answer the associated relationships and constraints that would be needed to express geo-spatial data. Give examples of how you might model the data and associated relationships and constraints using:

ii) A relational database
ii) An object-oriented database **(16 marks)**

b) Discuss the extended features of a relational query language (such as SQL) that will be needed to support the querying of geo-spatial data. Your answer should refer to the queries in the application described below.

An application that processes geo-spatial data.
A University has adopted a personal identity card (PID) system to improve security and to restrict access to different groups of people at certain times and dates. To enter a building or room a person swipes their PID card through a card reader outside the door of the building or room. The data collected from the card reader is stored in a database for later analysis. For example it is possible to calculate the route a person has taken, the time they enter and leave each area.
The database can support queries such as:
Display all the rooms that person X has entered within the geographic area of buildings Y and Z.
Display the nearest lecture room to lab Y. **(9 marks)**

5. The protagonists of XML and supporting technologies claim that:

XML separates data structure from data representation.
XML will revolutionise data interchange and presentation on the web.
XML will revolutionise search and authoring on the web.

Discuss. **(25 marks)**