

UCLA
Computer Science Department
Winter 2014

Instructor: J. Cho

Student Name and ID: _____

CS143 Midterm: Closed Book, 90 minutes

(IMPORTANT PLEASE READ **):**

- There are 3 problems on 10 pages for a total of 70 points to be completed in 90 minutes. *You should look through the entire exam before getting started, in order to plan your strategy.*
- *Simplicity and clarity of your solutions will count.* You may get as few as 0 point for a problem if your solution is far more complicated than necessary, or if we cannot understand your solution.
- If you need to make any assumption to solve a question, please write down your assumptions.
- To get partial credits, you may want to write down how you arrived at your answer step by step.
- You may use one-page double-sided cheat-sheet during exam. You are also allowed to use a calculator.
- Please, write your answers neatly. Attach extra pages as needed. Write your name and ID on the extra pages.

Problem	Score	
1	25	
2	25	
3	20	
Total	70	

Problem 1 (Queries): 25 points

1. The table `Sensor(date, time, temp, humidity)` stores the temperature and the humidity values every few hours in the Los Angeles airport. Write a relational algebra query that returns the highest temperature value of each day in the table. That is, your query should return `(date, highest temperature of the day)` pairs for each day in the table. Make your query simple and clear. Your answer will be graded based on simplicity as well as correctness. (10 points)

ANSWER:

$$\pi_{date,temp}(Sensor)$$

$$- \pi_{R1.date,R1.temp}(\sigma_{R1.date=R2.date \wedge R1.temp < R2.temp}(\rho_{R1}(Sensor) \times \rho_{R2}(Sensor)))$$

2. Consider the following database tables on car ownership and car accidents:

`Person(SSN, name, address)`

`Car(license, year, model)`

`Accident(license, accident-date, driver, damage-amount)`

`Owns(SSN, license)`

Note that the driver involved in a car accident may not always be the owner of the car. We assume that a car cannot get involved in more than one accident at a certain date.

For the following English statements, choose all queries in SQL or relational algebra that return the proper answer. Note that each question may have zero, one or more correct answers. List only the ones that are correct. Write 'None' if there is no correct answer.

- (a) Find the SSN of all persons who have had all of their cars involved in an accident. (5 points)

i. `(SELECT SSN FROM Owns, Accident
WHERE Owns.license = Accident.license)
EXCEPT
(SELECT SSN FROM Owns
WHERE license NOT IN (SELECT license FROM Accident))`

ii. $\pi_{SSN,license}(Owns \bowtie Accident) / \pi_{license}(Accident)$

Correct queries:

ANSWER:

(i) only. (ii) returns the owner(s) who own every car that has been in an accident.

- (b) Who is the driver who participated in the most costly accident? Return the driver and the amount of damage. (5 points)
- i. `SELECT driver, damage-amount FROM Accident
WHERE damage-amount IN (SELECT MAX(damage-amount) FROM Accident)`
 - ii. `SELECT driver, damage-amount FROM Accident
WHERE damage-amount = MAX(damage-amount)`
 - iii. `(SELECT driver, damage-amount FROM Accident)
EXCEPT
(SELECT A1.driver, A1.damage-amount
FROM Accident A1, Accident A2
WHERE A1.damage-amount < A2.damage-amount)`

Correct queries:

ANSWER:

(i), (iii) only. (ii) generates error because of MAX() in WHERE clause.

- (c) Find the license number of all cars that have been involved in more than one accident (DO NOT RETURN DUPLICATES) (5 points)
- i. `SELECT license FROM Accident
GROUP BY license
HAVING COUNT(incident-date) > 1`
 - ii. `SELECT A1.license FROM Accident A1, Accident A2
WHERE A1.license = A2.license AND A1.incident-date <> A2.incident-date`
 - iii. `SELECT DISTINCT A1.license FROM Accident A1
WHERE A1.license IN (SELECT A2.license
FROM Accident A2
WHERE A1.incident-date <> A2.incident-date)`

Correct queries:

ANSWER:

(i), (iii) only. (ii) may return duplicate license numbers.

