ELIZADE UNIVERSITY, ILARA-MOKIN, ONDO STATE
FACULTY OF ENGINEERING
DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING

SECOND SEMESTER EXAMINATION, 2017/2018 ACADEMIC SESSION

COURSE TITLE: Introduction to Databases Design and Implementation

COURSE CODE: ECT 320

EXAMINATION DATE: 2nd AUGUST 2018

COURSE LECTURER: Dr. B. S. Afolabi

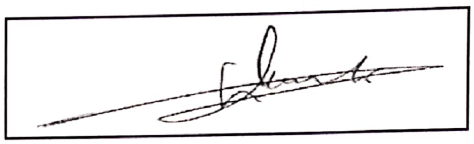

HOD's SIGNATURE

TIME ALLOWED: 2 HOURS

INSTRUCTIONS:

1. **ANSWER ALL QUESTIONS IN SECTION A AND ANY THREE QUESTIONS IN SECTION B.**
2. SEVERE PENALTIES APPLY FOR MISCONDUCT, CHEATING, POSSESSION OF UNAUTHORIZED MATERIALS DURING EXAM.
3. YOU ARE **NOT** ALLOWED TO BORROW ANY WRITING MATERIALS DURING THE EXAMINATION.

Section A (Attempt all questions in this Section, each questions carries one Mark)

1. A relational database consists of a collection of a) Tables b) Fields c) Records d) Keys
2. A _____ in a table represents a relationship among a set of values. a) Column b) Key c) Row d) Entry
3. The term _____ is used to refer to a row. a) Attribute b) Tuple c) Field d) Instance
4. The term attribute refers to a _____ of a table. a) Record b) Column c) Tuple d) Key
5. For each attribute of a relation, there is a set of permitted values, called the _____ of that attribute. a) Domain b) Relation c) Set d) Schema
6. Database _____ , which is the logical design of the database, and the database _____,which is a snapshot of the data in the database at a given instant in time. a) Instance, Schema b) Relation, Schema c) Relation, Domain d) Schema, Instance
7. Course(course_id,sec_id,semester)
   Here the course_id,sec_id and semester are _____ and course is a _____ . a) Relations, Attribute b) Attributes, Relation c) Tuple, Relation d) Tuple, Attributes
8. Department (dept name, building, budget) and Employee (employee_id , name, dept name,salary)
   Here the dept_name attribute appears in both the relations .Here using common attributes in relation schema is one way of relating _____ relations. a) Attributes of common b) Tuple of common c) Tuple of distinct d) Attributes of distinct
9. A domain is atomic if elements of the domain are considered to be _____ units. a) Different b) Indivisbile c) Constant d) Divisible
10. The tuples of the relations can be of _____ order. a) Any b) Same c) Sorted d) Constant
11. Which one of the following is a set of one or more attributes taken collectively to uniquely identify a record? a) Candidate key b) Sub key c) Super key d) Foreign key
12. Consider attributes ID, CITY and NAME. Which one of this can be considered as a super key? a) NAME b) ID c) CITY d) CITY , ID
13. The subset of super key is a candidate key under what condition? a) No proper subset is a super key b) All subsets are super keys c) Subset is a super key d) Each subset is a super key
14. A _____ is a property of the entire relation, rather than of the individual tuples in which each tuple is unique. a) Rows b) Key c) Attribute d) Fields
15. Which one of the following attribute can be taken as a primary key? a) Name b) Street c) Id d) Department
16. Which one of the following cannot be taken as a primary key? a) Id b) Register number c) Dept_id d) Street
17. An attribute in a relation is a foreign key if the _____ key from one relation is used as an attribute in that relation. a) Candidate b) Primary c) Super d) Sub
18. The relation with the attribute which is the primary key is referenced in another relation. The relation which has the attribute as primary key is called a) Referential relation b) Referencing relation c) Referenced relation d) Referred relation
19. The _____ is the one in which the primary key of one relation is used as a normal attribute in another relation. a) Referential relation b) Referencing relation c) Referenced relation d) Referred relation
20. A _____ integrity constraint requires that the values appearing in specified attributes of any tuple in the referencing relation also appear in specified attributes of at least one tuple in the referenced relation. a) Referential b) Referencing c) Specific d) Primary
21. Which of the following gives a logical structure of the database graphically? a) Entity-relationship diagram b) Entity diagram c) Database diagram d) Architectural representation
22. The entity relationship set is represented in E-R diagram as a) Double diamonds b) Undivided rectangles c) Dashed lines d) Diamond
23. The Rectangles divided into two parts represents a) Entity set b) Relationship set c) Attributes of a relationship set d) Primary key
24. Consider a directed line(->) from the relationship set advisor to both entity sets instructor and student. This indicates _____ cardinality a) One to many b) One to one c) Many to many d) Many to one
25. We indicate roles in E-R diagrams by labeling the lines that connect _____ to _____. a) Diamond , diamond b) Rectangle, diamond c) Rectangle, rectangle d) Diamond, rectangle
26. An entity set that does not have sufficient attributes to form a primary key is termed a _____. a) Strong entity set b) Variant set c) Weak entity set d) Variable set
27. For a weak entity set to be meaningful, it must be associated with another entity set, called the a) Identifying set b) Owner set c) Neighbour set d) Strong entity set
28. Weak entity set is represented as a) Underline b) Double line c) Double diamond d) Double rectangle
29. If you were collecting and storing information about your music collection, an album would be considered a(n) _____. a) Relation b) Entity c) Instance d) Attribute
30. What term is used to refer to a specific record in your music database; for instance; information stored about a specific album? a) Relation b) Instance c) Table c) Column

