Ahmadu Bello University, Zaria
Department of Mathematics

2009/2010 Second Semester Examination
COSC 406: Advanced Database Systems

Date: October 6, 2010 **Time Allowed**: 120 Minutes

Instructions:

1. Attempt ANY FOUR questions.
2. Write all your answers in the spaces provided on this Question Paper.

Student's Registration Number:..... **Signature**:.....

Date of Examination:.....**Time**:.....

Scores:

Question	Maximum Scores	Scores Obtained
1	20	
2	20	
3	20	
4	20	
5	20	
Total	**80**	

1. Answer both part (a) and (b) of the following questions:

a. (**10 marks**). Study each of the following statements carefully and indicate whether each is true or false. One mark each.

S/No.	Statement	True or false?
1	Each cell of relation may contain several values from the same domain.	
2	Each attribute has a distinct name.	
3	Values of an attribute are all from the same domain.	
4	Order of attributes has some significance.	
5	Physical data independence refers to immunity of the conceptual schema to changes in the internal schema.	
6	Order of tuples has no significance, theoretically.	
7	Four basic (ACID) properties of a transaction are: Automatic, Controllable, Integrity and Durability	
8	Functional Dependency describes the relationship between attributes in different relations	
9	First Normal Form (1NF) is a relation in which the intersection of each row and column contains one and only one value	
10	With NOT NULL, system rejects any attempt to insert a null in the column.	

b. (**10 marks**). Study each of the explanations in the right-hand table below and match it against the most appropriate term from the left-hand table. Write the corresponding number of the term in the left-hand table to the left of the matching explanation in the right-hand table. One mark each.

Term
1 – A relation
2 – Attribute
3 – Domain
4 – Degree
5 – Cardinality
6 – View
7 – Foreign key
8 – Entity integrity
9 – Referential integrity
10 – Base relation

Matching ID	**Definition**
	is a named column of a relation.
	is a set of allowable values for one or more attributes.
	is a number of tuples in a relation.
	is a number of attributes in a relation.
	is a table with columns and rows.
	Dynamic result of one or more relational operations operating on the base relations to produce another relation.
	An attribute or set of attributes within one relation that matches candidate key of some (possibly same) relation.
	A named relation, corresponding to an entity in conceptual schema, whose tuples are physically stored in database.
	In a base relation, no attribute of a primary key can be null.
	If foreign key exists in a relation, either the foreign key value must match a candidate key value of some tuple in its home relation or foreign key value must be wholly null.

2. (**20 marks**). Use the following relations to answer the questions that follow.

EMPLOYEE

Fname	Minit	Lname	Ssn	Bdate	Address	Sex	Salary	Super_ssn	Dno
John	B	Smith	123456789	1965-01-09	731 Fondren, Houston, TX	M	30000	333445555	5
Franklin	T	Wong	333445555	1955-12-08	638 Voss, Houston, TX	M	40000	888665555	5
Alicia	J	Zelaya	999887777	1968-01-19	3321 Castle, Spring, TX	F	25000	987654321	4
Jennifer	S	Wallace	987654321	1941-06-20	291 Berry, Bellaire, TX	F	43000	888665555	4
Ramesh	K	Narayan	666884444	1962-09-15	975 Fire Oak, Humble, TX	M	38000	333445555	5
Joyce	A	English	453453453	1972-07-31	5631 Rice, Houston, TX	F	25000	333445555	5
Ahmad	V	Jabbar	987987987	1969-03-29	980 Dallas, Houston, TX	M	25000	987654321	4
James	E	Borg	888665555	1937-11-10	450 Stone, Houston, TX	M	55000	NULL	1

DEPARTMENT

Dname	Dnumber	Mgr_ssn	Mgr_start_date
Research	5	333445555	1988-05-22
Administration	4	987654321	1995-01-01
Headquarters	1	888665555	1981-06-19

DEPT_LOCATIONS

Dnumber	Dlocation
1	Houston
4	Stafford
5	Bellaire
5	Sugarland
5	Houston

Write relational algebra expressions to answer the following queries.

a. (**4 marks**). To retrieve the first name, last name, and salary of all employees who work in department number 5.

b. (**4 marks**). To retrieve the social security numbers of all employees who either work in department 5 or directly supervise an employee who works in department 5.

c. (**4 marks**). Write the result of the following relational algebra query:

