

BCS THE CHARTERED INSTITUTE FOR IT

BCS HIGHER EDUCATION QUALIFICATIONS
BCS Level 6 Professional Graduate Diploma in IT

ADVANCED DATABASE MANAGEMENT SYSTEMS

Friday 23rd March 2018 - Afternoon

Answer **any** THREE questions out of FIVE. All questions carry equal marks.

Time: THREE hours

Answer any **Section A** questions you attempt in **Answer Book A**

Answer any **Section B** questions you attempt in **Answer Book B**

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

Calculators are NOT allowed in this examination.

Section A

Answer Section A questions in Answer Book A

Section A
Answer Section A questions in Answer Book A

A1

This question relates to concurrency control in the context of a multi-user on-line transaction processing environment.

- a) Explain, using examples from a database application (such as student registrations on courses), how the so-called **A.C.I.D.** properties of database transactions assist database integrity. **(8 marks)**
- b) Explain, with aid of examples, the difference between *serial* and *serialisable* schedules of transactions. Comment on whether (and if so, how) one is a superset of the other. **(5 marks)**
- c) Briefly explain how database transactions can be :-
(i) blocked.
(ii) deadlocked. **(3 marks)**
- d) Describe a range of techniques that can be used to counter or minimise the impact of deadlocks in a highly concurrent system. Comment on the effectiveness of each technique. **(5 marks)**
- e) Briefly describe row level locking and table level locking and comment on the pros and cons of applying these types of locks. **(4 marks)**

A2

This question relates to XML and/or related technology.

- a) Outline the main differences between the following data models:-
(i) Document Oriented.
(ii) Relational. **(4 marks)**
- b) Explain why the use of XML and/or any related technology is ideally suited to achieve the following objectives:-
(i) Represent document-oriented data (instead of using a relational database).
(ii) A means of data exchange over the WWW. **(8 marks)**

- c) What is an XSL stylesheet? Show the result of applying the XSL stylesheet in Fig A2 to the XML document in Fig A1. Give a brief explanation of how the result is obtained. (5 marks)

Fig A1 students.xml document:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<?xml-stylesheet href="mystylesheet.xsl" type="text/xsl"?>
<studentloan>
  <student>
    <name>
      <first>Nicky</first>
      <last>Melson</last>
    </name>
    <address>Forest Way 1, NR236HF, Burgh, Norfolk</address>
    <loan>
      <amount>5000</amount>
      <paymentdue>2020-01-01</paymentdue>
      <repayment amount="120" date="2017-05-01"/>
      <repayment amount="100" date="2017-06-02"/>
    </loan>
    <loan>
      <amount>9500</amount>
      <paymentdue>2021-01-01</paymentdue>
      <repayment amount="400" date="2017-01-01"/>
    </loan>
  </student>
  <student>
    <name>
      <first>Majeev</first>
      <last>Khan</last>
    </name>
    <address>Circus Street 8, YO438GT, Malton, N. Yorkshire</address>
  </student>
</studentloan>
```

Fig A2 xsl stylesheet

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<xsl:stylesheet version="1.0"
xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
<xsl:template match="studentloans">
<html>
<body>
<xsl:apply-templates select="//loan"/>
</body>
</html>
</xsl:template>
<xsl:template match="loan">
<xsl:apply-templates select="../name/last"/>,
<xsl:apply-templates select="../name/first"/>:
<xsl:apply-templates select="amount"/><br/>
</xsl:template>
</xsl:stylesheet>
```

- d) Write an XPath expression that returns all the names of students who have had at least one loan. (3 marks)
- e) Describe the function of the XQuery code listed in Fig A3 and determine the output produced when it is run against the XML document in Fig A1.

Fig A3 XQuery code listing:

```
<studentloan>
{
for $x in doc("atudents.xml")//student
return <student>
{$x/name}
<debt> {fn:sum($x/loan/amount) - fn:sum($x/loan/repayment/@amount)}
</debt>
</student>
}
</studentloan>
```

(5 marks)

A3

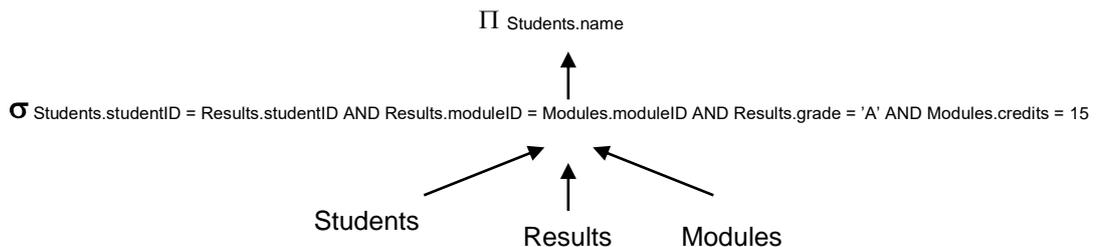
- a) Denormalization is often used as a way of tuning the performance of a database.
- (i) Discuss the advantages and disadvantages of denormalization. (4 marks)
- (ii) Using an example of a database design and query, show how denormalization can be used. (4 marks)

b) Given the following three linked tables:

```
Students (studentID, name, course)
Modules (moduleID, title, credits)
Results (resultID, studentID*, moduleID*, grade)
```

Suppose we have the following query and its corresponding initial parse tree:

```
SELECT Students.name
FROM Students, Modules, Results
WHERE Students.studentID = Results.studentID
AND Results.moduleID = Modules.moduleID
AND Results.grade = 'A'
AND Modules.credits = 15;
```



Transform the above parse tree into one that corresponds to the most efficient way of processing the query.

(11 marks)

c) Describe an appropriate database security control for each of the following security threats:

- Ransomware.
- SQL injection.
- An employee leaving the company.
- Brute-force password attacks.
- An employee making mistakes while entering data.
- A hacker intercepting data sent over the internet.

(6 marks)

Section B
Answer Section B questions in Answer Book B

B4

You are a **database consultant** bidding for a new contract with a prestigious blue-chip organization. Part of the selection process is a technical interview. Answer the following questions posed by the interview panel.

- a) With reference to a sample relation of your own choosing, explain and discuss the following relational model terminology, including its function and any related concepts.

- Tuple.
- Attribute.
- Domain.
- Constraint.

A diagram showing your sample relation is strongly suggested.

(10 Marks)

- b) Using your own SQL examples, illustrate the various types of joins that may be employed when extracting data from one or more database tables.

(10 Marks)

- c) Within the realm of distributed databases, explain the terms 'replication' and 'fragmentation'.

(5 Marks)

B5

You are a **data architect** bidding for a new contract with a prestigious blue-chip organization. Part of the selection process is a technical interview. Answer the following questions from the interview panel.

- a) Describe the defining characteristics of a *data warehouse* and how it differs in content and purpose from an 'OLTP' database system. You should use your own suitable examples and/or diagrams as needed.

(10 Marks)

- b) Explain, the term 'ETL' with respect to data warehouses, taking care to highlight common problems or issues in each stage. You should use your own suitable examples and/or diagrams as needed.

(10 Marks)

- c) Consider the phrase 'OLAP'. Explain what the term means, the underlying concepts involved, any associated benefits or limitations, typical applications and features along with any additional technical or implementation points you think appropriate to mention. You should support your discussion with suitable diagrams and/or examples.

(5 Marks)

END OF EXAM