Section B (Attempt any THREE question from this Section)

Question 1

We consider the following relation:
Articles(ID,title,journal,issue,year,startpage,endpage,TR-ID)
It contains information on articles published in scientific journals. Each article has a unique ID, a title, and information on where to find it (name of journal, what issue, and on which pages). Also, if results of an article previously appeared in a "technical report" (TR), the ID of this technical report can be specified. We have the following information on the attributes:

- For each journal, an issue with a given number is published in a single year.
- The endpage of an article is never smaller than the startpage.
- There is never (part of) more than one article on a single page.

The following is an instance of the relation:

your E/R diagram, including primary key and foreign key constraints.(1 Mark for each conversion up to a maximum of 10 Marks)

Question 3

Suppose we have the following relational schema

Person(pid:integer, name:string, street:string, postcode:string)
Car(cid:integer, year:integer, model:string)
OwnedBy(pid:integer, cid:integer)
AccidentReport(rid:integer, damage:integer, details:string)
ParticipatedIn(pid:integer, rid:integer, cid:integer)

where the underlined attributes represent the primary keys of the associated relation. The table OwnedBy implements a relationship between persons and cars using foreign keys. The table ParticipatedIn implements a relationship between persons, accident reports, and cars, where tuple (p; r; c) indicates that the person p was the driver of the car c associated with the accident report r.

(a) Write an SQL query to return those pid's of persons driving in at least one accident, with no duplicates. (3 Marks)
(b) Write an SQL query to return all tuples (pid; c), where c is the number of cars owned by person pid (records where c = 0 do not have to be generated). (3 Marks)
(c) Write an SQL query to return all tuples (cid; c), where c is the number of persons owning car cid (records where c = 0 do not have to be generated). (2 Mark)
(d) Write a (nested) SQL query to return all tuples (pid; rid) where pid was driving in the accident reported in rid, but the car driven by pid is not owned by pid. 2 Marks
(e) Write an SQL query to return all tuples (rid; c), where c is the number of drivers involved in the accident reported in by rid (records where c = 0 do not have to be generated). 2 Marks
(f) Write an SQL query to return all tuples (rid; c), where c is the number of cars involved in the accident reported in by rid (records where c = 0 do not have to be generated). 2Mark
(g) Do the functional dependencies implied by the schema imply that the results of queries in (e) and (f) will always be the same? Explain. 3Marks
(h) Perhaps there is something wrong with this schema. How would you fix the schema to ensure that results of queries (e) and (f) would always be the same? 3Marks

