

CSE 414 Final Exam

March 17, 2014

Name _____

Question 1	/ 21
Question 2	/ 9
Question 3	/ 12
Question 4	/ 12
Question 5	/ 20
Question 6	/ 15
Question 7	/ 18
Question 8	/ 16
Question 9	/ 12
Total	/ 135

The exam is open textbook but otherwise you may not use any notes, computers, electronics devices, phones of the smart or not-so-smart variety, telegraphs, telepathy, smoke signals, magic spells, or any other contraptions.

The exam lasts 110 min. Please budget your time so you get to all questions.

Please wait to turn the page until everyone has their exam and you are told to begin.

Relax. You are here to learn.

Question 1. (21 points, 7 each) SQL. We have a small database with information about people and events.

Person(pid, name)
Event(eid, name, start_time, end_time)
Invited(pid, eid, going)

Where:

- *Person* contains information about all people in the database. Each person has a unique integer id (*pid*), which is the key.
- *Event* is a list of all events. Each event has a unique integer event id (*eid*), which is the key. The start and end time of each event is given as an integer, which represents the number of seconds that have elapsed since 00:00:00 on January 1, 1970 (i.e., times are given as the number of seconds from a specific time in the past.) You may assume that the *start_time* of each event is always less than the event *end_time*.
- Each tuple in *Invited* indicates that the Person identified by *pid* has been invited to the Event identified by *eid*. The attribute *going* has the value 1 if the person will attend the event and is 0 if not. Attributes *pid* and *eid* are foreign keys referencing *Person* and *Event* respectively.

Write SQL queries that return the information described below.

(a) (7 points) Write a SQL query that returns the number of people who are invited to and are going to attend the event "Film Festival".

Person(<u>pid</u> , name) Event(<u>eid</u> , name, start_time, end_time) Invited(<u>pid</u> , <u>eid</u> , going)
--

Question 1. (cont) (b) (7 points) Write a SQL query that finds the event or events that the most people have been invited to attend. The query should return the event id(s) (eid) and event name(s) (name) of the event(s) that have the most persons invited, regardless of whether those people plan to attend.

(c) (7 points) Write a SQL query that returns the ids (pid) and names of all persons who have been invited to two distinct events that overlap in time. Two events overlap if one starts while the other is in progress. If the previous event ends at exactly the same time that a later one starts, the two events do not overlap. The query result should include each person (pid and name) only once, even if they have been invited to more than one overlapping event. Persons should be included in the result regardless of whether or not they plan to attend either or both events.

Question 2. (9 points, 3 each) Relational algebra. Suppose we have two relations A and B. Relation A has na tuples and relation B has nb tuples. For each of the following relational algebra expressions, give the minimum and maximum number of tuples that can appear in the result.

(You do not need to show your work, but it might be helpful in case we need to assign partial credit to a not-quite-right answer.)

(a) $A \cap B$

(b) $\sigma_c(A) - B$ (where c is some logical (Boolean) condition)

(c) $A \bowtie B$ (recall that \bowtie is the left-outer-join operator)

Question 3. (12 points, 4 each) XML. Suppose we have the following XML document describing a music catalog. The information is organized by label and lists for each label the artists and albums that have been released on that label. A single artist may have released albums on more than one label, so an artist name may appear under multiple label listings (A DTD for this data is given in the next problem. Feel free to reference that if it is useful. You also can remove this page for reference if you wish.)

```

<Catalog>
  <Label> <Name> Swan Song </Name>
    <Artist>
      <Name> Led Zeppelin </Name>
      <Album>
        <Title> Physical Graffiti </Title>
        <Song> Custard Pie </Song>
        <Song> Kashmir </Song>
        <Price> 19.99 </Price>
        <Year> 1975 </Year>
      </Album>
    </Artist>
  </Label >
  <Label> <Name> Albert </Name>
    <Artist>
      <Name> AC/DC </Name>
      <Album>
        <Title> T.N.T. </Title>
        <Song> Live Wire </Song>
        <Song> T.N.T. </Song>
        <Song> High Voltage </Song>
        <Price> 19.99 </Price>
        <Year> 1975 </Year>
      </Album>
      <Album>
        <Title> Back in Black </Title>
        <Song> Shoot to Thrill </Song>
        <Song> Back in Black </Song>
        <Song> Rock and Roll Ain't Noise Pollution </Song>
        <Price> 25.99 </Price>
        <Year> 1980 </Year>
      </Album>
    </Artist>
  </Label>
  <Label> <Name> RCA </Name>
    <Artist>
      <Name> Justin Timberlake </Name>
      <Album>
        <Title> The 20/20 Experience </Title>
        <Song> Suit & Tie </Song>
        <Song> Mirrors </Song>
        <Price> 12.99 </Price>
        <Year> 2013 </Year>
      </Album>
    </Artist>
  </Label>
</Catalog>

