

CSE 344 Final Examination

December 12, 2012, 8:30am - 10:20am

Name: _____


Question	Points	Score
1	30	
2	20	
3	30	
4	20	
Total:	100	

- This exam is open book and open notes but NO laptops or other portable devices.
- You have 1h:50 minutes; budget time carefully.
- Please read all questions carefully before answering them.
- Some questions are easier, others harder; if a question sounds hard, skip it and return later.
- Good luck!


1 E/R Diagrams, Constraints, Conceptual Design

1. (30 points)


- (a) (10 points) For each E/R diagram, indicate if the CREATE TABLE statement that follows the diagram is consistent with the diagram or not. If it is not consistent, make the necessary changes. Only make changes that are necessary.




```
CREATE TABLE A (a1 int PRIMARY KEY,
                a2 int,
                b int,
                c1 int NOT NULL REFERENCES C)
```



```
CREATE TABLE C (p1 int PRIMARY KEY,
                c1 int)
```



```
CREATE TABLE G (h1 int REFERENCES H,
                k1 int REFERENCES K,
                l1 int REFERENCES L,
                g1 int,
                PRIMARY KEY(h1,k1,l1 )
```



```
CREATE TABLE O (o1 int,
                m1 int REFERENCES M,
                PRIMARY KEY(m1,o1))
```

Answer (Corrections, if any, to the above CREATE TABLE statements):
Write below or edit directly above.

Solution:

```
CREATE TABLE A (a1 int PRIMARY KEY,
                a2 int,
                a3 int,
                b int,
                c1 REFERENCES C)
CREATE TABLE C (p1 int PRIMARY KEY REFERENCES P,
                c1 int)
CREATE TABLE G is correct
CREATE TABLE O (o1 int PRIMARY KEY)
```

- (b) (10 points) You are working with a customer who just designed his E/R diagram and started to write the following CREATE TABLE statements.

```
CREATE TABLE Product (pid INT PRIMARY KEY,  
                        name VARCHAR(20))
```

```
CREATE TABLE Inventory(pid INT PRIMARY KEY,  
                        quantity INT,  
                        FOREIGN KEY (pid) REFERENCES Product)
```

```
CREATE TABLE Supplier(sid INT PRIMARY KEY,  
                       name VARCHAR(20),  
                       address VARCHAR(50))
```

```
CREATE TABLE Supplies(sid INT REFERENCES Supplier,  
                      pid INT REFERENCES Product,  
                      PRIMARY KEY (sid, pid))
```

```
CREATE TABLE PurchaseOrder (pid INT REFERENCES Product,  
                              quantity INT)
```

The customer would like to add certain constraints to his database. Advise your customer on the appropriate implementation for each of the following constraints. If appropriate, show how to modify the CREATE TABLE statements:

- Ensure that attribute `quantity` from table `Inventory` is always greater than or equal to 0.

Answer (Your suggestion):

Solution:

```
CREATE TABLE Inventory(pid INT PRIMARY KEY,  
                        quantity INT CHECK (quantity >= 0),  
                        FOREIGN KEY (pid) REFERENCES Product)
```

- Whenever someone deletes a tuple in **Supplier**, any tuple in **Supplies** that referred to it should also be deleted.

Answer (Your suggestion):

Solution:

```
CREATE TABLE Supplies(sid INT REFERENCES Supplier
                        ON DELETE CASCADE,
                        pid INT REFERENCES Product,
                        PRIMARY KEY (sid, pid))
```

- If the quantity of a product in inventory goes down to zero, automatically insert a tuple associated with this product in the **PurchaseOrder** relation (explain your suggestion in English):

Answer (Your suggestion):

Solution: Create a row-level trigger.

- **Event** is an update to the quantity attribute of the **Inventory** relation.
- **Condition** is that the quantity is now zero
- **Action** is to insert a tuple into **PurchaseOrder**

- (c) (10 points) Consider the following relational schema and set of functional dependencies.

$R(A,B,C,D,E,F,G)$ with functional dependencies:

$A \rightarrow D$

$D \rightarrow C$

$F \rightarrow EG$

$DC \rightarrow BF$

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).

Answer (Decompose R into BCNF):

Solution: Watch-out! The first FD does NOT violate BCNF so we need to pick another one to decompose. We try the second one:

Try $\{D\}^+ = \{B, C, D, E, F, G\}$. Decompose into $R_1(B, C, \underline{D}, E, F, G)$ and $R_2(\underline{A}, D)$.

R_2 has two attributes, so it is necessarily in BCNF.

For R_1 , again not all FDs violate BCNF so we need to be careful.

Try $\{F\}^+ = \{E, F, G\}$. Decompose into $R_{11}(E, \underline{F}, G)$ and $R_{12}(B, C, \underline{D}, F)$.

