CSE 444 Final Examination

June 10th, 2010

Name: ______________________________

Question	Points	Score
1	15	
2	20	
3	10	
4	10	
5	10	
6	10	
7	15	
8	10	
Total:	100	

This exam is an open book exam. You have 1 hour 50 minutes; budget time carefully. Intermediate steps are rarely required but often useful for partial credit. Good luck!

1 SQL

1. (15 points)

   Consider the following Flickr-type database:

   ```
   Users(uid, name)
   Picture(pid, owner, size)
   Comment(cid, auth, pict, text)
   ```

   Where:

   - `User.uid`, `Picture.pid`, `Comment.cid` are keys.
   - `Picture.owner`, `Comment.auth` are foreign keys into `Users`.
   - `Comment.pict` is a foreign key to `Picture`.

   (a) (5 points) Write a SQL query that counts, for every user, the number of other users who have commented on their pictures. That is, for each user you need to compute the total number of distinct users (including herself) who have commented on any of her pictures. Your answer should return only those answers where the count is ≥ 1 (i.e. you don't need to return users that have no comments on any of their pictures).

```
Users(uid, name)
Picture(pid, owner, size)
Comment(cid, auth, pict, text)
```

(b) (5 points) A *spammer* is a user who comments on all pictures that are not owned by him (her). Write a SQL query that returns all spammers.

```
Users(uid, name)
Picture(pid, owner, size)
Comment(cid, auth, pict, text)
```

(c) (5 points) Two users are *friends* if they comment on each others' pictures: that is, x, y are friends if x made a commented on some picture of y, and y made a comment on some picture of x. Write a SQL query that computes, for each user, the number of his/her friends. Your query should return only those answers where the count is ≥ 1.

2 Transactions

2. (20 points)

   Consider a concurrency control manager by timestamps. Below are several sequences of events, including start events, where sti means that transaction Ti starts and coi means Ti commits. These sequences represent real time, and the timestamp-based scheduler will allocate timestamps to transactions in the order of their starts. In each case below, say what happens with the last request.

   You have to choose between one of the following four possible answers:

   1. the request is accepted,
   2. the request is ignored,
   3. the transaction is delayed,
   4. the transaction is rolled back.

   (a) (4 points) st1; st2; st3; r1(A); w1(A); r2(A);
   The system will perform the following action for r2(A): ______________________________

   (b) (4 points) st1; st2; r2(A); co2; r1(A); w1(A)
   The system will perform the following action for w1(A): ______________________________

   (c) (4 points) st1; st2; st3; r1(A); w2(A); w3(A); r2(A);
   The system will perform the following action for r2(A): ______________________________

   (d) (4 points) st1; st2; r1(A); r2(A); w1(B); w2(B);
   The system will perform the following action for w2(B): ______________________________

   (e) (4 points) st1; st2; st3; r1(A); w3(A); co3; r2(B); w2(A)
   The system will perform the following action for w2(A): ______________________________

3 Conceptual Design

3. (10 points)

   (a) (5 points) Decompose in BCNF the relation $R(A, B, C, D, E)$ that satisfies the following functional dependencies. Show your steps, and show the keys in the decomposed relations.

$$\begin{aligned} A &\rightarrow B \\ CD &\rightarrow E \end{aligned}$$

(b) (5 points) For each of the statements below, indicate whether they are true or false. You do not need to justify your answers.

- Every relation with only two attributes is in BCNF.
- If X and Y are super-keys then $X \cup Y$ is also a super-key.
- If X and Y are super-keys then $X \cap Y$ is also a super-key.
- If $X \to A$ and $Y \to A$, then $(X \cap Y) \to A$.
- If $X^+ = X$ and $Y^+ = Y$ then $(X \cap Y)^+ = (X \cap Y)$.

4 Indexes

4. (10 points)

Consider the following database about word occurrences in Webpages:

```
Webpage(url, author)
Occurs(url, wid)
Word(wid, text, language)
```

where:

- `Webpage.url` and `Word.wid` are keys.
- `Occurs.url` and `Occurs.wid` are foreign keys to `Webpage` and `Word` respectively.

Assume the following statistics

$$
\begin{aligned}
T(\texttt{Webpage}) &= V(\texttt{Occurs}, \texttt{url}) = 10^9 \\
T(\texttt{Occurs}) &= 10^{12} \\
T(\texttt{Word}) &= V(\texttt{Occurs}, \texttt{wid}) = 10^6 \\
V(\texttt{Webpage}, \texttt{author}) &= 10^7 \\
V(\texttt{Word}, \texttt{language}) &= 100
\end{aligned}
$$

Assume ten records can be fit in one block, hence $B(\texttt{Webage}) = T(\texttt{Webpage})/10$ and similarly for all other tables.