Problem 2 (Integrity, View, and Authorization): 25 points

For all questions in this problem, we refer to three tables created by the following statements:

```
CREATE TABLE Vendor(
  vid INT,
  name CHAR(20),
  PRIMARY KEY(vid));

CREATE TABLE Product(
  pid INT,
  name CHAR(20),
  vid INT,
  PRIMARY KEY(pid),
  FOREIGN KEY(vid) REFERENCES Vendor(vid)
    ON DELETE CASCADE
    ON UPDATE CASCADE);

CREATE TABLE Sales(
  pid INT,
  price INT,
  quantity INT,
  sales-time DATETIME
  PRIMARY KEY(pid, sales-time),
  FOREIGN KEY(pid) REFERENCES Product(pid));
```

As you can imagine from their names, Vendor table contains the vendor information, Product table contains the product information and its vendor, and the tuples in the Sales table correspond to the actual sales of a product made at a particular date and time.

Assume that the three tables contain the following tuples:

Vendor

vid	name
1	Motorola
2	LG
3	Samsung

Product

pid	name	vid
1	Razr	1
2	Vue	2
3	Blackjack	3
4	Slide	3

Sales

pid	price	quantity	sales-time
1	100	1	2008-10-25 10:00:00
3	150	1	2008-10-26 11:00:00

1. Write a *single* SQL constraint or SQL trigger statement that ensures that the sales-time of any newly added Sales tuple should not be before that of any existing Sales tuple with the same pid. That is, any new sale of a product cannot precede a prior sale of the product. Note that the aggregate function MAX() returns NULL if it is applied to an empty set of tuples. If your answer is a PRIMARY KEY, a FOREIGN KEY, or a CHECK constraint, just write the constraint definition part and the name of the table for which the constraint is defined. If your answer is a trigger, write its full definition. (5 points)

ANSWER:

```
CREATE TRIGGER NO-EARLIER-SALE AFTER INSERT ON Sales REFERENCING NEW ROW
AS nr FOR EACH ROW WHEN (nr.sales-time < (SELECT MAX(sales-time) FROM Sales
WHERE pid = nr.pid) DELETE FROM Sales WHERE pid = nr.pid AND sales-time
= nr.sales-time
```

2. For each of the following statement, briefly state what will happen when the statement is issued. If some tuples are deleted or updated due to the statement, clearly indicate the exact list of tuples that are affected (and their new values if they are updated). If the statement is rejected, briefly explain why. Do *NOT* assume the constraints/triggers that you defined for question 1. For each statement, assume the table instances given at the beginning of this problem. Consider each statement separately from others.

- (a) DELETE FROM Vendor WHERE vid = 3; (2.5 points)

ANSWER:

The statement is rejected because it will try to delete (3, 'Blackjack', 3) from Product table, but this will violate the FOREIGN KEY constraint in Sales table.

- (b) DELETE FROM Vendor WHERE name = 'LG'; (2.5 points)

ANSWER:

(2, 'LG') is deleted from the Vendor table. (2, 'Vue', 2) is also deleted from the Product table

- (c) UPDATE Vendor SET vid = 4 WHERE vid = 1; (2.5 points)

ANSWER:

(1, 'Motorola') is updated to (4, 'Motorola') in Vendor. (1, 'Razr', 1) is also updated to (1, 'Razr', 4) in Product.

- (d) UPDATE Product SET pid = 5, vid = 3 WHERE pid = 2; (2.5 points)

ANSWER:

(2, 'Vue', 2) is updated to (5, 'Vue', 3)

3. Assume that u_1 is the owner of all three tables, Vendor, Product, and Sales. Now u_1 wants to allow another user u_2 to create the ProductSales view by executing the following statement:

```
CREATE VIEW ProductSales AS
  SELECT pid, SUM(quantity*price) total_sales
  FROM Sales
  GROUP BY pid
```

Assume that u_1 wants to allow u_2 to give out the SELECT privilege for the ProductSales view to any other user that u_2 chooses. Write a single GRANT statement that provides the *minimum* set of privileges that have to be given to u_2 to allow this. That is, u_2 should be able to run the above statement and allow other users to SELECT on ProductSales view. (5 points)