```

(continued on next page)

Question 3. (cont.) Give the results of the following XPath expressions when they are applied to the XML document on the previous page.

(a) `//Price`

(b) `Catalog/Label//Album[//Year/text() >= 1980]/Title`

(c) `Catalog/Label[count(//Album/Song) > 2]/Name`

Question 4. (12 points, 6 each) For reference, the DTD of the previous question’s XML data is:

```
<!DOCTYPE Catalog [  
<!ELEMENT Catalog (Label+)>  
<!ELEMENT Label (Name, Artist+)>  
<!ELEMENT Artist (Name, Album+)>  
<!ELEMENT Album (Title, Song+, Price, Year)>  
<!ELEMENT Name(#PCDATA)>  
<!ELEMENT Title(#PCDATA)>  
<!ELEMENT Song(#PCDATA)>  
<!ELEMENT Price(#PCDATA)>  
<!ELEMENT Year(#PCDATA)>  
>
```

Below, give XQuery expressions that return the specified results. The queries should work on any valid XML document described by the above DTD, not just the sample data given in the previous problem. The results should be well-formed XML. If needed, you can assume the data is in a file named “catalog.xml”.

(a) Return a list giving the names of every artist included in the catalog and, for each artist, the titles of all of the albums they have released on any label.

(b) Return a list showing for each year the names of artists who released an album during that year.

Question 5. (20 points) Database design. We would like to design a database to describe the members of the Marvel/DC comic book and movie universe. This universe has the following properties that we want to model.

- A citizen has a name and unique social security number (ssn)
- A superhero has a name and a unique superhero id number (sidn)
- A superhero power has a type and an effect, and that combination is unique
- A city has a name, population, latitude and longitude
- A superhero may have a secret identity as a citizen
- A superhero may have one or more powers
- Citizens and superheroes reside within a city (i.e. Gotham, Metropolis)

(a) (10 points) Draw an E/R diagram that captures this information.

(continued next page)

Question 5. (cont.) (b) (10 points) Give a series of SQL Create Table statements to create tables to hold the information in the E/R diagram from part (a) of this question. Your SQL statements should identify the key or keys of each table and include any appropriate constraints, including foreign key constraints.

Question 6. (15 points) Database design. We have a database with information about university courses. Unfortunately, it was designed by Mr. Frumble, and it has only one table with this schema:

Course(CourseID, Course_name, Instructor, Instructor_email, Max_students, Room)

These functional dependencies exist between attributes in this table:

CourseID -> Course_name
CourseID -> Instructor
CourseID -> Room
Instructor -> Instructor_email
Room -> Max_students

Is this schema in BCNF? If so, give a justification for your answer and identify the key(s) in the schema. If not, identify the bad functional dependencies and decompose the schema into two or more tables that are in BCNF. Identify the key(s) of each table.

(more space for your answer on the next page, if you need it)

Question 6. (cont.) Additional space for your answer if needed.

Question 7. (18 points, 6 each) Serializability. For each of the following schedules, draw the precedence (conflict) graph and decide if the schedule is conflict-serializable. If the schedule is conflict-serializable, give an equivalent serial schedule by listing the transactions in an order they could occur (you do not need to write down all of the read-write operations, just give the order of the transactions). If the schedule is not conflict-serializable, explain why not.

