CSE 344 Midterm

Monday, February 6, 2012, 9:30-10:20

Name: _____

This is a closed book exam. You have 50'. Please write your answers in the space provided.

Question	Points	Score
1	40	
2	40	
3	20	
Total:	100	

1 SQL

1. (40 points)

   Consider a photo sharing Website, where users can post pictures, as well as comment and rate other user's pictures. The schema is:

   ```
   Users(uid, name)
   Comment(uid, pid, score, txt)
   Picture(pid, author, img)
   ```

   The database has the following constraints:

   - `Comment.uid` is a foreign key to `Users`

   - `Comment.pid` is a foreign key to `Picture`

   - `Picture.author` is a foreign key to `Users`

   - `Comment.score` is an integer number between 1 and 10

   - All attributes are `NOT NULL`

   (a) (10 points) Write a SQL query that returns all users who have given a score of 8 or higher to 50 pictures or more. For each user, your query should return the user ID and the name.

```
Users(uid, name)
Comment(uid, pid, score, txt)
Picture(pid, author, img)
```

(b) (10 points) A picture is considered *highly rated* if it received at least one score of 10, from a user other than its author. A *cautious* user is a user who commented only on highly rated pictures. (A user who did not comment at all is also cautious.) Write a SQL query that finds all cautious users. Your query should return a list of `uid, name` pairs.

```
Users(uid, name)
Comment(uid, pid, score, txt)
Picture(pid, author, img)
```

(c) (10 points) A hacker found a way to break into the system and ran the following commands:

```
insert into Comment(uid, pid, score, txt)
  select x.uid, y.pid, 0 as score, 'worst picture i ever saw' as txt
  from Users x, Picture y
  where x.uid = y.author;

update Picture
set author = (select x.uid
              from Comment x
              where x.pid = Picture.pid
              order by x.score desc
              limit 1);
```

That is, he inserted some fake comments, and messed up all the picture authors. Luckily, his two queries were logged by the system (that' why we know them) and no other user activities were performed on the database during or after the breakin. As usual, the hacker was quickly caught, arrested, and sent to jail, but he refused to cooperate in repairing the database. Your job is to repair the database. Assume that the database was in a consistent state right before the hacker's attack, and write SQL commands (one or more) that will restore the database to its original state. Your answer should consists of a sequence of INSERT/UPDATE/DELETE statements, and should not rely on any log or backup of the database. For example, you might turn in something like this (which are correct SQL statement, but are not the real answer):

```
delete from Comment
where exists (select * from Users x where x.uid = uid and x.name='Hacker Joe')

update Picture
set author = (select x.uid from Comment x
              where pid = x.pid and x.score = 10 limit 1)
```

```
Users(uid, name)
Comment(uid, pid, score, txt)
Picture(pid, author, img)
```

(d) (10 points) For each statement below, indicate whether it is true of false.

   i. An index may help a select-from-where SQL query run faster, or may not affect its running time, but it can never make a query run slower.

   i. _____

   True or false?

   ii. An index may help an update (insert, delete, or update) SQL query run faster, or may not affect its running time, but it can never make a query run slower.

   ii. _____

   True or false?

   iii. Consider a selection operation $\sigma_{\text{price}>90 \wedge \text{price}<100}(\text{Product})$. Using an <u>unclustered</u> index on price will make the query at least as fast as scanning the entire table Product.

   iii. _____

   True or false?

   iv. Consider a selection operation $\sigma_{\text{price}>90 \wedge \text{price}<100}(\text{Product})$. Using a <u>clustered</u> index on price will make the query at least as fast as scanning the entire table Product.

   iv. _____

   True or false?

   v. A large table Product(pid, name, price) is queried intensively and is never updated. Then we should create three clustered indexes, on Product(pid), Product(name), and Product(price).

   v. _____

   True or false?