Both R_{11} and R_{12} are in BCNF.

2 Transactions

2. (20 points)

- (a) (10 points) Consider the following transaction schedules. For each schedule, indicate if it is **conflict-serializable** or not:

r1(A); r2(B); r1(B); w2(B); w1(A); w1(B); r2(A); w2(A); c1; c2

Answer (YES/NO):

Solution: NO

r1(A); r2(B); r3(A); r2(A); r3(C); r1(B); r3(B); r1(C); r2(C); c1; c2; c3

Answer (YES/NO):

Solution: YES

w1(A); w2(A); w1(B); w3(B); w1(C); w3(C); w2(C); c1; c2; c3

Answer (YES/NO):

Solution: YES

- (b) (10 points) Your friend Bob just wrote a Java application that talks to a back-end SQL Server DBMS. The application enables students to register for courses. Looking through the code, you notice that Bob's application performs the following sequence of operations.

Prompt the user for a student ID and password.

Start a new transaction.

Look up the student in the database.

If the student ID is not in the database or the password is incorrect, abort the transaction.

Look up the courses recommended for the student.
Display the courses on the screen.

While the user does not choose QUIT

 Prompt the user to select a course.

 If the course is available, then register the student.

End while

Commit the transaction.

Explain the problem in Bob's design. Explain why this is a problem. Explain how to fix the problem. You do not need to write the updated application pseudo-code.

Answer (Be careful to answer all three parts of the question):

Solution:

- What is the problem? Bob put too many operations, including waiting for user input, inside a single, very long transaction.
- Why is this a problem? This is a problem because it will reduce the throughput of the application considerably. Each transaction will hold locks for an unnecessary long period of time, preventing concurrent accesses to the data by other users. Additionally, if a failure occurs, the user will lose a lot of work as all his or her course registrations will get rolled back.
- How to fix the problem? The solution is to use finer-grained transactions: Check the login and password as one transaction. Look-up the recommended courses as a second transaction. Then each attempt to register for one course should be executed as its own transaction.

3 Parallel Data Processing


3. (30 points)

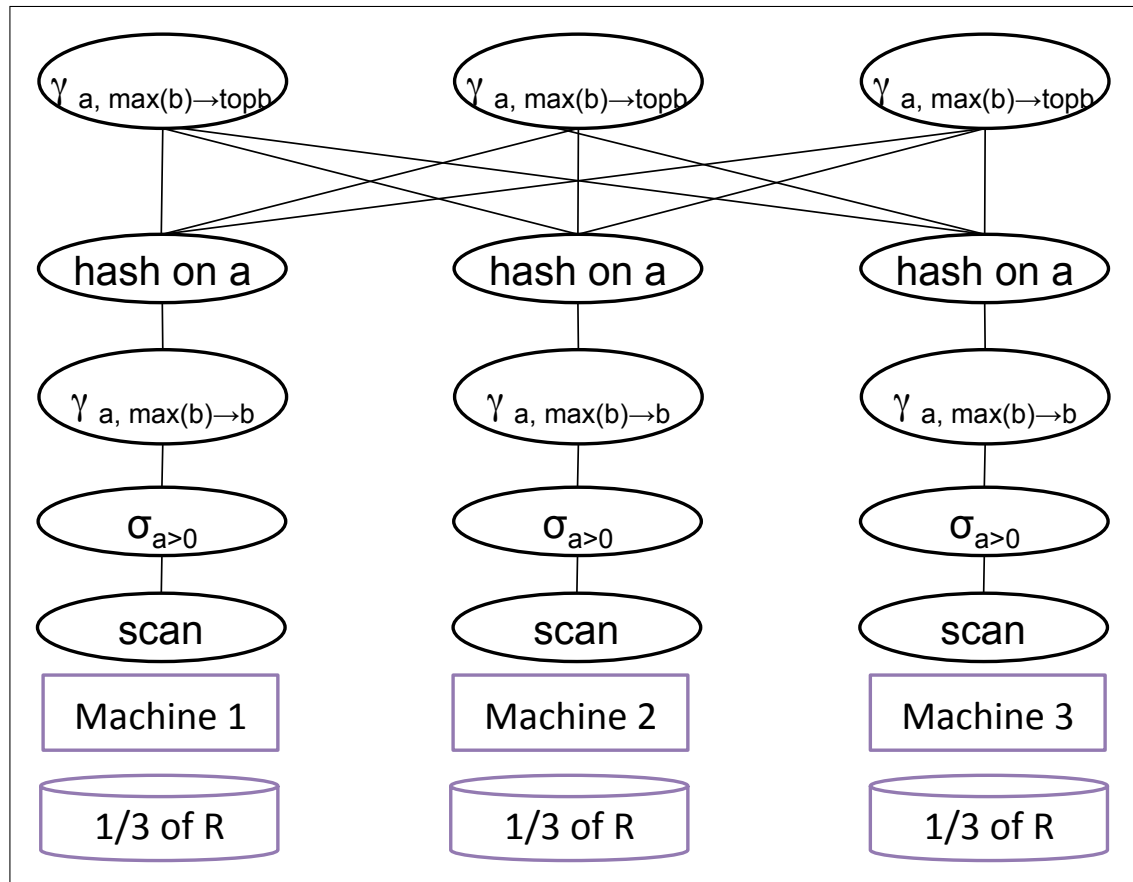
- (a) (15 points) Consider a relation $R(a, b)$ that is horizontally partitioned across $N = 3$ machines as shown in the diagram below. Each machine locally stores approximately $\frac{1}{N}$ of the tuples in R . The tuples are randomly organized across machines (i.e., R is block partitioned across machines).

