Professor: Mario Nascimento

Student name:_____ **Student ID:** _____

General Guidelines: *Exam duration: 120 min. Closed book, and no collaboration. All questions need to be answered. Use the spaces after the questions in this set of pages for your answers (additional pages may be used - write your name and id on them - but should not be necessary). Marked exams will be available, at the instructor's office, on Dec 18. Deadline for appeals is Dec 18 at 5 pm.*

Question 1 (10 points)

Consider the ER model below. Show the corresponding relational model, after the ER-relational mapping, be sure to indicate all primary and foreign keys constraints.

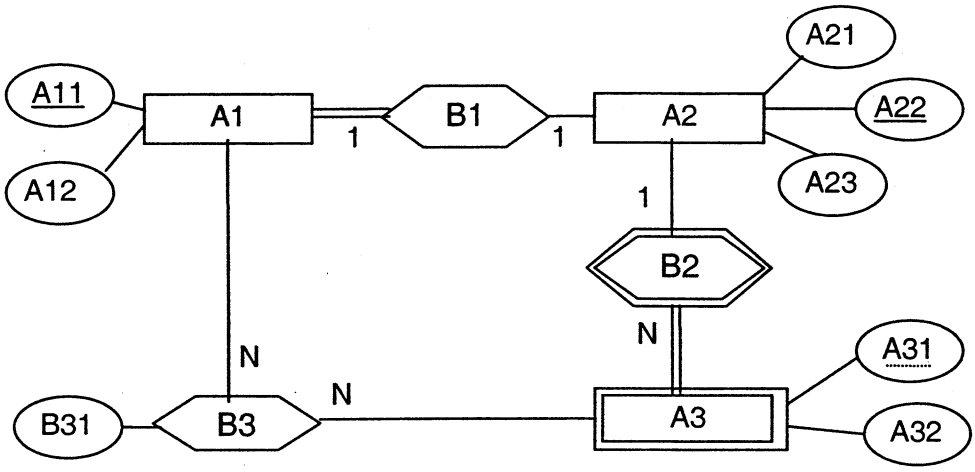

Question 2 (6 points - each item worth 3 points if your answer is correct, or -3 points if your answer is incorrect. No penalty for not answering)

Considering only the original ER model shown in Question 1:

a) The values of attribute A31 are unique, similarly to the way a primary key is unique
( ) TRUE ( ) FALSE

b) If an entity in A3 is deleted then at least one entity in A2 is also deleted
( ) TRUE ( ) FALSE

Student name:_____Student ID: _____

Question 3 (20 points)

Consider the following relations (only denoted primary keys are necessarily not null):

Employee (name, sin, address, salary, supervsin, dno)
dno refers to Department.dnumber
supervsin refers to Employee.sin
Department (dnumber, mgrsin, dname)
mgrsin refers to Employee.sin
Project (pname, pnumber, plocation, dnum)
dnum refers to Department.dnumber
WorksOn (esin, pno, hours)
esin refers to Employee.sin
pno refers to Project.pnumber

Explain in plain English exactly what the following SQL statements are doing (do not explain how they are processed).

Example: *SELECT COUNT(*)*
FROM Employee, Department
WHERE dno = dnumber AND dname = 'Research'
Good answer: This query retrieves the number of employees in the Research dept.
Bad answer: This query joins the table Employee and Department based on the attributes dno and dnumber and selects the tuples where dname is "Research", then it counts the number of selected tuples

a)
```
(SELECT pnumber
FROM Project, Department, Employee
WHERE dnum = dnumber AND mgrsin = sin AND name = "Smith")
UNION
(SELECT pnumber
FROM Project, WorksOn, Employee
WHERE pnumber = pno AND esin = sin AND name = "Smith")
```

b)
```
SELECT COUNT (DISTINCT salary)
FROM employee
```

c)
```
SELECT name
FROM Employee
WHERE ((SELECT pno
        FROM WorksOn
        WHERE sin = esin)
       CONTAINS
       (SELECT pnumber
        FROM Project
        WHERE dnum = 5))
```

d)
```
SELECT E.name, S.name
FROM Employee E LEFT OUTER JOIN Employee S ON E.supervsin = S.sin
```

Question 4 (25 points)

Considering the same relations for Question 3, write SQL statements corresponding to the following queries:

a) Retrieve the SIN of all employees who work in project number 1, 2 or 3

b) Find the names of all employees whose salary is greater than the salary of all employees in department 5.

c) Show the sum of all salaries of employees working on the 'PX' project if every such employee was given a 10% raise

d) List all projects which have no employee assigned to it.

e) For each department having more than 5 employees, retrieve the department number and the number of employees makring more than $40,000

Question 5 (9 points - each item worth 3 points if your answer is correct, or -3 points if your answer is incorrect. No penalty for not answering)

a) When using hashing tables with a bucket scheme, overflows are avoided
( ) TRUE ( ) FALSE

b) B+-trees are very likely to be better than Hashing tables for range queries
( ) TRUE ( ) FALSE

c) ISAM trees as well as B+-trees have guaranteed logarithmically search time
( ) TRUE ( ) FALSE

Student name:_____Student ID: _____

Question 6 (10 points)

Assume a hypothetical disk with the following specification:

rotational speed: R rps
seek time: S secs
transfer time: T blocks/sec
Number of used blocks per file: U
Number of blocks per track: B

Assume further that:

U < B
All files have the same size
The order in which the records are read does not matter

Stating your assumptions clearly, show formulas for the minimum time required to read all records stored in a single file in this disk in the cases where:

a) the disk has a single surface and a single read/write head
b) the disk has F surfaces and F read/write heads (F > U) and is capable of performing parallel I/O, i.e., all heads can read/write at the same time

Student name:_____Student ID: _____

Question 7 (10 points)

Consider the following set of keys: { 1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21 }. Show the resulting B+-tree, of order 2, when all the key values are inserted in ascending order. Does the resulting B+-tree configuration depend on the key insertion order ? Justify your answer.

Question 8 (5 points)

Consider the relation R = { A, B, C, D, E } and the following functional dependencies: FD1 = (B, C $\rightarrow$ A, D, E) and FD2 = (C $\rightarrow$ E). What is the key for relation R ? Is R in second normal form ? Why ?

Question 9 (5 points)

Discuss the following statement: "relational algebra and SQL are equivalent in the sense that any query expressed in SQL can be expressed in relational algebra and vice-versa".