

THE BRITISH COMPUTER SOCIETY

THE BCS PROFESSIONAL EXAMINATION Professional Graduate Diploma

ADVANCED DATABASE MANAGEMENT SYSTEMS

23rd April 2004, 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. Extensible Markup Language (XML) is becoming the most commonly used data format for storing semi-structured data. Describe a scheme for storing XML data in a relational database. Discuss the performance implications of your answer for an application that retrieves a given XML element (e.g. given an XML file which contains product details, retrieve information about a given product). **(25 marks)**

2. “We do not believe that the database language SQL is capable of providing a firm foundation for the future.” (Hugh Darwen and Chris Date).

Critically evaluate this statement with reference to SQL’s support for the relational model of data. **(25 marks)**

3. a) Explain what is meant by the terms *granule*, *schedule*, *permutable actions* and *serialisability*. Show how these concepts can be used to determine if two database transactions may execute concurrently. **(10 marks)**

b) Demonstrate how the principles you have described in your answer to part a) have been embodied in the following three concurrency techniques:
 - i) Timestamp algorithms
 - ii) Optimistic concurrency control
 - iii) Locking**(15 marks)**

4. “An objective of a distributed database system is that to a user it should behave in exactly the same way as a non-distributed database system”.
 - a) With the aid of examples explain why it is difficult to achieve the above objective in practice. **(10 marks)**

 - b) With the aid of diagrams describe the technology needed to co-ordinate and manage the physical distribution of data using ‘replication’. Discuss the trade-offs that are needed to configure a distributed database application that supports replication. Apply your answer to a distributed database application with which you are familiar. **(15 marks)**

Turn over]

5. a) Explain the differences between Object Oriented and Relational data models in the way that they model the following:
- i) Object/Tuple identity
 - ii) Object/Tuple relationship (8 marks)
- b) Many database applications have constraints called 'business rules', which are programmed to prevent a database update taking place if a rule is violated.

Assume a RDBMS is deployed on the server side of an n-tier client-server platform and the following business rules are required.

Business Rule#1: *From a library database:*

Borrowers with status = 'student' cannot have any further loans if they have currently borrowed 5 items already because 5 is the loan limit for this type of borrower.

Business Rule#2: *From an exams database:*

Students are assessed on a course by taking 4 exams. Exam marks for a particular student are entered one after the other, if the mark entered is less than 30% then an overall grade of FAIL is returned and no further marks can be entered for this student. When all 4 marks are entered then the average mark is calculated and a grade of PASS or FAIL returned. For a PASS the average mark must NOT be less than 40%. Otherwise the grade given is FAIL.

- i) Explain how each business rule could be programmed to check for compliance following user interaction on a client interface. Illustrate your answer with sample code or pseudocode.
- ii) Discuss the factors that influence the decision as to whether business rules are implemented in middleware or on the database server in an n-tier client server platform. (17 marks)