



THE UNIVERSITY OF THE WEST INDIES

EXAMINATION OF APRIL/MAY 19 2001

Code and Name of Course: CS35A - DATABASE MANAGEMENT SYSTEMS

Paper:

Date and Time: Wednesday May 09, 2001 9:00 - 11:00 a .m.

Duration: 2HRS

INSTRUCTIONS TO CANDIDATES: This paper has 5 page(s) and 4

questions

Answer questions 1 and any two others.

Question 1 [30 Marks]

a) Consider the following information:

The University of the West Indies has decided to invest in a database to maintain information of the assignment of Students to the various halls on campus (e.g. Taylor Hall, Chancellor Hall, etc.). All students are assigned to halls, whether they live on or not. This is done since students engage in activities for their halls.

Each hall has a name, location and telephone number. These halls are either co-ed or single sex. For those halls that are co-ed, information about the number of male rooms and the number of female rooms must be kept (since males and females are not assigned to a single room), while for those halls that are single sex, the sex must be kept.

For students who live off campus a contact number and address where they will be living during the semester must be recorded. For those who live on campus, a name and phone number of a next of kin must be recorded.

Each room has a number, cost and a type (single occupancy, double occupancy) associated with it. All rooms of a hall that are of the same type cost the same. Students are assigned rooms and the date the student is assigned the room must also be recorded.

.../2

- i. Draw an Entity-Relationship Diagram for the information above. State any assumptions that you made and any constraints that you were unable to represent. [9 Marks]
 - ii. Convert the diagram into a set of relations normalised up to third normal form. You must state why these relations are in third normal form. [7 Marks]
- b) Discuss, using examples, the importance of declaring primary and foreign keys when implementing a database. [4 Marks]
- c) Assuming a relational database with the following relations, where the primary key of each relation has been underlined:

CUSTOMER

Field Name	Type	Description
<u>CID</u>	Integer	Unique customer id.
Cname	Character	Customer's name
Cphone	Character	Customer's telephone number

ORDER

Field Name	Type	Description
<u>OID</u>	Integer	Unique order id.
CID	Character	Id. of customer making order
ODate	Date	Date customer made order

CONSISTSOFT

Field Name	Type	Description
<u>OID</u>	Integer	Order id.
<u>PID</u>	Integer	Product id.
Qty	Integer	Quantity of product on order

PRODUCT

Field Name	Type	Description
<u>PID</u>	Integer	Unique product id.
Pname	Character	Name of product
Pcost	Integer	Cost of product

Formulate the following queries in SQL:

- i. Find the id. of the order that 'John Brown' made on 'December 3rd 2000'.
[2 Marks]
- ii. Find the customer id's of those customers that have made orders that consist of 'soap'.
[2 Marks]
- iii. Find the customer id's of those customers that have made more than three orders.
[2 Marks]
- iv. For each order, find the total cost.
[2 Marks]
- v. Increase the cost of the product with id. 123 by \$50.
[2 Marks]

Question 2 [15 Marks]

Using the same set of tables specified in Question 1 b), do the following:

- a) Write a relational calculus expression for the following query: "Find the name and phone numbers of those customers that have ordered product 111 (id) and the order was made after Jan. 1st 2001.
[2 Marks]
- b) Convert the relational calculus expression in a) into an equivalent relational algebra expression.
[2 Marks]
- c) What is the fundamental difference between the two languages (i.e. relational algebra and relational calculus)? Explain. Which category does SQL fall under? Give reasons.
[3 Marks]
- d) Find an alternative relational algebra expression to that derived in b). Compare the two expressions in terms of efficiency of query processing.
[4 Marks]
- e) Discuss how the decision of query processing is affected if the information is stored in a distributed database. Use examples to emphasise your points.
[4 Marks]

.../4

Question 3 [15 Marks]

A representative from the Sandals chain of hotels is seeking your advice. Sandals are planning to implement a database for their operations and are currently trying to choose between a centralised and a distributed database. After some consultation you obtain the following information: *Sandals hotels are located throughout Jamaica as well as other islands in the Caribbean and it is expected that the number of hotels will continue to grow. The head office is here in Jamaica. The database is expected to store information about the hotel's customers. Customers are typically repeat customers i.e. once they visit a particular Sandals hotel it is likely that they will return to that hotel on subsequent visits. Each hotel needs to keep information about their employees. Within each hotel there are a number of sections e.g. dining, entertainment, reservations, finances, etc. The various employees tend to work in specific sections. Management will need to have an idea of how well both each hotel and the entire chain is performing.*

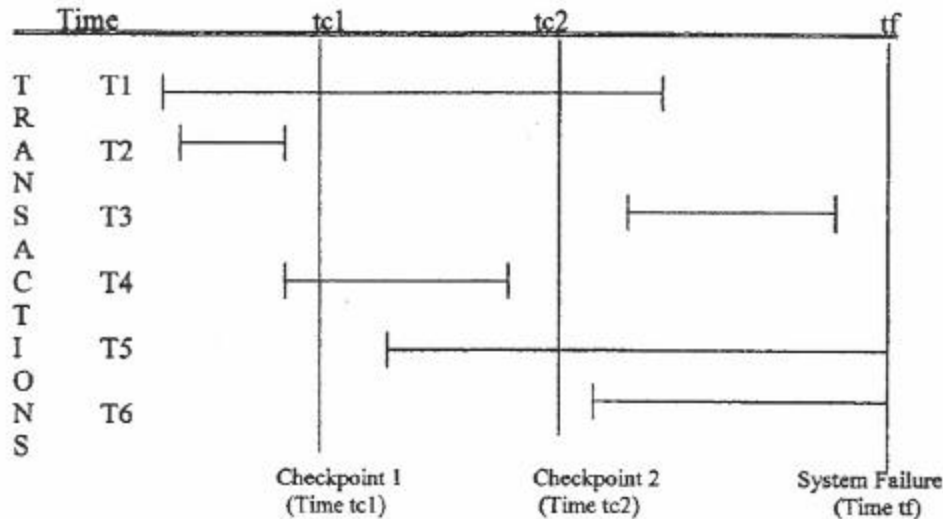
The company is willing to spend the money on the technology necessary to make both options feasible.

- a) Would you advise the company to choose a centralised or distributed database? Give reasons. **[5 Marks]**
- b) Give an example of how you would vertically fragment the data if the company opted for a distributed database. **[3 Marks]**
- c) Give an example of how you would horizontally fragment the data if the company opted for a distributed database. **[3 Marks]**
- d) Would you advise the company to replicate the data? If so why and to what extent? If not, give reasons. **[4 Marks]**

.../5

Question 4 [15 Marks]

a) Given the following information:



Where:

- A system failure has occurred at time tf .
- The two most recent checkpoints prior to time tf were taken at time $tc1$ and $tc2$, where $tc1$ is earlier than $tc2$.
- Transaction T1 started prior to time $tc1$ and completed (successfully) after time $tc2$ and before time tf .
- Transaction T2 completed (successfully) prior to time $tc1$.
- Transaction T3 started after time $tc2$ and completed (successfully) before time tf .
- Transaction T4 started prior to time $tc1$ and completed (successfully) prior to time $tc2$.
- Transaction T5 started prior to time $tc2$ but did not complete by time tf .
- Transaction T6 after time $tc2$ but did not complete by time tf .

For each transaction describe what needs to be done, after the system failure, to ensure that the database is restored to a consistent state. [6 Marks]

- b) *Concurrent database transactions must be serializable.* Describe what this sentence means. [4 Marks]
- c) Explain, using an example, what a deadlock is. Describe, using the example, one technique that can be used to remove deadlocks. [5 Marks]