Question 4

An online picture sharing company uses a database with the following schema:

```
create table Usr (
        uid int primary key,
        uname text not null,
        city text not null);
create table Picture (
        pid int primary key,
        uid int not null references Usr(uid),
        size int not null,
        pdf text);
```

Every user has a key (uid), a name (uname) and a city. Every picture has a key (pid), an author (uid) that is a foreign key to Usr, a size, and the pdf content (which is plain text).

a) Write a SQL query that retrieves all users who have only pictures of less than 1MB (size < 1000000). Your query should return the users' uid and uname (3 Marks)
b) Write a Relational Algebra expression that is equivalent to the following SQL query:

```
select distinct x.city
from usr x
where not exists (select *
                  from usr y, picture z
                  where x.uname = y.uname
                  and x.city = y.city
                  and y.uid = z.uid
                  and z.size < 3000000);
```

(4 Marks)

c) Consider the following query:

```
select x.uid, x.uname,
        (select count(*)
        from Picture y
        where x.uid = y.uid and y.size > 1000000)
from Usr x
where x.city = 'Benin;
```

ID	Title	Journal	Issue	Year	Startpage	Endpage	TR-ID
42	Cuckoo Hashing	JAlg	51	2004	121	133	87
33	Deterministic Dictionaries	JAlg	41	2001	69	85	62
33	Deterministic Dictionaries	JAlg	41	2001	69	85	56
39	Dictionaries in less space	SICOMP	31	2001	111	133	47
57	P vs NP resolved	JACM	51	2008	1	3	99
77	What Gödel missed	SICOMP	51	2008	1	5	98
78	What Gödel missed	Nature	2222	2008	22	22	98

a) Indicate for each of the following expressions whether it is a valid SQL statement or not. A valid statement, as described in GUW, should be accepted by a standard SQL interpreter, whereas an invalid statement should result in an error message. (1 Mark each)
   i. SELECT * FROM Articles WHERE endpage-startpage>10;
   ii. SELECT * FROM Articles WHERE endpage-startpage<0;
   iii. SELECT SUM(title) FROM Articles;
   iv. SELECT AVG(year) FROM Articles WHERE title LIKE 'C%';
   v. SELECT COUNT(*) FROM Articles GROUP BY year;
   vi. SELECT year,COUNT(*) FROM Articles WHERE COUNT(*)>10 GROUP BY year;

b) Indicate for each of the following queries, how many tuples would be returned if it was run on the instance of Articles (1 Mark each)
   i. SELECT ID FROM Articles WHERE year<2006;
   ii. SELECT DISTINCT ID FROM Articles WHERE year<2006;
   iii. SELECT AVG(year) FROM Articles GROUP BY journal;
   iv. SELECT ID FROM Articles WHERE title LIKE '%d';

Consider the relations Authors(auID,name) and Authoring(articleID,authorID), containing information on names of authors, and who is authoring which papers, respectively. Use this as additional information to answer the remainder of the questions

c) Write an SQL query that returns for each article, its ID, title and the number of authors. (2 Marks)
d) Write an SQL query that returns the titles of articles authored by '*Robert Tarjan*'. (2 Marks)
e) Write an SQL query that returns the number of co-authors of '*Robert Tarjan*'. (i.e., the number of authors who have written at least one article together with him.) (2 Marks)
f) Write SQL statements that correspond to the following two relational algebra expressions. Duplicate elimination should be performed.
g) $\pi_{title,year}(\sigma_{year=2005}(Articles))$ (2 Marks)
h) $\Upsilon_{year,COUNT(ID)}(Articles)$ (2 Marks)

Question 2

The academic world is an interesting example of international cooperation and exchange. This problem is concerned with modeling of a database that contains information on researchers, academic institutions, and collaborations among researchers. A researcher can either be employed as a professor or a lab assistant. There are three kinds of professors: Lecturer, Associate Professors, and Professors. The following should be stored:

