(University of Choice)

MASINDE MULIRO UNIVERSITY OF SCIENCE AND TECHNOLOGY (MMUST)

MAIN CAMPUS

UNIVERSITY EXAMINATIONS

2022/2023 ACADEMIC YEAR

THIRD YEAR FIRST SEMESTER EXAMINATIONS

FOR THE DEGREE OF BACHELOR OF COMPUTER SCIENCE

COURSE CODE: BCS 314

COURSE TITLE: DATABASE SYSTEMS II

DATE: 05/12/2022 **TIME:** 8:00-10:00AM

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) For each of the following protocols, describe aspects of practical applications that would lead you to suggest using the protocol, and aspects that would suggest not using the protocol: **[9 Marks]**
   i. Two-phase locking
   ii. Timestamp ordering
   iii. Validation

b) Explain what a Trigger is used for and how it differs from a traditional declarative integrity constraint. **(5 Marks)**

c) With relevant examples, using MMUST ERP, explain how the following differ: fragmentation transparency, replication transparency, and location transparency. [8 Marks]

d) Briefly answer the following questions based on this schema:

Emp(eid: integer, ename: string, age: integer, salary: real)

Works(eid: integer, did: integer, pct time: integer)

Dept(did: integer, budget: real, managerid: integer)

i. Suppose you have a view SeniorEmp defined as follows: [4 **Marks]**

```
CREATE VIEW SeniorEmp (sname, sage, salary)
AS SELECT E.ename, E.age, E.salary
FROM Emp E
WHERE E.age > 50
```

Explain what the system will do to process the following query:

```
SELECT S.sname
FROM SeniorEmp S
WHERE S.salary > 100,000
```

ii. Give an example of a view on Emp that could be automatically updated by updating Emp. **[4 Marks]**

QUESTION TWO

Consider the following relational schema. An employee can work in more than one department; the *pct_time* field of the Works relation shows the percentage of time that a given employee works in a given department.

Emp(*eid:* integer, *ename:* string, *age:* integer, *salary:* real)
Works(*eid:* integer, *did:* integer, *pct time:* integer)
Dept(*did:* integer, *budget:* real, *managerid:* integer)

Write integrity constraints (domain, key, foreign key, or CHECK constraints; or assertions) or triggers to ensure each of the following requirements, considered independently. **[5 Marks]**

a) The total percentage of all appointments for an employee must be under 100%.
b) A manager must always have a higher salary than any employee that he or she manages.

[5 Marks]

c) Whenever an employee is given a raise, the manager's salary must be increased to be at least as much. **[4 Marks]**
d) Whenever an employee is given a raise, the manager's salary must be increased to be at least as much. Further, whenever an employee is given a raise, the department's budget must be increased to be greater than the sum of salaries of all employees in the department. **[6 Marks]**

QUESTION THREE

a) Briefly discuss why data fragmentation could be useful in such distributed database systems, and explain each of the following fragmentation rules: **[9 Marks]**
   i. Completeness
   ii. Reconstruction
   iii. Disjointness
b) Provide any **THREE** difficulties that database designer has to face when implementing a distributed database system. **[6 Marks]**
c) The components of a Distributed Database include the following: Network Processor, Remote data processor, Data dictionary, and local data processor. Explain the function of each of the above components. **[5 Marks]**

QUESTION FOUR

a) Give a specific example of when you might be willing not to enforce serializability for a particular transaction. **[3 Marks]**
b) Why do most DBMSs allow transactions to lock a page (as opposed to just individual tuples)? **[4 Marks]**
c) Consider the following schedule where, for clarity, we have included all the operations that will appear in the log, plus additional helpful detail.

Operation	**Column 2**
T1 STARTS T1 reads item B T1 writes item B with old value 11, new value 12 T2 STARTS T2 reads item B	

T2 writes item B with old value 12, new value 13 T3 STARTS T3 reads item A T3 writes item A with old value 29, new value 30 T2 reads item A T2 writes item A with old value 30, new value 31 T2 COMMITS T1 reads item D T1 writes item D with old value 44, new value 45 *** T3 COMMITS T1 COMMITS	

i. What serial schedule is this equivalent to? If none, then explain why. **[4 Marks]**

ii. Is this schedule consistent with two phase locking? If your answer is yes, then in column 2, insert into the schedule a minimal set of additional operations that will make the schedule no longer consistent with two phase locking. Do not introduce any new transactions. Make sure to show exactly where your new operations should be inserted in the original schedule. If your answer is no, then in column 2, remove from the schedule a minimal set of operations, so that the revised schedule is consistent with two phase locking. **[5 Marks]**

iii. Consider the version of the schedule that is consistent with two phase locking (either the original one or your revised version, depending on your answer to the previous part). Is that version consistent with strict two phase locking? Why or why not? **[4 Marks]**

QUESTION FIVE

Consider the following schema:

Suppliers(*sid:* integer, *sname:* string, *address:* string)

Parts(*pid:* integer, *pname:* string, *color:* string)

Catalog(*sid:* integer, *pid:* integer, *cost:* real)

The key fields are underlined, and the domain of each field is listed after the field name. Therefore *sid* is the key for Suppliers, *pid* is the key for Parts, and *sid* and *pid* together form the key for Catalog. The Catalog relation lists the prices charged for parts by Suppliers.

Write the following queries in relational algebra:

i. Find the ***sids*** of suppliers who supply some red part and some green part. **[3 Marks]**

ii. Find the ***sids*** of suppliers who supply every red part or supply every green part. **[4 Marks]**

iii. Find pairs of ***sids*** such that the supplier with the first ***sid*** charges more for some part than the supplier with the second ***sid***. **[4 Marks]**

iv. Find the *pid*s of parts supplied by at least two different suppliers. **[4 Marks]**

v. Find the ***pids*** of the most expensive parts supplied by suppliers named Riggy G. **[5 Marks]**