

UNIVERSITY OF BOLTON/RAK CAMPUS

**SCHOOL OF BUSINESS AND CREATIVE
TECHNOLOGIES**

MSc INFORMATION TECHNOLOGY

SEMESTER 2 EXAMINATION 2009/2010

MODERN DATABASE SYSTEMS

MODULE NO: MIT4100

Date: Wednesday 2 June 2010

Time: 1.00 pm - 3.00 pm

INSTRUCTIONS TO CANDIDATES:

There are FOUR questions.

Answer ANY THREE questions.

All questions carry equal marks.

Marks for parts of questions are shown
in brackets.

QUESTION 1

- a) Define a distributed database management system and give its functions. **6 marks**
- b) By means of a diagram show the reference architecture for a distributed database management system. Show the various schemas present in the architecture and explain the reasons for each of these schemas. **8 marks**
- c) i) Discuss the reasons for fragmenting a relation. **3 marks**
ii) Explain the different rules that must be followed during the fragmentation. **4 marks**
- d) i) Explain briefly the need for a 2PC protocol for a distributed database management system involving at least 2 database servers. **2 marks**
ii) Draw state diagrams for the coordinator. (states INITIAL, WAITING, DECIDED, COMPLETED) and the participant servers (states INITIAL, PREPARED, ABORTED, COMMITTED) . **4 marks**
iii) How should recovery proceed if
1) one of the participant servers fails in the PREPARED state?
2) one of the participant servers fails in the COMMITTED state?
3) the coordinator server fails in its DECIDED state? **6 marks**

QUESTION 2

- a) i) Define the ACID properties of a transaction. Give the reasons why each property is vital for all transactions. **8 marks**
ii) By means of a diagram show the high level database module that handles transactions, concurrency control and recovery. Explain the various components present. **4 marks**

Question 2 continued over page

Business & Creative Technologies
 MSc Information Technology
 Semester 2 Examination 2009/10
 Modern Database Systems
 Module no: MIT4100

Question 2 continued

b) i) Explain the lost update problem referring to the following transactions

Time	T1	T2	
t 1		begin_transaction	
t 2	begin_transaction	read(balX)	
t 3	read(balX)		
t 4		balX = balX + 100	
t 5	balX = balX - 50		
t 6		write(balX)	
t 7	write(balX)	commit	
t 8	commit		2 marks

ii) What is 2-Phase Locking? Give its rules.

4 marks

iii) Write a solution to the problem occurring in b (i) using 2PL and explain it.

7 marks

c) i) With the help of transactions T3 and T4, explain the deadlock problem that can occur

with locking based mechanism for concurrency control. Illustrate your answer with a list of operations of each transaction.

4 marks

ii) What is time stamping? How do time stamping based protocols for concurrency control differ from locking based protocols?

4 marks

Please turn the page

QUESTION 3

- a) i) Define database replication. **1 mark**
- ii) Discuss the benefits of database replication to an organization. **4 marks**
- iii) Explain the basic components of the database replication environment—
replication objects, replication groups, and replication sites. **4 marks**
- b) i) Compare and contrast synchronous with asynchronous replication. **3 marks**
- ii) With the help of suitable examples explain any 2 possible forms of data
ownership. **6 marks**
- iii) Database triggers is one of the issues in the implementation of replication. Give
the drawbacks of this approach. **4 marks**
- c) Replication is a main feature of mobile databases. Discuss 3 items of extra
functionality required of a mobile DBMS. **6 marks**
- d) What is the need for a replication server? Give the different replication server
functionality. **5 marks**

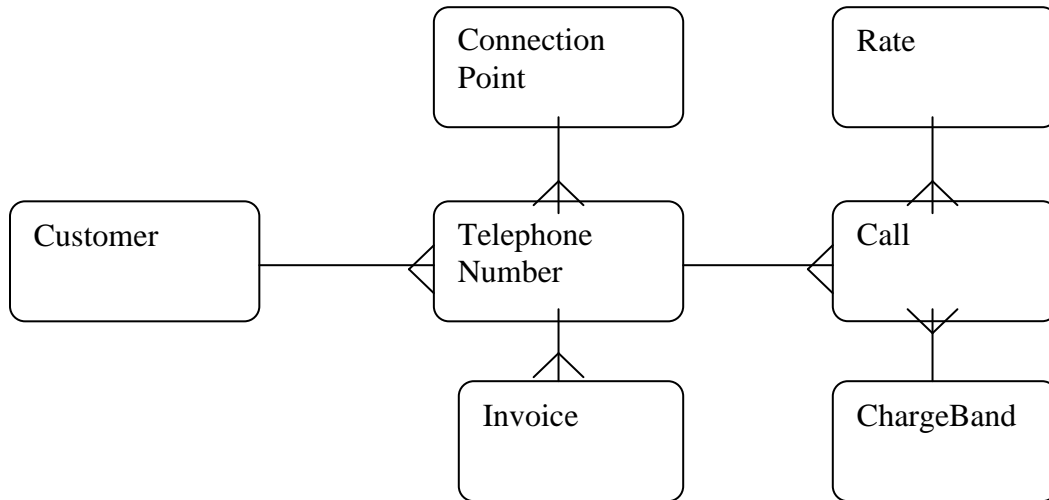
QUESTION 4

- a) i) Define what is meant by a 'data warehouse'. Explain three ways in which data warehouses differ from operational (OLTP) databases **4 marks**
- ii) The typical architecture of a data warehouse contains three major pieces of software – the *load manager*, the *warehouse manager* and the *query manager*. Outline the functions of each. **4 marks**
- b) i) There are five information flows within a data warehouse, namely: inflow, metaflow, downflow, upflow and outflow. Explain the middle three. **3 marks**
- ii) Explain what is meant by data warehouse *metadata* with examples. **3 marks**
- iii) Define a *data mart*. Discuss the relationship between data marts and data warehouses. **4 marks**
- c) The Open Tele Case Study referred to in the course has an E-R diagram as shown below.
- The directors wish to trace the customers' usage of the network, and try to answer questions such as the following:
- What makes customers churn?
 - Are some customers more likely to churn than others ?
 - How can we identify those that may churn before they do?
- i) For an Open Telecoms data warehouse for network usage, identify with reasons, its facts and dimensions. Draw up a star schema for the warehouse and give details of each entity in the data warehouse. **9 marks**
- ii) Outline briefly for one of the three sample questions above, the SQL that would go some way towards answering the questions (exact syntax not required). **3 marks**
- iii) Describe how the data for the warehouse could be set up initially, and then on a day to day basis. **3 marks**

Question 4 continued over page

Business & Creative Technologies
 MSc Information Technology
 Semester 2 Examination 2009/10
 Modern Database Systems
 Module no: MIT4100

Question 4 continued



The entity types in the E-R model are as follows:-

Customer (AccountNumber, Name, Category, Address, PostCode)

ConnectionPoint (NodeNumber, ClassOfSwitch)

Rate (RateCode, RatePerSecond)

Call (TelNumber, StartDate, StartTime, NumberCalled, CallDuration, CallCost)

LineRental (RentalClass, BillFrequency, RentalAmount)

ChargeBand (BandCode, BandDescription, StartTime, StopTime, ChargeFactor)

TelephoneNumber(TelNumber , TypeOfLine)

END OF QUESTIONS