Show a relational algebra plan for this query and how it will be executed across the $N = 3$ machines. Pick an *efficient* plan that leverages the parallelism as much as possible. Include operators that need to re-shuffle data and add a note explaining how these operators will re-shuffle that data.

```
SELECT a, max(b) as topb
FROM R
WHERE a > 0
GROUP BY a
```

Answer (Draw the parallel query plan):





- (b) (10 points) Explain how the query would be executed in MapReduce (**not Pig**). Make sure to specify the computation performed in the map and the reduce functions.

Answer (Describe the map and reduce functions):

Solution: We describe the functions and the overall execution.

Each map task scans a block of R and calls the map function for each tuple. The map function applies the selection predicate to the tuple. For each tuples that passes the selection, it outputs a record with key=a and value=b.

```
map(String key, String value):  
// key: irrelevant  
// value: tuple  
if a > 0  
    EmitIntermediate(a, b);
```

The MapReduce engine reshuffles the output of the map phase and groups it on the intermediate key, which is attribute a.

Each reduce task computes the aggregate value for each group assigned to it (by calling the reduce function) and outputs the final result.

```
reduce(String key, Iterator values):  
// key: attribute a  
// values: all values of b in the group  
  
topb = compute max value in the group  
Emit(topb);
```

It is possible to use the *combiner* to perform a local aggregation before the data gets re-shuffled.

- (c) (5 points) What would change if we hash-partitioned R on $R.a$ *before* executing the above query? Please discuss the case of the parallel DBMS and the case of MapReduce.

Answer (Explain the difference for a parallel DBMS and for MapReduce):

Solution: The parallel DBMS would avoid the data re-shuffling step. It would compute the aggregates locally.

Hash-partitioning the data would not make any difference for the MapReduce job as MapReduce wouldn't know that the data is hash-partitioned. Additionally, MapReduce treats map and reduce functions as black-boxes and does not perform any optimizations on them.

4 NoSQL and Data Integration

4. (20 points)

- (a) (10 points) In order to scale a DBMS, one approach is to horizontally partition each relation across a set of machines in a cluster. Why is it expensive to execute transactions that touch arbitrary data in such a configuration?

Answer:

Solution: It is difficult to ensure ACID properties for transactions that touch data on multiple machines because the DBMS instances running on those machines must agree to either COMMIT or ABORT the whole transaction. Such an agreement requires that the machines involved in the transaction run the two-phase commit (2PC) protocol, which requires the exchange of multiple messages and is thus expensive.

Additionally, if a transaction is allowed to touch arbitrary data, it may include expensive SQL queries that require a large amount of data to be re-shuffled between machines, which can also be very expensive.

- (b) (10 points) Jack is looking at customer data from two different branches of his store. Luckily, Jack made sure that the database in each branch followed the same schema. Each branch has a table CustomersBranch holding the ID, name, and state of each customer. He thus integrates the data by defining the following view

```
CREATE VIEW CustomersData AS
SELECT B1.name, B1.state
FROM CustomersBranch1 B1
UNION
SELECT B2.name, B2.state
FROM CustomersBranch2 B2
```

He now wants to analyze this data and he executes the following SQL query:

```
SELECT state, count(*)
FROM CustomersData
GROUP BY state
```

But the results are different from what he expected: Both the number of groups and the counts of customers are much higher than what he expected. Give a possible explanation for each of these two problems..

Answer (Give two reasons):

Solution:

- Problem with too many groups: State names might be misspelled or perhaps one store uses abbreviations while the other one doesn't (value heterogeneity).
- Problem with too many customers: Customer names may appear with slight variations in both databases even though both refer to the same physical person (duplicate records).