

Free University of Bozen
Faculty of Computer Science
Alessandro Artale

Databases 2
Revision Exercise
Recovery Failure and Concurrency

1 Transactions and Log Files

You have been given the database elements X, Y, Z with starting values: $X = 4, Y = 5, Z = 1$, and two transactions T_1, T_2 :

$T_1.$ $X := 2 * X; Y := Y + X$

$T_2.$ $Z := 2 * Z; X := X - Z$

Assume that transaction T_2 starts after T_1 finishes, and that an UNDO logging technique is used to recover from a system failure. Answer the following questions:

- 1.a. Give a table containing: a) the primitive actions for both T_1 and T_2 in the appropriate sequential order (when giving the actions use the primitive actions READ, WRITE, OUTPUT, FLUSH LOG, together with the required assignment statements for the local variables.); b) the values of the local variables; c) the values for X, Y, Z both in main memory and on disk; d) the status of the log in main memory.
- 1.b. Give a table like in part 1.a assuming that the recovery technique is a REDO logging technique, and the execution order is inverted, i.e., transaction T_1 starts after T_2 finishes.

2 Recovery from System Failure

Consider the following log file (to be read from left to right):

< START T_1 > < $T_1, X, 0$ > < $T_1, Y, 0$ > < START T_2 > < $T_2, Z, 0$ > < $T_1, W, 5$ > < COMMIT T_1 >
< START T_3 > < $T_3, Y, 10$ > < START T_4 > < $T_4, X, 5$ > < COMMIT T_2 > < $T_4, W, 10$ >
< $T_4, V, 5$ > < COMMIT T_4 > < $T_3, V, 10$ > < COMMIT T_3 >

Suppose that we start a Nonquiescent Checkpoint under an UNDO logging technique immediately after one of the following log records has been written:

- i. < $T_1, Y, 0$ >
- ii. < $T_2, Z, 0$ >
- iii. < START T_3 >
- iv. < COMMIT T_2 >
- v. < $T_4, V, 5$ >

- 2.a. For each of the above points, complete the log file with the insertion of START CKPT and END CKPT log records.

Consider now the log file as after the above completion under point ii. Suppose there is a crash. Using the primitive actions: WRITE, OUTPUT, ABORT, give the actions that need to be done to recover from the crash in each of the following situations:

- 2.b. The last record in the log file is < COMMIT T_4 >.

- 2.c. The last record in the log file is < $T_4, X, 5$ >.

Furthermore, for each of the above points, give the new log file after deleting useless portions.

Suppose now that we start a Nonquiescent Checkpoint under a REDO logging technique and the log file is as follows:

< START T_1 > < $T_1, X, 3$ > < $T_1, Y, 5$ > < START T_2 > < $T_2, Z, 12$ > < $T_1, W, 7$ > < COMMIT T_1 >
< START T_3 > < $T_3, Y, 16$ > < START CKPT(T_2, T_3) > < START T_4 > < $T_4, X, 32$ > < COMMIT T_2 >
< $T_4, W, 10$ > < END CKPT > < $T_4, V, 28$ > < COMMIT T_4 > < $T_3, V, 14$ > < COMMIT T_3 >

Using the primitive actions: WRITE, OUTPUT, ABORT, give the actions that need to be done to recover from a crash in each of the following situations:

2.d. The last record in the log file is $\langle \text{COMMIT } T_4 \rangle$.

2.e. The last record in the log file is $\langle T_4, X, 32 \rangle$.

As in the UNDO case, for each of the above points, give also the new log file after deleting useless portions.

3 Concurrency Control

For each of the following two schedules:

S1. $r_3(X); r_1(Y); r_2(Z); r_2(W); r_3(Y); r_1(Z); w_1(Y); r_3(W); w_3(W); w_2(Y); w_2(Z); w_3(X)$.

S2. $r_4(Y); r_2(X); w_2(X); r_1(X); w_4(Y); r_1(Y); r_2(Z); w_1(Y); w_2(Z); r_1(Z); r_3(Y); w_1(Z); w_3(X)$.

Answer the following questions:

3.a. Give the precedence graph and check whether the schedule is conflict-serializable. Justify each precedence relation in the graph, i.e., for each arch in the obtained graph show a conflicting pair of actions which gives rise to such an arch.

3.b. If the schedule is conflict-serializable show all the admissible (given the precedence graph) serial schedules.

Given now the following schedule:

S. $r_2(X); r_4(Y); r_1(X); w_4(Y); r_2(Z); r_1(Y); r_3(Y); w_3(Y); w_2(X)$.

3.c. Assuming that the scheduler deals with Two-Phase Locking with shared and exclusive locks and allowing upgrading (i.e., a shared lock can be upgraded to an exclusive lock), add shared and exclusive locks and unlocks to **S**. Insert each shared and exclusive lock as delayed as possible, while unlock as soon as possible (always in accordance with the Two-Phase Locking strategy). Give **S** in a table with four columns T_1, T_2, T_3, T_4 .

Given now the following schedule:

S. $r_4(Y); r_2(X); r_1(X); w_4(Y); r_1(Y); r_2(Z); w_2(X); w_1(Y); r_1(Z); r_3(Z); r_3(X); w_1(Z); w_3(Z)$

3.d. Assume that the scheduler deals with Two-Phase Locking with shared and exclusive locks and allowing upgrading (i.e., a shared lock can be upgraded to an exclusive lock). Adding shared and exclusive locks and unlocks to **S**, show that the schedule end with a deadlock and must be aborted.