DEP_LOCS ← DEPARTMENT ∗ DEPT_LOCATIONS

Dname	Dnumber	Mgr_ssn	Mgr_start_date	Location
Headquarters	1	888665555	1981-06-19	Houston
Administration	4	987654321	1995-01-01	Stafford
Research	5	333445555	1988-05-22	Bellaire
Research	5	333445555	1988-05-22	Sugarland
Research	5	333445555	1988-05-22	Houston

d. (**4 marks**). Write the result of the following relational algebra query:

$$\rho_{R(Dno,NoOfEmployees,AvaerageSal)}({}_{Dno}\Im_{COUNT\ Ssn,AVERAGE\ Salary}(EMPLOYEE))$$

e. (**4 marks**). Given R and S below, show the result of: $T \leftarrow R \div S$

A	B
a1	b1
a1	b2
a1	b3
a1	b4
a2	b1
a2	b3
a2	b4
a3	b1
a3	b2
a3	b4
a4	b1
a4	b3

A
a1
a2
a4

3. (**20 marks**). Consider the following information about a university database:

- Professors have an SSN, a name, an age, a rank, and a research specialty.
- Projects have a project number, a sponsor name (e.g., NSF), a starting date, an ending date, and a budget.
- Graduate students have an SSN, a name, an age, and a degree program (e.g., M.S. or Ph.D.).
- Each project is managed by one professor (known as the project's principal investigator).
- Each project is worked on by one or more professors (known as the project's co-investigators).
- Professors can manage and/or work on multiple projects.
- Each project is worked on by one or more graduate students (known as the project's research assistants).
- When graduate students work on a project, a professor must supervise their work on the project. Graduate students can work on multiple projects, in which case they will have a (potentially different) supervisor for each one.
- Departments have a department number, a department name, and a main office.
- Departments have a professor (known as the chairman) who runs the department.
- Professors work in one or more departments, and for each department that they work in, a time percentage is associated with their job.
- Graduate students have one major department in which they are working on their degree.
- Each graduate student has another, more senior graduate student (known as a student advisor) who advises him or her on what courses to take.

Design and draw an ER diagram that captures the information about the university. Use only the basic ER model here; that is, entities, relationships, and attributes. Be sure to indicate any key and participation constraints.

4. Study each of the following schema diagrams carefully.

a. (**10 marks**). Map the following ER schema into a relational schema indicating all primary keys and foreign keys.

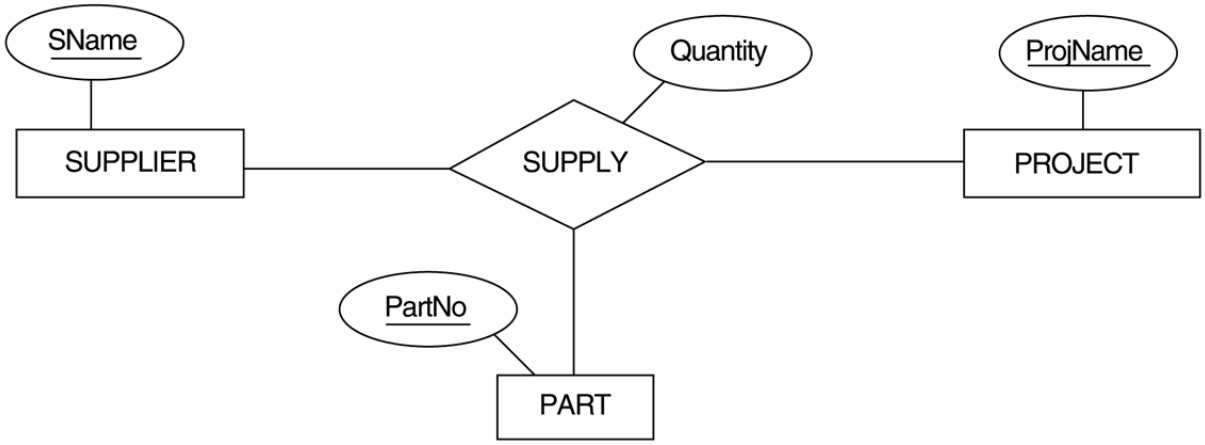

b. (**10 marks**). Map the following EER schema into a relational schema indicating all primary keys and foreign keys.

