COSC 3480: Final Exam (Sample Questions)

Name: _____

D. Mirkovic, Fall 2004

Maximum number of points possible: 100. This exam counts for 30% of your overall grade. Questions vary in difficulty, and the number of points associated with a question is not necessarily an indication of the difficulty of the question. Do not get stuck on one question. This is a collection of typical questions not an actual exam.

1. (5) Explain briefly why would you choose a database system instead of simply storing data in operating system files? When would it make sense not to use a database system?

2. (5) What is a transaction?

3. (5) Explain the following terms briefly: attribute, domain, relationship.

4. (5) Explain one-to-many relationship.

5. (5) What is the primary key?

6. (15) Consider the following schema:

   Books(*bid:* integer, *author:* string, *title:* string, *year:* integer)

   a) (5) Write an SQL statement required to create the Books relation including the appropriate primary key constraint.

b) (5) Write an SQL command to insert a tuple (1234, 'John Grisham', 'The Brethren', 2000).

c) (5) Write an SQL command to delete all books printed before 1900.

d) (5) Write an SQL statement change the name of the author to 'Virginia Woolf' in all books written by 'Virginia Stephen'.

7. (15) Answer each of the following questions briefly. The questions are based on the following relational schema:

Faculty(*fid:* integer, *fname:* string, *dept:* string, *salary:* real)
Class(*cid:* integer, *cname:* string, *time:* string, *room:* string, *teacherid* integer)

a) (10) Write the SQL statements required to create the above relations, including appropriate definitions of all primary and foreign key integrity constraints.

b) (10) Which of the following database modifications are legal?

i. `INSERT`
`INTO` Faculty (fid, fname, dept, salary)
`VALUES` (1234, 'John', 'Art', 4323.2)

ii. `INSERT`
`INTO` Faculty (fid, fname, dept, salary)
`VALUES` (null, 'John', 'Art', 4323.2)

iii. `INSERT`
`INTO` Class (cid, cname, time, room, teacherid)
`VALUES` (301, 'Music 101', 'TTH 4', 'PGH 303', null)

iv. `UPDATE` Faculty F
`SET` F.dept=1234
`WHERE` F.name='Mike'

v. `UPDATE` Class C
`SET` C.room='PGH 201'
`WHERE` C.teacherid=1234

8. (15) Consider the following schema:

Aircraft(*aid:* integer, *aname:* string, *cruisingrange:* integer)
Certified(*eid:* integer, *aid:* integer)
Employees(*eid:* integer, *ename:* string, *salary:* integer)

The key fields are underlined, and the domain of each field is listed after the field name. Write the following queries in Relatinal Algebra:

a) (5) Find the *enames* of pilots certified for some Boeing aircraft.

b) (5) Find the *enames* of pilots who can operate planes with a range greater than 3,000 miles but are not certified on any Boeing aircraft.

c) (5) Find the *eids* and *enames* of pilots certified for all aircrafts.

d) (5) Find the *enames* of pilots who are certified on Boeing and Airbus aircrafts.

9. (20) Write the following queries in SQL using the schema from problem 4.

a) Find the *enames* of pilots certified for some Airbus aircraft.

_____ (5)

b) Find the *eids* of employees who make the highest salary.

_____ (5)

c) Find the *eids* of employees who make the second highest salary.

_____ (5)

d) Find the *eids* of employees who are certified for the largest number of aircraft.

_____ (5)

10. (10) Assume we have two relations R(A,B) and S(B,C). All three attributes (A, B, and C) are integer attributes. Assume that Relation R contains the following tuples: (1,2), (2,3), and (3,4). Assume that Relation S contains the following tuples (2,2), (2,3) and (5,1).

a) Is A a key for R. Answer (yes/no): _____ (5)

b) How many tuples are in the result of the cross-product between R and S? Answer (a number): _____ (5)

c) (5) How many tuples are in the result of the natural join between R and S? Answer (a number): _____ (5)

11. (15) Your Vet has decided to store information about his patients and their owners in a database. He has wisely chosen to hire you as a database designer. The data he would like to store is as follows.

a) Each pet has a name, a species, age and belongs to only one customer.

b) Each customer has an SSN, a name, an address, and a phone number and can own more then one pet.

c) A pet must be identified uniquely by its name and the SSN of its owner.

a) Design a conceptual schema for the data above and draw an ER diagram for your schema. Be sure to indicate all key and cardinality constraints and any assumptions you make.

b) (5) Write SQL statements to create the relations corresponding to the above ER diagram and capture as many of the constraints as possible.

12. (10) Write a C embedded SQL code that will extract the *pname* and *color* of all parts supplied by a supplier named 'GWC' based on 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 names and colors should be stored in a cursor, which should be open, read from in a loop, and closed after all data from the cursor is processed and printed to the standard output. Make sure you declare all variables appropriately. Use the following template to formulate your answer.

```
char SQLSTATE[6];

EXEC SQL _____________________________________________  (1)

char pname[20], color[40];

_____________________________________________________  (1)

EXEC SQL _____________________________________________  (1)

   SELECT ____________________________________________  (1)

   FROM _____________________________________________  (1)

   WHERE ____________________________________________  (1)

_____________________________________________________  (1)

do{

   EXEC SQL _________________________________________  (1)

   printf(_____________________________________);

}while(_____________________________!='020000');      (1)

_____________________________________________________  (1)
```

13. (5) What is a stateless protocol? Give one example of such protocol.

14. (5) Where are stored procedures executed? How can you write stored procedures?

15. (5) Give one example of the three-tier architecture. Describe the tasks performed by the different tiers.

16. (5) What is an index on a file of records? What is a search key for an index? Why do we need indexes?

17. (15) Consider the following schema with the Sailors relation:

Sailors(*sid:* integer, *sname:* string, *rating:* integer, *age:* real)

For each of the following indexes, list whether the index matches the given selection conditions. If there is a match, list the primary conjuncts.

1. A B+-tree index on the search key $< Sailors.sid >$.

   a) $\sigma_{Sailors.sid<20,000}$(Sailors)

   _____ (2)

2. A hash index on the search key $< Sailors.sid >$.

   a) $\sigma_{Sailors.sid<20,000}$(Sailors)

   _____ (2)

   b) $\sigma_{Sailors.sid=20,000}$(Sailors)

   _____ (2)

3. A B+-tree index on the search key $< Sailors.sid, Sailors.age >$.

   a) $\sigma_{Sailors.age=21}$(Sailors)

   _____ (2)

   b) $\sigma_{Sailors.sid=20,000}$(Sailors)

   _____ (2)

18. (5) Consider a relation stored as a randomly ordered file for which the only index is an unclustered index on a field called *sal*. If you want to retrieve all records with *sal ¿ 20*, is using the index always the best alternative? Explain.

19. (5) What is an access path? When does an index match an access path? What is a primary conjunct, and why is it important?

20. (5) Describe Strict 2PL.

21. (5) Define *support* and *confidence* for an association rule.