(a) (5 points) Consider the following plan:

$$(\sigma^{\texttt{index-lookup}}_{\texttt{author='John'}}(\texttt{Webpage}) \bowtie^{\texttt{index-join}}_{\texttt{url=url}} \texttt{Occurs})$$

$$\bowtie^{\texttt{main-memory-hash-join}}_{wid=wid} \sigma^{\texttt{index-lookup}}_{\texttt{language='French'}}(\texttt{Word})$$

Compute the cost of the plan in each of the following cases:

1. We have the following indexes:

`Webpage.url`	=	primary index
`Webpage.author`	=	secondary index
`Occurs.url`	=	secondary index
`Occurs.wid`	=	primary index
`Word.wid`	=	primary index
`Word.language`	=	secondary index

2. We have the following indexes:

`Webpage.url`	=	secondary index
`Webpage.author`	=	primary index
`Occurs.url`	=	primary index
`Occurs.wid`	=	secondary index
`Word.wid`	=	secondary index
`Word.language`	=	primary index

(b) (5 points) Consider the following plan:

$$(\texttt{Webpage} \bowtie^{\texttt{merge-join}}_{\texttt{url=url}} \texttt{Occurs}) \bowtie^{\texttt{merge-join}}_{wid=wid} \texttt{Word}$$

Choose a set of indexes that minimizes the total number of disk I/O's for the plan. Each index should be on a single attribute, and you can choose as many (or as few) indexes as you want. Indicate for each index if it is primary or secondary. For each merge-join we assume to have sufficient main memory to complete it in two pass. (There are no index-join operators in this plan: you need to figure out how indexes can help at all in this query plan.)

5 Query Optimization

5. (10 points)

   For the following questions, we consider the schema $R(A, B)$, $S(C, D)$, $T(E, F)$, $U(G, H)$.

   (a) (5 points) Write a logical plan for the following query:

```
select R.A, sum(T.F)
from R, S, T
where R.B = S.C and S.D = T.E
group by R.A
having count(*) > 20
```

(b) (5 points) Consider the following query:

```
select *
from R, S, T, U
where R.B = S.C and S.D = T.E and S.D = U.G
```

- Write **all** left-deep, cartesian-free join trees for this query. Assume that the join operator is not commutative: that is, you should report both $R \bowtie S$ and $S \bowtie R$ as distinct plans. Note: you need to turn in several plans.

- Write a full semijoin reducer for this query.

6 Statistics

6. (10 points)

Consider the relations $R(A, B), S(C, D), T(E, F)$ and the following histograms on $R.A$ and $T.F$:

$R.A$	$0 \dots 999$	$1000 \dots 1999$	$2000 \dots 2999$	$3000 \dots 3999$	$4000 \dots 4999$
	10^4	$2 \cdot 10^4$	$3 \cdot 10^4$	$3 \cdot 10^4$	$2 \cdot 10^4$

$T.F$	$0 \dots 2499$	$2500 \dots 2699$	$2700 \dots 3999$	$4000 \dots 7999$
	10^4	10^4	10^4	10^4

(a) (1 point) What kind of histogram is $R.A$ and what kind of histogram is $T.F$?

(b) (4 points) Estimate the number of tuples returned by $\sigma_{500 \le A \le 3499}(R)$. Show your work for partial credit.

(c) (5 points) Estimate the number of tuples returned by the following query:

```
SELECT *
FROM R, S, T
WHERE R.A = 2432 and R.B = S.C and S.D = T.E and T.F = 1234
```

assuming the two histograms above, plus the following statistics:

$$
\begin{aligned}
T(R) &= 10^5 \quad T(S) = 6 \cdot 10^6 \quad T(T) = 4 \cdot 10^5 \\
V(R, B) &= V(S, C) = 3 \cdot 10^3 \\
V(S, D) &= V(T, E) = 2 \cdot 10^4
\end{aligned}
$$

Show your work for partial credit.

7 Parallel Databases

7. (15 points)

   Consider the following relations $R(A, B), S(C, D)$ and the following query:

```
select R.A, sum(S.D)
from R, S
where R.B = S.C
group by R.A
```

   (a) (5 points) Write the query in Pig Latin.

   (b) (5 points) How many map-reduce tasks does Pig Latin require to evaluate the query ?

   (c) (5 points) Give the total I/O cost of your Pig Latin program, as a function of $B(R), B(S), B(R \bowtie S)$, and P (the number of processors). Assume that each processor's main memory is $\geq \max(B(R), B(S), B(R \bowtie S))/P$. Note: your analysis should be based on the actual implementation of map-reduce, and should not assume any other optimizations.

8 Bloom Filters

8. (10 points)

   (a) (1 point) Is the following statement true or false ? Bloom filters are used to improve cardinality estimation.

   (a) ____________

   True or false ?

   (b) (2 points) A Bloom filter using a hash map of 1k Bytes has a false positive error rate of 19%. In order to improve the error rate, the systems administrator decides to double the size of the hash map to 2k Bytes, and, at the same time, to double the number of hash functions used by the Bloom filter. What is the new false positives error rate ?

   (b) ____________

   The new false positive error rate is:

   (c) (2 points) A regular hash map of 1k Bytes has a false positive error rate of 19%. In order to improve this rate, the systems administrator decides to double the size of the hash map to 2k Bytes. What is the new false positive error rate ?

   (c) ____________

   The new false positive error rate is:

(d) (5 points) Data supplier S1 has $n = 1M$ $(= 10^6)$ documents. Data supplier S2 has also $n = 1M$ documents. Each document has $1k$ bytes. They have 50 documents in common and they want to compute these. They will proceed as follows:

- S1 computes a hash map M with cn bits, where c=8 and sends it to S2
- S2 checks its items in M and sends all matches to S1
- S1 computes the result and sends the matching 50 documents to S2

Indicate the total number of bytes transferred over the network in each step assuming (a) the hash map is a standard hash table, (b) the hash map is a bloom filter.