- For each researcher, his/her name, year of birth, and current position (if any).
- For each institution, its name, country, and inauguration year.
- For each institution, the names of its schools (e.g. School of Law, School of Business, School of Computer Science,. . .). A school belongs to exactly one institution.
- An employment history, including information on all employments (start and end date, position, and what school).
- Information about co-authorships, i.e., which researchers have co-authored a research paper. The titles of common research papers should also be stored.
- For each researcher, information on his/her highest degree (BSc, MSc or PhD), including who was the main supervisor, and at what school.
- For each professor, information on what research projects (title, start date, and end date) he/she is involved in, and the total amount of grant money for which he/she was the main applicant.

a) Draw an E/R diagram for the data set described above. Make sure to indicate all cardinality constraints specified above. The E/R diagram should not contain redundant entity sets, relationships, or attributes. Also, use relationships whenever appropriate. If you need to make any assumptions, include them in your answer. (10 Marks)
b) Convert your E/R diagram from question a) into relations, and write SQL statements to create the relations. You may make any reasonable choice of data types. Remember to include any constraints that follow from the description of the data set or

For each query below indicate if it is correct AND equivalent to the given query. You should answer 'yes' only of the query returns exactly the same answers. (1 Mark for each correct answer)

i. Is this an equivalent query?

```
select x.uid, x.uname, count(*)
from Usr x, Picture y
where x.uid = y.uid and y.size > 1000000 and x.city = 'Benin'
group by x.uid, x.uname;
```

ii. Is this an equivalent query?

```
select x.uid, x.uname, count(*)
from Usr x, Picture y
where x.uid = y.uid and x.city = 'Benin'
group by x.uid, x.uname
having y.size > 1000000;
```

iii. Is this an equivalent query?

```
select x.uid, x.uname, count(y.pid)
from Usr x left outer join Picture y on x.uid = y.uid and y.size > 1000000
group by x.uid, x.uname, x.city
having x.city = 'Benin';
```

iv. Is this an equivalent query?

```
select x.uid, x.uname, count(*)
from Usr x left outer join Picture y on x.uid = y.uid and y.size > 1000000
group by x.uid, x.uname, x.city
having x.city = 'Benin';
```

d) Consider the following query:

```
select distinct x.uid, x.uname
from Usr x, Picture u, Picture v, Picture w
where x.uid = u.uid and x.uid = v.uid and x.uid = w.uid
      and u.size > 1000000 and v.size < 3000000 and w.size = u.size;
```

For each query below indicate if it is correct AND equivalent to the given query: (1 Mark for each correct answer)

i. Is this an equivalent query?

```
select distinct x.uid, x.uname
from Usr x, Picture u, Picture v, Picture w
where x.uid = u.uid and x.uid = v.uid and x.uid = w.uid
      and u.size > 1000000 and v.size < 3000000 and w.size > 1000000;
```

ii. Is this an equivalent query?

```
select distinct x.uid, x.uname
from Usr x, Picture u, Picture v
where x.uid = u.uid and x.uid = v.uid
      and u.size > 1000000 and v.size < 3000000;
```

iii. Is this an equivalent query?

```
select distinct x.uid, x.uname
from Usr x, Picture u, Picture w
where x.uid = u.uid and x.uid = w.uid
      and u.size > 1000000 and u.size < 3000000 and u.size = w.size;
```

iv. Is this an equivalent query?

```
select distinct x.uid, x.uname
from Usr x, Picture u, Picture v, Picture w
where x.uid = u.uid and x.uid = v.uid and x.uid = w.uid
      and u.size > 1000000 and v.size < 3000000 and w.size = v.size;
```

e) The company merges with a local adertising company in Benin, which has a database of artistic pictures. Each picture has a title, the pdf image, and may have several authors. All authors are from Benin.

```
create table Artwork (
        wid int primary key,
        title text,
        pdf text);
create table Author (
        aid int primary key,
```

wid int not null references Artwork(wid),
aname text not null);

Some users and pictures occur in both databases. Write a SQL query that finds all common user,picture pairs. Two users are considered to be the same if their names and cities are the same; two pictures are considered to be the same if their pdf's are the same. Your query should return the user name (which is the same in both databases), its two keys in Usr and Author, and the two keys of the identical p0icture in Picture and Artwork. (5 Marks)