110

(University of Choice)

MASINDE MULIRO UNIVERSITY OF SCIENCE AND TECHNOLOGY (MMUST)

MAIN CAMPUS

UNIVERSITY EXAMINATIONS

2021/2022 ACADEMIC YEAR

SECOND YEAR SECOND SEMESTER EXAMINATIONS

FOR THE DEGREE

OF

BACHELOR OF SCIENCE IN INFORMATION TECHNOLOGY

COURSE CODE: BIT 221

COURSE TITLE: DATABASE MANAGEMENT SYSTEMS

DATE: Wednesday 20/04/2022 **TIME:** 8:00a.m-10:00a.m

INSTRUCTIONS TO CANDIDATES

Question ONE (1) is compulsory
Attempt any TWO (2) questions

TIME: 2 Hours

MMUST observes ZERO tolerance to examination cheating

This Paper Consists of 3 Printed Pages. Please Turn Over.

QUESTION ONE [30 MARKS]

a) In recent years there has been a trend to represent traditional relational data in a semi-structured data format such as XML.
   i. List and explain three of the main motivations for this trend over recent years. **[6 Marks]**
   ii. State an application and explain why the stated application warrants the need to store and process semi-structured data. **[6 Marks]**

b) Describe what you understand by the term 'Relational Model' and identify the main components of relational Model **[6 Marks]**

c) Today, instead or purchasing physical *machines*, an alternate approach to running a database management system is to run it in the cloud. One of the key promises of the cloud is the illusion of infinite resources and the ability for users to elastically grow and shrink the resource consumption of their DBMS.
   i. Why can it be difficult to scale a relational DBMS? **[6 Marks]**
   ii. Indicate one approach that the so-called "NoSQL" systems take to overcome the above challenge (Why can it be difficult to scale a relational DBMS). **[6 Marks]**

QUESTION TWO [20 Marks]

The following database contains information about actors, plays, and roles performed.
Actor(actor_id, name, year_born)
Play(play_id, title, author, year_written)
Role(actor_id, character_name, play_id)

Where:
Actor is a table of actors, their names, and the year they were born. Each actor has a unique actor_id, which is a key.
Play is a table of plays, giving the title, author, and year written for each play. Each play has a unique play_id, which is a key.
Role records which actors have performed which roles (characters) in which plays. Attributes actor_id and play_id are foreign keys to Actor and Play respectively. All three attributes make up the key since it is possible for a single actor to play more than one character in the same play.

Write a SQL query that returns the names of all actors who have performed some play by the author "Putin" and have never performed in any play written by author "Biden". The list should not contain duplicates but does not need to be ordered.

[20 Marks]

QUESTION THREE [20 Marks]

Assume a journal library database holds the following Tables (with sample data) at a central site on a database server called (MMUST_LIB_SRV)

JOURNAL **Author**

JournalID	JournalName
3125	Database Weekly
3126	Database Journal
3127	Oracle News
3128	ACM TODS

AuthorID	AuthorFName	AuthorSname
23	James	Ouma
18	Grace	Mukulu
67	John	Murkomen

ARTICLE

ArticleID	ArticleTitle	AuthorID	ournalID
3115	ObjectOriented Analysis	23	3126
2409	Oracle indexing	18	3127
1398	DBA performance tools	23	3126
1289	Pioneers of Databases	23	3125
2554	Query optimisation	67	3126
1678	Daplex	18	3128
4561	Niam Fact model	18	3128

LOAN

ArticleID	BorrowerID	LoanDate	ReturnDate
2409	23	3/1/15	4/2/15
1398	23	3/1/15	24/1/15
1289	17	6/2/15	8/2/15
2554	26	1/2/15	12/2/15
2409	43	14/2/15	
2554	52	14/2/15	

BORROWER

BorrowerID	BorrowerFname	BorrowerLname	BorrowerTelNo
52	Jane	Mukami	07123456
43	Fred	Onsongo	072345678
17	Henry	Maua	072134567
26	Jonas	Mutua	073123456

The above tables are currently held on a database at the single site called MMUST_LIB_SRV. Following reorganization it is intended to distribute the journals held at the central database across 3 new branch libraries located at remote sites called MM_SRV, BG_SRV and NR_SRV. The central library becomes the HQ (an administrative centre) meaning that it no longer keeps or loans out any journals itself. Instead journals are made available for loan to borrowers registered at any of the 3 new sites.

a) Briefly Describe in general terms THREE different approaches of achieving database distribution. **[8 Marks]**

b) Describe in detail THREE different proposals for data distribution of the central database (MMUST_LIB_SRV). **[12 Marks]**

Hint: Show the distribution/replication of table fragments/partitions and explain any trade-off and pros/cons you think are relevant.

QUESTION FOUR [20 Mrks]

a) Briefly describe the various locking options available to a DBMS as part of the concurrency control function, paying particular attention to the granularity of the locking, the nature of the locks and the allowable operations and what happens when deadlock occurs. You should supply any examples or diagrams as you deem suitable *(hint: think about the performance issues of different lock types too).* **[8 Marks]**

b) Explain the distinction between the terms *serial schedule* and *serializable schedule*. **[4 Marks]**

c) Explain the usefulness of each transaction properties with the aid of examples. **[8 Marks]**

QUESTION FIVE [20 Marks]

Given the following three linked tables in which primary key columns are underlined and foreign key columns are labelled with an asterisk (*) :

Customers (custID, name, country)
Products (prodID, price)
Orders (orderID, custID*, prodID*, o_date)

and the following query:

```
SELECT Customers.name
FROM Customers, Orders, Products
WHERE Customers.custID = Orders.custID
AND Orders.prodID = Products.prodID
AND Orders.o_date = '15-Sep-2016'
AND Products.price > 100;
```

a) Suppose this query is run by executing the following sequence of steps:

1. R1 = Join of Customers and Orders
2. R2 = Join of Products and R1
3. R3 = Selection (date = '15-Sep-2016') from R2
4. R4 = Selection (price > 100) from R3

5. R5 = Projection (name) from R4

i. What is the problem caused if the query is executed based on the sequence above. **[2 Marks]**

ii. Suggest a new sequence that will make the query more efficient. You would need to introduce extra steps and not simply re-arrange the existing steps. **[8 Marks]**

iii. Suppose there is an index on the column "country" of the "Customers" table above. Explain how this index could be used when executing each of the following queries: SELECT *
FROM Customers
WHERE country = 'Kenya'; **[2 Marks]**

b) The following query returns information about all persons who have acted in a play that they have written:

```
SELECT a.name, p.title, r.character_name
FROM Actor a, Play p, Role r
Where a.name= p.author AND a.actor_id = r.actor_id AND p.play_id = r.play_id
```

Give a relational algebra query plan drawn as a tree that correctly computes this query. **[8 Marks]**