2 Relational Algebra and Relational Calculus

2. (40 points)

```
Users(uid, name)
Comment(uid, pid, score, txt)
Picture(pid, author, img)
```

Consider the same database schema as in the first question:

```
Users(uid, name)
Comment(uid, pid, score, txt)
Picture(pid, author, img)
```

(a) (10 points) Write a Relational Algebra expression that is equivalent to the SQL query below:

```
select distinct u.uid
from Users u, Picture x, Comment y
where u.uid = x.author and x.pid = y.pid and y.score > 8
group by u.uid, x.pid
having count(*) > 10
```
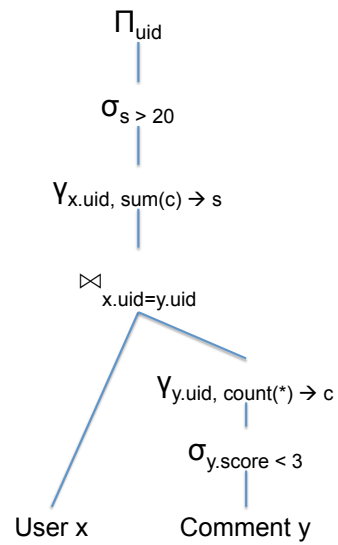
```
Users(uid, name)
Comment(uid, pid, score, txt)
Picture(pid, author, img)
```

(b) (10 points) Write a Relational Algebra expression that is equivalent to the SQL query below:

```
select x.pid
from picture x
where not exists
    (select *
     from comment y
     where x.pid = y.pid and y.score <5)
```

```
Users(uid, name)
Comment(uid, pid, score, txt)
Picture(pid, author, img)
```

(c) (10 points) Consider the Relational Algebra expression below:

$$\Pi_{uid}$$

$$\sigma_{s > 20}$$

$$\gamma_{x.uid, \, sum(c) \to s}$$

$$\bowtie_{x.uid=y.uid}$$

$$\gamma_{y.uid, \, count(*) \to c}$$

$$\sigma_{y.score < 3}$$

User x   Comment y

Write an equivalent SQL query *without* using any subqueries.

```
Users(uid, name)
Comment(uid, pid, score, txt)
Picture(pid, author, img)
```

(d) (10 points) Consider the following query in the relational calculus:

$$Q(u, n) = \texttt{Users}(u, n) \wedge \forall p. \forall i. (\texttt{Picture}(p, u, i) \Rightarrow \forall v. \forall s. \forall t. (\texttt{Comment}(v, p, s, t) \Rightarrow s > 5))$$

Write an equivalent expression that does not use a universal quantifier.

3 XML and XPath

3. (20 points)

   (a) (10 points) Consider the XML document below:

```
<a>
  <a>
    <a> 1 </a>
    <a> 2 </a>
  </a>
  <a>
    <a> 3 </a>
    <a> 4 </a>
  </a>
</a>
```

For each of the XPath queries below, indicate how many elements the expression returns:

    i. `//a`

                                                        i. _____

    Answer

    ii. `//a[a/text()=3]/a`

                                                        ii. _____

    Answer

    iii. `//a[a/a/text()=3]/a/a`

                                                        iii. _____

    Answer

    iv. `//a[a/text()=2][a/text()=3]`

                                                        iv. _____

    Answer

    v. `//a[a/a/text()=2][a/a/text()=3]`

                                                        v. _____

    Answer

(b) (10 points) For each statement below indicate whether it is true of false.

  i. XML databases are replacing relational databases in today's enterprise applications.

                                                          i. _____

True or false?

  ii. Compared to relational data, XML data is easier to exchange between applications or between users.

                                                        ii. _____

True or false?

  iii. The XPath expressions `/a/b[@c='5'][@d='2']` and `/a/b[@d='2'][@c='5']` are equivalent.

                                                       iii. _____

True or false?

  iv. The XPath expressions `/a/b[@c='5'][2]` and `/a/b[2][@c='5']` are equivalent.

                                                       iv. _____

True or false?

  v. The XPath expressions `/a/*//b` and `/a//*/b` are equivalent.

                                                       v. _____

True or false?