(a) R1(A) W4(C) R3(C) W2(B) R4(B) R3(B)

(b) R3(B) R1(A) R1(B) W1(C) R2(A) R1(D) R2(C) W1(B)

(c) R1(A) R2(A) W2(A) R3(B) R2(B) W3(B) W2(B) W1(A)

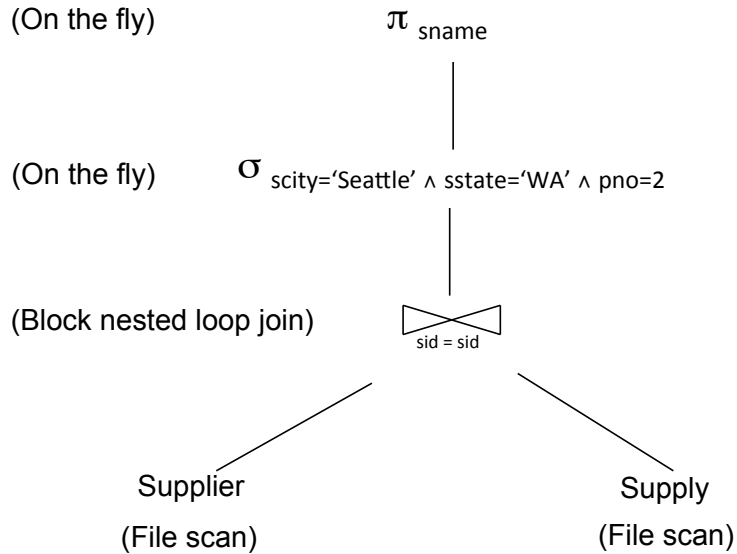
Question 8. (16 points) Query estimation. One of our favorite queries involves tables of suppliers and parts. The two tables have the following schemas:

Supplier(sid, sname, scity, sstate)
 Supply(sid, pno, quantity)

We assume that these tables have the following characteristics:

T(Supplier) = 1000 T(Supply) = 10000
 B(Supplier) = 100 B(Supply) = 100
 V(Supplier, scity) = 20 V(Supply, pno) = 2500
 V(Supplier, state) = 10

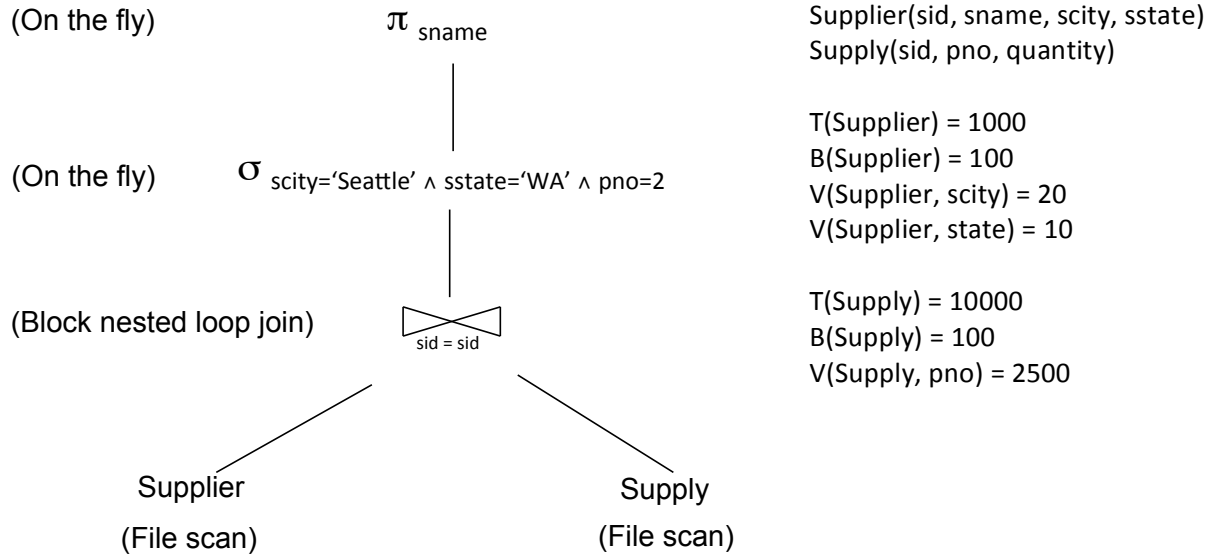
Now, consider the following query, where the physical operators to be used are shown in parentheses.



