

CPS 196.3 Fall 2002

Optional Problem Set #2

Assigned: Friday, December 6

Due: Wednesday, December 11

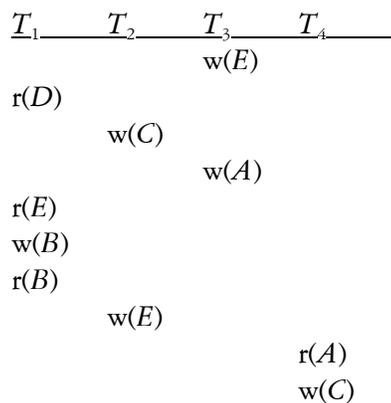
Problem 1.

Consider the join $R \bowtie_{R.A=S.B} S$, given the following information about the tables to be joined. The cost metric is the number of disk I/O's and the cost of writing out the result should be uniformly ignored.

- R contains 10,000 rows and has 10 rows per page.
 - S contains 2,000 rows and also has 10 rows per page.
 - $S.B$ is a key of S .
 - Both tables are stored compactly on disk in no particular order.
 - No indexes are available.
 - 52 memory blocks are available for query processing.
- (a) What is the expected cost of joining R and S using a block-based nested-loop join, with R as the outer table?
- (b) What is the expected cost of joining R and S using a block-based nested-loop join, with S as the outer table?
- (c) What is the expected cost of joining R and S using a sort-merge join? What is the minimum number of memory blocks required for this cost to remain unchanged?
- (d) What is the expected cost of joining R and S using a hash join? What is the minimum number of memory blocks required for this cost to remain unchanged?

Problem 2.

Consider the following schedule.



- (a) Draw the precedence graph.
- (b) Is this schedule conflict-serializable? If not, explain why not. Otherwise, give the equivalent serial schedule.

Problem 3.

For each schedule below, tell whether it is conflict-serializable. If yes, also tell:

- Whether it is recoverable;
- Whether it avoids cascading rollbacks;
- Whether it is possible under *strict* 2PL.

- (a) $T_1.write(B), T_2.read(A), T_2.write(A), T_1.read(A), T_1.write(A), T_1.commit, T_2.commit$
 (b) $T_1.write(B), T_2.read(A), T_2.write(A), T_1.read(A), T_1.write(A), T_2.commit, T_1.commit$
 (c) $T_1.write(B), T_2.read(A), T_2.write(A), T_2.commit, T_1.read(A), T_1.write(A), T_1.commit$
 (d) $T_1.write(B), T_2.read(A), T_1.read(A), T_2.write(A), T_1.write(A), T_2.commit, T_1.commit$
 (e) $T_2.write(B), T_2.read(A), T_2.write(A), T_1.write(B), T_2.commit, T_1.read(A), T_1.commit$

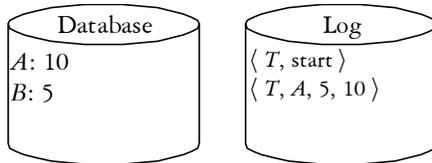
Problem 4.

Suppose we have a transaction T that performs the following two actions:

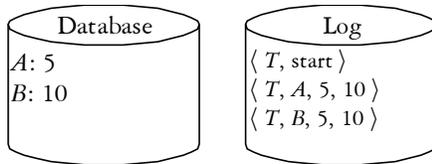
$$A := A + 5; B := B + 5.$$

Say that undo/redo logging is in use, and that initially, $A = 5$ and $B = 5$. For each hypothetical disk state shown below, state whether it is a legal (possible) state for undo/redo logging. If it is not a legal state, explain why not. Recall that the log entries are in the format $\langle transaction_id, variable_id, old_value, new_value \rangle$.

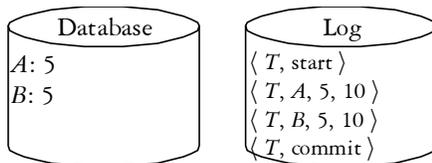
(a)



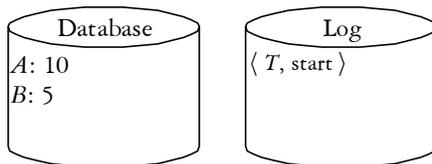
(b)



(c)



(d)



Problem 5.

Consider the following transaction log from the start of the run of a database system that uses undo/redo logging with fuzzy checkpointing:

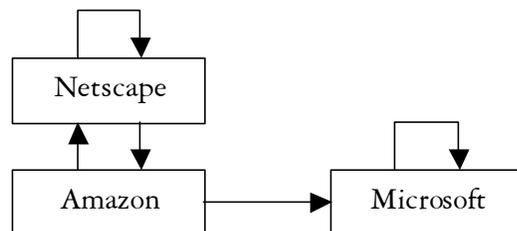
1. $\langle T1, \text{start} \rangle$
2. $\langle T1, A, 45, 10 \rangle$
3. $\langle T2, \text{start} \rangle$
4. $\langle T2, B, 5, 15 \rangle$
5. $\langle T2, C, 35, 10 \rangle$
6. $\langle T1, D, 15, 5 \rangle$
7. $\langle T1, \text{commit} \rangle$
8. $\langle T3, \text{start} \rangle$
9. $\langle T3, A, 10, 15 \rangle$
10. $\langle \text{begin-checkpoint } \{T2, T3\} \rangle$
11. $\langle T2, D, 5, 20 \rangle$
12. $\langle T2, \text{commit} \rangle$
13. $\langle \text{end-checkpoint} \rangle$
14. $\langle T4, \text{start} \rangle$
15. $\langle T4, D, 20, 30 \rangle$
16. $\langle T3, C, 10, 15 \rangle$
17. $\langle T3, \text{commit} \rangle$
18. $\langle T4, \text{commit} \rangle$

What is the value of the data items A , B , C , and D on disk after recovery:

- (a) if the system crashes just before line 6 is written to disk?
- (b) if the system crashes just before line 10 is written to disk?
- (c) if the system crashes just before line 12 is written to disk?
- (d) if the system crashes just before line 13 is written to disk?
- (e) if the system crashes just before line 16 is written to disk?
- (f) if the system crashes just before line 18 is written to disk?

Problem 6.

Consider the following mini-Web consisting of three pages.



- (a) Compute the naïve PageRank values for these three pages using the fixed-point iteration. Start with all PageRank values equal to 1 and carry out at least 10 iterations. Show your intermediate results.
- (b) Compute the practical PageRank values for these three pages (using a delay factor of 0.8). You may solve for the PageRank values directly instead of using the fixed-point iteration.