

CS 533 Fall 2013

Homework Assignment 1 - Due October 7 Monday

1. Consider the following sequences of actions listed in the order they are submitted to the DBMS:

- **Sequence S1:** R1[X], W2[X], W2[Y], W3[Y], W1[Y], C1, C2, C3
- **Sequence S2:** R1[X], W2[Y], W2[X], W3[Y], W1[Y], C1, C2, C3

For each sequence and for each of the following concurrency control mechanisms, describe how the concurrency control mechanisms handle the sequence. Add lock and unlock requests to the previous sequence of actions as per the locking protocol. The DBMS processes actions in the order shown. If a transaction is blocked, assume that all its actions are queued until it is resumed; the DBMS continues with the next action (according to the listed sequence) of an unblocked transaction.

- (a) Strict 2PL with deadlock detection. (Show the waits-for graph in case of deadlock.)
 - (b) Conservative (and Strict) 2PL.
2. Consider a database organized in terms of the following hierarchy of objects: The database itself is an object (D), and it contains two files ($F1$ and $F2$), each of which contains 1000 pages ($P1 \dots P1000$ and $P1001 \dots P2000$, respectively). Each page contains 100 records, and records are identified as $p : i$, where p is the page identifier and i is the slot of the record on that page. Multiple granularity locking is used, with RL , WL , IRL , IWL , and $RIWL$ locks, and database-level, file-level, page-level, and record-level locking. For each of the following operations, indicate the sequence of lock requests that must be generated by a transaction that wants to carry out just these transactions.
- (a) Read record $P1200 : 5$
 - (b) Read records $P1200 : 98$ through $P1205 : 2$
 - (c) Read pages $P10$ through $P980$
 - (d) Read all pages in file $F1$ and modify 10 pages
 - (e) Delete record $P1200 : 98$ (This is a blind write.)

- (f) Delete all records.
3. Consider the 3 phases of recovery shown on Slide 30 of Lecture 8. Let C be the LSN where the *Analysis* phase starts, B be the LSN where the *Redo* phase starts, A be the LSN where the *Undo* phase end. Answer the following questions with respect to this slide.
- (a) Can A come after C ? If no, explain why not. If yes, then give an example scenario (show the log records explicitly) where A comes after C .
 - (b) Can A come after B ? If no, explain why not. If yes, then give an example scenario where A comes after B .
 - (c) Can B come after C ? If no, explain why not. If yes, then give an example scenario where B comes after C .
4. Sometimes the logged action is not redone during the *Redo* phase of the Aries recovery algorithm as described in Slide 61 of Lecture 8. Specifically, it speaks of 3 scenarios when the logged action is not redone.
- (a) Why is it cheaper to test the first two conditions?
 - (b) Describes an execution that illustrates the use of the first condition.
 - (c) Describes an execution that illustrates the use of the second condition.

Please keep in mind:

- This assignment is to be done individually. The honor code is in effect.
- Distance students must submit your homework through RamCT. On campus students must submit your assignment via email (cs533@cs.colostate.edu and cc to tmoataz@cs.colostate.edu).
- Assignments must be submitted before 11:59 p.m. October 7, 2013. No late work will be accepted except in exigent situations.
- Your answers must be typed.