CSE 444 Final Examination

June 7, 2012, 8:30am - 10:20am

Name: _____

Question	Points	Score
1	20	
2	15	
3	20	
4	15	
5	30	
Total:	100	

- This exam is an open book exam.
- 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.
- Relax! You are here to learn.

1 Indexes

1. (20 points)

   (a) (10 points) Consider the following B+ tree. Draw the modified tree after deleting value 19 and then 50. We only ask you to draw the final tree but you can draw both for partial credit. For the parts of the tree that remain the same, you can simply write "same":

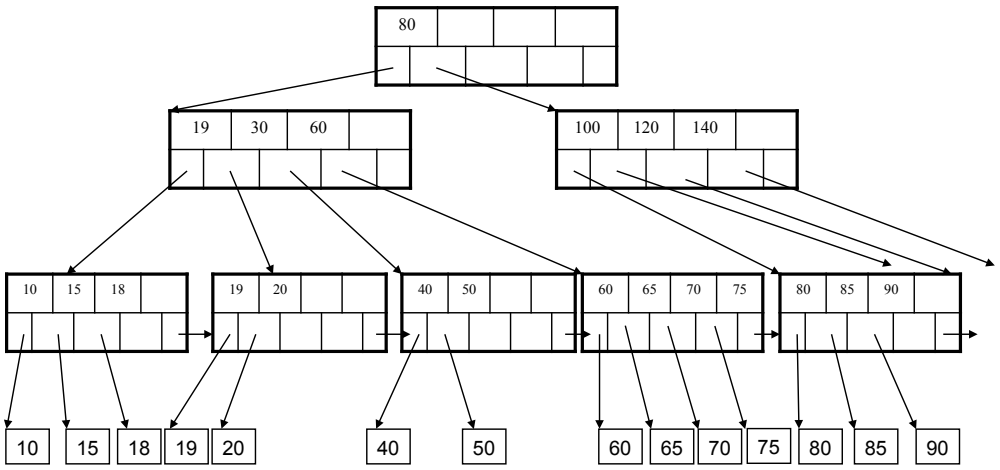

Answer (Draw an updated B+ tree):

(b) (10 points) Consider a relation R(a,b,c) with T(R) = 10,000 records, B(R) = 1,000 pages (10 records fit on each page), and where a is a non-negative integer primary key. How many pages will be read from disk to answer the selection query $\sigma_{a<2500}(R)$ in each of the following scenarios. Assume that 100 records match the selection predicate.

1. Relation R is stored in a heap file.
2. Relation R is stored in a sequential file sorted on a and there is a B+ tree index with search key a. All index pages are already in memory.
3. Relation R is stored in a heap file. There also exists an unclustered B+ tree index with search key a. All index pages are already in memory.
4. Relation R is stored in a heap file. There also exists an unclustered hash-based index with search key a. None of the index pages are in memory.

Answer (Compute the cost of the query in all the cases):

2 Query Optimization

2. (15 points)

Consider the following four relations:

Relation	Records	Pages	Unique Values
R(a,b,c)	T(R) = 10,000	B(R) = 1,000	
S(d,e)	T(S) = 100,000	B(S) = 10,000	V(S,d) = 100 and V(S,e) = 100
U(f,g)	T(U) = 100,000	B(U) = 10,000	V(U,f) = 10,000.
W(h,i,j)	T(W) = 1,000	B(W) = 100	

Additionally, consider that there exists an unclustered B+ tree index on R.a, which is the primary key of R.

(a) (15 points) What is the cost of the query plan below assuming 2,000 pages of memory? Count only the number of page I/Os.

Hints:

- As you compute the cost of the plan in a bottom up fashion, at the output of each operator consider whether the result can fit in memory and be used directly by the next operator or if you need to write anything to disk.
- Consider whether you need one-pass or two-pass algorithms.

