Name:_____

CSE 444, Fall 2010, Midterm Examination
10 November 2010

Rules:

- Open books and open notes.
- No laptops or other mobile devices.
- Please write clearly.
- Relax! You are here to learn.

Question	Max	Grade
1	30	
2	30	
3	40	
Total	100	

Initials:_____

1. (**30** points) **SQL**

   Consider a database with the following three relations.

```
Sensor( sid, type, location, description )
Temperature ( sid, time, tempvalue )
Pressure ( sid, time, pressvalue )
```

   `Temperature.sid` is a foreign key that references `Sensor.sid`.

   `Pressure.sid` is a foreign key that references `Sensor.sid`.

   This database holds a set of sensor readings (temperature and pressure) recorded by a set of sensors deployed at various locations. Some sensors record only temperature readings. Others record only pressure readings. Some record both. Each sensor produces a reading every minute (pressure reading, temperature reading, or both). Each sensor is associated with a given type, location, and description.

   (a) (**8** points) Write a SQL query that returns the type and description of all sensors that have produced both temperature readings and pressure readings. Each unique combination of type and description should appear only once in the output.

Initials:_____

(b) (**8** points) Write a SQL query that computes the highest ever pressure reading among all sensors that recorded both pressure and temperature readings.

Initials:_____

(c) (**14** points) Consider the set of sensors that record BOTH temperature and pressure (your query will have to compute that set). Write a SQL query that computes, for these sensors only, the `sid` and average `tempvalue` for all sensors that recorded the highest ever pressure reading. Note that more than one sensor could have recorded the same, highest pressure reading.

Initials:_____

2. (**30** points) **Conceptual Design**

   (a) (**12** points) Draw an E/R diagram describing the following domain:

   - **Users** have a single attribute **uid** (key).
   - **AnonymousUsers** are a type of **Users** with attribute **country**.
   - **NamedUsers** are a type of **Users** with attributes **name** and **email**.
   - **Blogs** have attributes **bid**(key), **title**, and **topic**.
   - A **NamedUser** can be a **friend-of** zero or more other **NamedUsers**. Each friend-of relationship has an associated start **date**.
   - A **NamedUser** can **own** many **Blogs**, but each blog is owned by exactly one user.
   - A **User** can **post** to zero or more **Blogs**. Each post has a given **content**. Many users can post to a blog.

   Your answer should consist of an E/R diagram, which includes entity sets, attributes, relationships, ISA relations. Indicate the type of each relationship with appropriate arrows (one-one, one-many, etc.).

Initials:_____

(b) (**8** points) Consider the following relational schema and set of functional dependencies. (a) List **all** superkey(s) for this relation. (b) Which of these superkeys form a key (i.e., a minimal superkey) for this relation? Justify your answer in terms of functional dependencies and closures.

R(A,B,C,D,E) with functional dependencies CD $\rightarrow$ E and A $\rightarrow$ B.

Initials:_____

(c) (**10** points) Decompose R into BCNF. Show your work for partial credit. Your answer should consist of a list of table names and attributes and an indication of the keys in each table (underlined attributes).

Initials:_____

3. (**40** points) **Transactions**

   (a) (**6** points) In the ARIES method, assuming NO checkpoints have been used, explain what happens during the ANALYSIS pass of recovery. Your answer should indicate at least (1) what part of the log the system reads and in what direction and (2) what data structures the system rebuilds.

   (b) (**6** points) After the analysis pass, the protocol performs a REDO pass. Explain (1) where does REDO start reading the log and in what direction it reads the log, (2) what happens during the REDO pass.

Initials:_____

(c) (**6** points) The last pass during recovery is the UNDO pass. Explain (1) where does the UNDO start reading the log and in what direction it reads the log, (2) what happens during the UNDO pass.

Initials:_____

(d) (**7** points) Consider the following two transactions and schedule (time goes from top to bottom). Is this schedule conflict-serializable? Explain why or why not.

Transaction T_0	Transaction T_1
$r_0[A]$	
$w_0[A]$	
	$r_1[A]$
	$r_1[B]$
	c_1
$r_0[B]$	
$w_0[B]$	
c_0	

Initials:_____

(e) (**7** points) Show how 2PL can ensure a conflict-serializable schedule for the same transactions above. Use the notation $L_i[A]$ to indicate that transaction i acquires the lock on element A and $U_i[A]$ to indicate that transaction i releases its lock on A.

Initials:_____

(f) (**8** points) Show how the use of locks without 2PL can lead to a schedule that is NOT conflict serializable.