ANSWER:

```
GRANT SELECT ON Sales TO 'u2' WITH GRANT OPTION;
```

4. For this question, assume that u_2 has already created ProductSales view after obtaining the necessary privileges from u_1 . In addition, another user, u_3 , has also created the following TotalSales view after u_2 granted the necessary privileges to u_3 :

```
CREATE VIEW TotalSales AS
  SELECT SUM(total_sales) revenue
  FROM ProductSales
```

Now u_1 wants to control the access to the Sales table more tightly and decides to run the following REVOKE command:

```
REVOKE GRANT OPTION FOR SELECT ON Sales FROM 'u2' CASCADE;
```

Briefly explain what will happen to all views, ProductSales and TotalSales, after the above command is executed. If the command generates an error, briefly explain why. (5 points)

ANSWER:

ProductSales stays. TotalSales gets dropped.

Problem 3 (Disk, File, and Index): 20 points

Consider the Sales table created with the following statement:

```
CREATE TABLE Sales(
  sales-id      INT,
  product-id    INT,
  product-name  CHAR(80),
  price         INT,
  quantity      INT,
  sales-time    TIME,
  PRIMARY KEY(sales-id))
```

We store one million tuples in the Sales table in the following disk:

- 1 platter (2 surfaces)
- 2,000 cylinders
- 100 sectors per track
- 2048 bytes per sector
- 6,000 RPM rotational speed
- 10ms average seek time

One disk sector corresponds to one disk block. INT and TIME datatypes require 4 bytes each and CHAR(n) datatype requires n bytes. Each Sales tuple is stored as a fixed-length record and the tuples are *NOT SPANNED*. We do not store any additional header in each block, so all block space is available to store tuples.

1. What is the (burst) transfer rate of our disk? (5 points)

ANSWER:

$(2048 \text{ bytes/sector}) * (100 \text{ sector/track}) * (6,000 \text{ track/min}) * (1 \text{ min}/60 \text{ sec}) = 20\text{M bytes/sec.}$

2. What is the minimum number of blocks that we need to store the Sales table? (5 points)

ANSWER:

$\lceil \frac{2048\text{bytes/block}}{100\text{bytes/tuple}} \rceil = 20\text{tuples/block}$. $\lceil \frac{1,000,000\text{tuple/table}}{20\text{tuple/block}} \rceil = 50,000\text{block/table}$. We need at least 50,000 blocks.

3. We need to retrieve the sales tuple with 'sales-id = 307' by scanning the table. The tuples are *NOT* sequenced by sales-id. Assume that the Sales table is stored in 100,000 disk blocks. The disk blocks for the table are allocated sequentially. Assume that the sustained transfer rate of our disk is 20MB/sec. Note that the sustained transfer rate is the average rate at which we can sequentially read a large volume of data that may span over multiple tracks. (These numbers may or may not be the same as your previous answers.) For other parameters, assume the numbers we gave in the beginning. How long does it take to retrieve the tuple? Compute the *expected (or average)* time. (5 points)

ANSWER:

It takes 10ms for the initial seek and 5ms for the rotational delay. Once we start reading the table block, we will have to read half of the table on average to find the tuple. Half of the table (50,000 block) is 100MB, so it takes $100/20 = 5$ sec to read the half. In total, it takes 5.015 sec (5 sec) on average.

4. Now assume that the table has been sequenced by the sales-id attribute and we retrieve the Sales tuple with 'sales-id = 307' through the binary search algorithm. For all parameters, assume the numbers we gave in the beginning. How long does it take to retrieve the tuple? Compute the *expected (or average)* time. For simplicity, you may assume that every retrieval of a tuple during binary search incurs one random disk IO. For your information, $\log_2(10) = 3.322$. (5 points)

ANSWER:

Binary search requires $\lceil \log_2(10^6) \rceil = \lceil 6 \times 3.322 \rceil = \lceil 19.932 \rceil = 20$ comparisons. Each tuple retrieval requires 10ms + 5ms + 0.1ms. Therefore, it takes $20 \times 15.1\text{ms} = 302\text{ms}$.