(a) (4 points) Translate this query into SQL.

Question 8 (cont.) (b) (6 points) What is the expected cost of this physical query? Remember that the cost of a query is measured in terms of the expected number of I/O operations, not including writing the result. Also, you should assume that there is enough main memory to hold all necessary data, if that is a consideration.

Query diagram and relation information repeated for convenience:



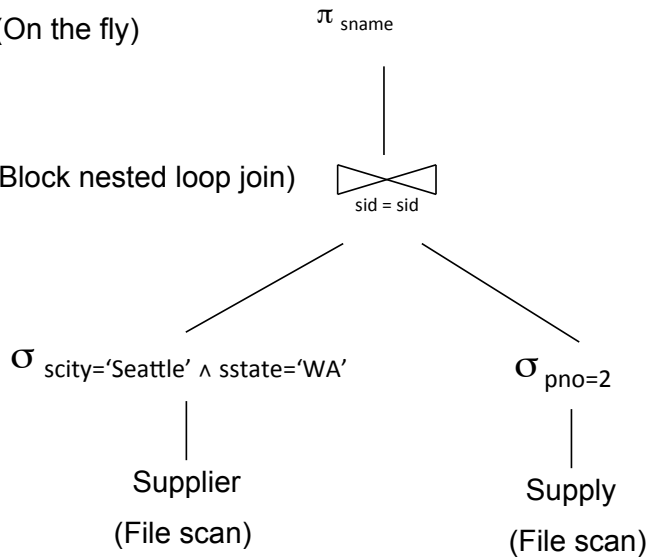
Write your cost estimate below. It would help to show enough of your work so we can follow it, in case it is necessary to award partial credit.

(continued next page)

Question 8. (cont) (c) (6 points) Now suppose we use this different plan for the same SQL query.

(On the fly)

(Block nested loop join)



Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

What is the cost estimate for this different plan? (Again, it would help to show enough of your work so we can follow it, but be sure the final answer is clearly indicated.)

Question 9. (12 points) Map-Reduce. One of the original applications of the map-reduce framework was to compute the *page rank* of each page on the web. The basic idea behind page rank is that pages are more likely to contain useful information if many other pages on the web contain links to them.

For this problem we'd like to compute a simplified version of page rank by counting the number of other pages that link to each page in our data. The input is a large collection of key-value data pairs (pageURL, linkURL) where pageURL is the key of the input data file and linkURL is the associated value. A (pageURL, linkURL) pair indicates that the page identified by pageURL contains a link to the page identified by linkURL. We assume that the data contains no duplicates, i.e., there is only one key-value pair for links from one page to another even if the page identified by pageURL contains multiple links to linkURL.

Describe a sequence of one or more map-reduce jobs (not Pig programs) that will report the number of pages that link to each page that appears as a linkURL in the input data. The output should be a collection of key-value pairs (linkURL, count), where count is the number of pages that contain a link to linkURL.

You need to clearly describe the (key, value) pairs that are input to and output from each map and reduce phase of the map-reduce job(s) needed. Be sure to clearly describe the input and output (key, value) pairs from each map and reduce stage, and explain (concisely) how the output of each map or reduce stage is computed from its input. If it takes more than one map-reduce job to compute the final result you should show how the output of each job is used as input to subsequent ones.

(additional space on the next page if needed – but don't be alarmed if you don't need it)

Question 9. (cont.) Additional space for your answer if needed.