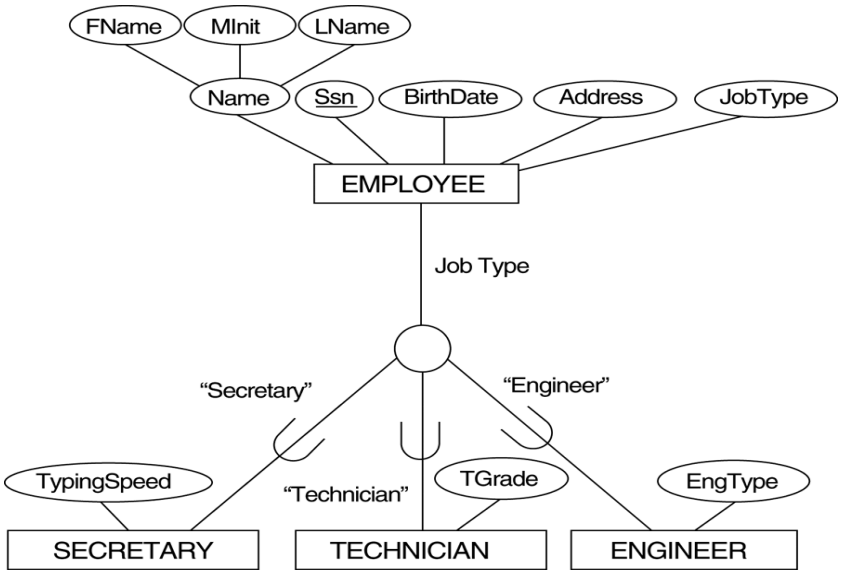

5. (**20 marks**). Study the following questions carefully and select the most appropriate option. Write your answer in the table provided below. Each question is worth two marks.

i	ii	iii	iv	v	vi	vii	viii	ix	x

i. A given relation is known to be in third normal form. Select the statement which can be inferred from this:
   A. All attributes contribute to the primary key
   B. Each non-key attribute determines the primary key
   C. Each non-key attribute is determined by the primary key
   D. Every determinant is a candidate key
   E. The relation is not in fourth normal form.

ii. There are two relations X and Y. Relation X has arity 1 and cardinality 2, relation Y has arity 3 and cardinality 4. Indicate the result of the SQL statement SELECT COUNT(*) FROM X,Y.
   A. 4
   B. 6
   C. 8
   D. Cannot be calculated from the information given
   E. None of the above.

iii. Which of the following is generally a benefit of normalisation?
   A. Performance is improved
   B. Insertion anomalies are avoided
   C. Selection anomalies are avoided
   D. Number of tables is reduced
   E. None of the above.

iv. A lack of normalisation can lead to which one of the following problems
   A. Lost Updates
   B. Deletion of data
   C. Insertion problems
   D. Deferred updates
   E. Deadlock

v. To transform a relation from first normal form to second normal form we must remove which one of the following?
   A. All partal-key dependencies
   B. All inverse partial-key dependencies
   C. All repeating groups
   D. All transitive dependencies
   E. None of the above

vi. To transform a relation from second normal form to third normal form we must remove which one of the following?
   A. All partal-key dependencies
   B. All inverse partial-key dependencies
   C. All repeating groups
   D. All transitive dependencies
   E. None of the above

vii. Each of the following is an argument which might be used to support the use of relations which are not fully normalised. Select the **weakest** argument.

A. A fully normalised database may perform too slowly
B. Full normalisation may compromise existing applications/systems
C. A fully normalised database may have too many tables
D. Full normalisation may make some queries too complicated
E. A fully normalised database may result in tables which are too large

viii. Consider the following functional dependencies

a,b => c,d e,g,h => f,j
a,c => b,d p,q => r,s
e,f,g => h,i s => t
f,g => j q => u
g,h => i

Which of the following relational schemas might be the result of normalising R(s,q,t,u)?

A. The schema R1(s,q) R2(s,t) R3(q,u)
B. The schema R1(s,q) R2(q,t) R3(t,u)
C. The schema R1(s,q) R2(s,t) R3(q,u)
D. The schema R1(s,q,t) R2(s,q,u)
E. The schema R(s,q,t,u)

ix. Consider the following functional dependencies

a,b => c,d e,g,h => f,j
a,c => b,d p,q => r,s
e,f,g => h,i s => t
f,g => j q => u
g,h => i

Which of the following relational schemas might be the result of normalising R(a,b,c,d)?

A. The schema R1(a,b) R2(a,c) R3(a,d)
B. The schema R1(a,b) R2(b,c) R3(c,d)
C. The schema R1(a,b) R2(a,c) R3(b,d)
D. The schema R1(a,b,c) R2(a,b,d)
E. The schema R(a,b,c,d)

x. Consider the following functional dependencies

a,b => c,d e,g,h => f,j
a,c => b,d p,q => r,s
e,f,g => h,i s => t
f,g => j q => u
g,h => i

Which of the following best describes the relation R(e,f,g,h,i,j)?

A. First Normal Form
B. Second Normal Form
C. Third Normal Form
D. Forth Normal Form
E. Boyce Codd Normal Form