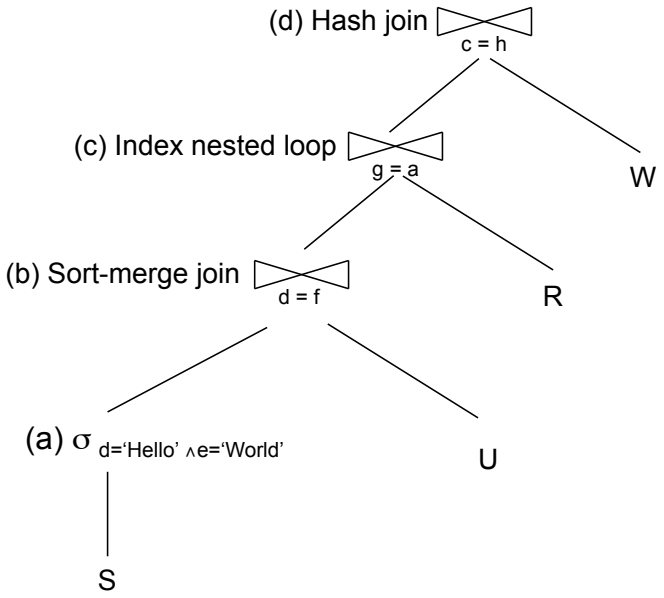

Answer (Compute the cost of the plan):

3 Transactions Concurrency Control

3. (20 points)

   (a) (10 points) Consider the following schedule. Explain what happens when transactions try to execute as per this schedule and the DBMS uses timestamp-based concurrency control.

   $ST_1 \rightarrow ST_2 \rightarrow ST_3 \rightarrow ST_4 \rightarrow R_1(X) \rightarrow R_2(X) \rightarrow W_2(X) \rightarrow W_1(X) \rightarrow W_3(Y) \rightarrow W_2(Y) \rightarrow C_3 \rightarrow W_4(Z) \rightarrow C_4 \rightarrow R_2(Z)$

   Answer (Fill in the table below showing what happens as the transactions execute):

T_1	T_2	T_3	T_4	X	Y	Z
1	2	3	4	RT=0	RT=0	RT=0
				WT=0	WT=0	WT=0
				C=true	C=true	C=true
$R_1(X)$						
...						

(b) (10 points) Consider the following schedule. Explain what happens when transactions try to execute as per this schedule and the DBMS uses **multiversion** concurrency control:

$ST_1 \rightarrow ST_2 \rightarrow ST_4 \rightarrow R_1(X) \rightarrow R_2(X) \rightarrow W_2(X) \rightarrow W_1(X) \rightarrow W_4(Z) \rightarrow R_2(Z)$

Answer

(Fill in the table below showing what happens as the transactions execute):

T_1	T_2	T_4	X_0 ...
1	2	4	
$R_1(X)$			RT=1
...			

4 Transactions Recovery

4. (15 points)

A DBMS uses the ARIES recovery algorithm. The DBMS just crashed. Upon restart, the state of the database on disk is the following:

Log on disk:

LSN	1	2	3
Transaction ID	T1	T2	T2
Previous LSN	-	-	2
Type	Update	Update	Commit
Page ID	P1	P2	-
Description	Updated A...	Updated B...	-

Pages on disk

Page 1 PageLSN=1
A_1

Page 2 PageLSN=0
B_0

(a) (5 points) Show the transactions table and the dirty pages table at the end of the analysis phase:

(b) (5 points) Explain what will happen during the Redo phase: Where will Redo begin? What updates will be performed? What changes will be made to the pages?

(c) (5 points) Explain what will happen during the Undo phase: What updates will be undone? What log entries will be written? What changes will be made to the pages?

5 Distributed and Parallel DBMSs

5. (30 points)

   (a) (10 points) In the presumed-abort two-phase commit protocol, what happens if the coordinator sends PREPARE messages and all but one subordinate vote to commit the transaction. The last subordinate wants to commit the transaction also but it crashes before receiving the PREPARE message from the coordinator.

   Answer (Describe the sequence of operations at the coordinator, at the subordinates that did not crash, and at the crashed subordinate after it recovers):

(b) (10 points) Consider the following relations:

```
Purchases(cid, pid, time)
Products(pid,description,price)
```

where cid is the unique identifier of each customer; pid is the unique identifier of each product; time is the time when a customer bought a product; description is the description of the product, and price is its price.

Assuming that relation Purchases is stored at site A and that Products is stored at site B, describe how a semijoin can serve to efficiently answer the following query:

```
SELECT *
FROM Purchases, Products
WHERE Purchases.pid = Products.pid
AND Products.price > 50
```

Answer (Describe how the query will be answered using a semijoin):

(c) (10 points) What does MapReduce do when some tasks take significantly longer to complete their executions than others. When is this approach helpful? When is it NOT helpful.

Answer (All three subquestions):