

Ollscoil na hÉireann, Gaillimh
National University of Ireland, Galway

GX1496

Spring Examinations 2003/2004

| | |
|----------------------|--|
| Exam Code(s) | 3IF1 |
| Exam(s) | 3 th BSc. In Information Technology |
| Module Code(s) | CT332 |
| Module(s) | Database Systems II |
| Paper No. | |
| Repeat Paper | |
| External Examiner(s) | Prof. P. Nixon |
| Internal Examiner(s) | Prof. G. Lyons |
| | Mr. C. O'Riordan |
| Instructions | Answer any 3 questions All questions carry equal marks. |
| Duration | 3 hours |
| No. of Answer books | 1 |
| | |
| | |
| | |
| | |
| | |
| No. of Pages | 4 |
| Department(s) | Information Technology |

OLLSCOIL NA hÉIREANN GAILLIMH
NATIONAL UNIVERSITY OF IRELAND GALWAY

SPRING EXAMINATIONS 2004

Third University Examination in Information Technology (3IF)

Database Systems II (CT332)

1. (a) Discuss the properties of a well designed relational schema. (6)
- (b) Describe a procedure you might adopt to map the following components of an EER diagram to a suitable relational schema: relationships (binary (1:1, 1:N, N:M), n-ary), specialisations and categorisations. Include examples, where appropriate, to illustrate your answer. (12)
- (c) Explain how the process would differ if you were developing a schema suitable for an object-oriented database. (8)
- (d) Given $R = \{A, B, C, D, E, F, G, H, I, J, K\}$ and the following functional dependencies:
- $\{A, B, C\} \rightarrow \{F, G, H\}$
 - $\{D\} \rightarrow \{A\}$
 - $\{A, B\} \rightarrow \{D\}$
 - $\{C\} \rightarrow \{E\}$
 - $\{H\} \rightarrow \{I\}$
 - $\{I\} \rightarrow \{J, K\}$
- decompose R to be a set of relations such that all relations satisfy BCNF. (7)

2. (a) Given the following database schema and interpretation:

SLSREP: Rep_No, Rep_Name, Addr, Status
CUSTOMER: Cust_No, Cust_Name, Addr, Balance, Rep_No
ORDER_INX: Ord_No, Cust_No, Date
ORDER_LINE: Ord_No, Item_No, quantity, Ord_Price
ITEM: Item_No, Item_Desc, Item_Price, Colour

(Keys are underlined).

The SLSREP relation stores information regarding sales representatives—name, address, status and unique number. The CUSTOMER relation is used to maintain information on customers. The attribute Rep_No in CUSTOMER acts as a foreign key to the SLSREP relation. ORDER_INX stores information on orders placed by customers. Any order

may contain a number of items and this information is stored in the ORDER_LINE relation. The final relation, ITEM, stores information on items—number, description, price and colour.

Develop SQL queries to satisfy the following information needs:

- List all customers (with balance less than 200) who have ordered a red-coloured item.
- List all customers (number and name) and the number of orders containing at least three items placed by those customers.

(6)

- (b) Outline the process of heuristic optimisation. Develop an operator tree that represents an efficient evaluation strategy for the query satisfying the information need:

“List all customers (with balance less than 200) who have ordered a red-coloured item.”

(11)

- (c) Outline and compare algorithms that may be used to implement the join operator. (8)

- (d) Explain how the optimised query tree developed in part (b) may be further optimised in a parallel database. (8)

3. (a) The two-phase locking protocol is a commonly adopted approach to ensure conflict serializability of schedules. Explain the terms *two-phase locking protocol* and *conflict serializability*. (6)

- (b) Given the following schedule, show how it would proceed under two-phase locking.

| TA | TB | TC |
|----------|----------|----------|
| | read(Y) | |
| | write(Y) | |
| read(X) | | |
| write(X) | | |
| | | read(Z) |
| | | read(Z) |
| | read(X) | |
| | write(X) | |
| read(Z) | | |
| | | read(Y) |
| | | write(Y) |

(10)

- (c) Describe the main entries in the system log. Assuming the immediate update protocol, show how the recovery process would proceed given the following log entry.

```
[start_transaction, T1]
[read, A, T1]
[write, A, 10, 12, T1]
[start_transaction, T2]
[read, A, T2]
[write, A, 12, 13, T2]
[commit, T1]
[read, B, T2]
[write, B, 50, 60, T2]
[checkpoint]
[start_transaction, T3]
[read, C, T3]
[write, C, 45, 60, T3]
[read, D, T3]
[read, C, T2]
[read, D, T2]
[write, C, 60, 61, T2]
[write, D, 100, 0, T2]
[commit, T2]
<crash>
```

How would the log differ under the deferred update protocol? (9)

- (d) Describe how the two phase commit may be used in a distributed database to ensure the atomic nature of transactions. (8)

4. (a) With respect to deductive databases, explain the terms: *fact*, *rule*, *rule-safety*, and *backward chaining*. (6)

- (b) Show with examples, how the operators (select, project, join and union) can be implemented in Datalog. (11)

- (c) Outline the structure of a B-tree. Calculate the number of blocks required to store a B-tree given the following data:
60,000 records to be stored, Average fill factor is 69%; Search field is unique and requires 9 bytes.
Block pointers require 6 bytes. Block size is 512 bytes. (8)

- (d) Write a note discussing the security issues in statistical databases. (8)