

Name: \_\_\_\_\_

**CSE 444, Spring 2009, Midterm Examination**  
**11 May 2009**

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	40	
2	30	
3	30	
Total	100	

Initials: \_\_\_\_\_

1. (40 points) **SQL**

Consider a database with the following four base relations and view. For the view, we give the SQL statement that created it.

Customer		
<u>cid</u>	name	address
1	Anna	Capitol Hill
2	Bob	University
3	Charlie	View Ridge

Owns	
<u>cid</u>	<u>aid</u>
1	1
1	2
2	1
2	2
3	3
3	4

Account		
<u>aid</u>	amount	type
1	900	checking
2	2000	savings
3	900	checking
4	100	savings

ActivityLog					
<u>opid</u>	<u>cid</u>	optype	<u>aid1</u>	<u>aid2</u>	amount
101	1	deposit	1	NULL	100
102	2	withdrawal	1	NULL	200
103	3	transfer	3	4	100

Owns.cid is a foreign key that references Customer.cid.

Owns.aid is a foreign key that references Account.aid.

ActivityLog.cid is a foreign key that references Customer.cid.

ActivityLog.aid1 and ActivityLog.aid2 are foreign keys that reference Account.aid.

```
CREATE VIEW GoldenCustomer AS
  SELECT DISTINCT C.cid, C.name
  FROM Customer C, Owns O, Account A
  WHERE C.cid = O.cid AND O.aid = A.aid AND A.amount > 1000
```

Initials: \_\_\_\_\_

- (a) (15 points) What is the output of the following query? Please draw the relation instance produced by this query in the form of a table. Include the table header to clearly show the schema of the relation. The order of the tuples in the table does not matter.

```
SELECT count(*) as total_activity
FROM Account A2, ActivityLog L2
WHERE A2.aid = L2.aid1 OR A2.aid = L2.aid2
GROUP BY A2.aid
```

Initials: \_\_\_\_\_

- (b) (15 points) What is the output of the following query? Notice that this query uses the previous one as a subquery.

```
SELECT C.name
FROM   Account A, ActivityLog L, Customer C, Owns O
WHERE  (A.aid = L.aid1 OR A.aid = L.aid2)
AND    A.aid = O.aid AND O.cid = C.cid
GROUP BY C.name, A.aid
HAVING count(*) >=
        ALL (SELECT count(*)
              FROM   Account A2, ActivityLog L2
              WHERE  A2.aid = L2.aid1 OR A2.aid = L2.aid2
              GROUP BY A2.aid)
```

Initials: \_\_\_\_\_

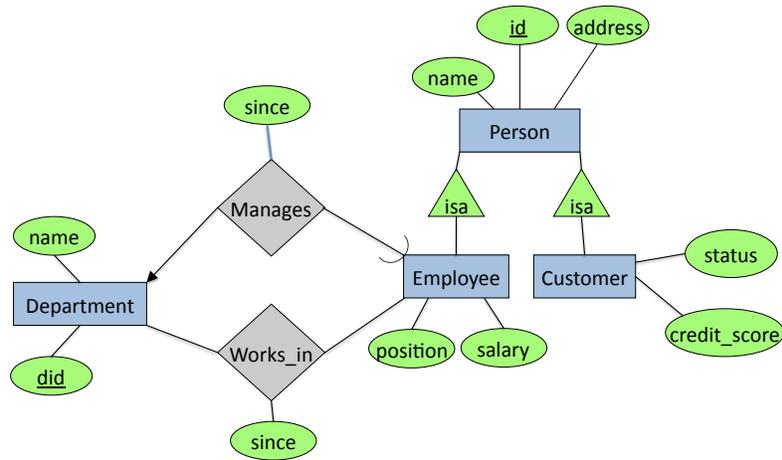
- (c) (10 points) What is the output of the following query? For this query, you can show the attributes in any order. If two attributes have the same name, indicate the relation where they come from.

```
SELECT *
FROM GoldenCustomer G, Owns O
WHERE G.cid = O.cid
```

Initials: \_\_\_\_\_

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

- (a) (10 points) Consider the following E/R diagram. Show a valid conversion of this diagram into relations. For each relation, indicate the **primary key** (you can simply underline the attributes that form the primary key) and any **foreign keys**. Note: you do NOT need to show any SQL statements. You only need to show the schema of the relations. You do NOT need to specify the types of the attributes, just their names.



Initials: \_\_\_\_\_

- (b) (10 points) If a relation is in 1NF but not in BCNF (nor in 3NF), what type of problem (or anomaly) can occur? Please **name one** anomaly and provide a **one to two sentence description** of this anomaly.

- (c) (10 points) Consider the following relational schema and set of functional dependencies. Decompose R into BCNF. Show your work for partial credit.

R(A,B,C,D,E,F,G,H) with functional dependencies  $A \rightarrow BG$ ,  $C \rightarrow D$ , and  $EF \rightarrow CH$ .

Initials: \_\_\_\_\_

3. (30 points) **Transactions**

After a system crash, the **redo-log** using non-quiescent checkpointing contains the following data:

```
< START T1 >
< T1, A, 10 >
< START T2 >
< T2, B, 5 >
< T1, C, 7 >
< START T3 >
< T3, D, 12 >
< COMMIT T1 >
< START CKPT ???? >
< START T4 >
< T2, E, 5 >
< COMMIT T2 >
< T3, F, 1 >
< T4, G, 15 >
< END CKPT >
< COMMIT T3 >
< START T5 >
< T5, H, 3 >
< START CKPT ???? >
< COMMIT T5 >
```

- (a) (10 points) What are the correct values of the two <START CKPT ????> records? You have to provide two correct values for the two ????s.

First START CKPT: \_\_\_\_\_

Second START CKPT: \_\_\_\_\_

Initials: \_\_\_\_\_

(b) (10 points) Indicate and **explain** what fragment of the log the recovery manager needs to read.

(c) (10 points) Assuming that the two < START CKPT ??? > records are correctly stored in the log, according to your answer above, show which elements are recovered by the redo recovery manager and compute their values after recovery.