

Problem 1

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

< START T₁ > < T₁, A, 5 > < T₁, B, 3 > < START T₂ > < T₂, B, 6 >
< START T₃ > < T₃, B, 18 > < T₂, D, 0 > < T₁, A, 10 > < COMMIT T₁ >
< T₃, C, 5 > < COMMIT T₃ > < START T₄ > < T₄, A, 20 > < T₂, E, 0 >
< COMMIT T₂ > < T₄, F, 0 > < COMMIT T₄ >

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

1. < T₁, B, 3 >;
2. < T₂, B, 6 >.

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

Nonquiescent Checkpointing

The steps in a *nonquiescent checkpointing* are:

1. Write $\langle \text{START CKPT}(T_1; \dots; T_k) \rangle$ record to log, where T_i are all the active (uncommitted) transactions;
2. Wait until all transactions $T_1; \dots; T_k$ commit or abort, *but do not prohibit new transactions*;
3. Write $\langle \text{END CKPT} \rangle$ record to log, and flush the log.

If recovery is necessary, we know that all transactions prior to a recorded checkpoint have committed and need not be undone.

Problem 2

Consider now the log file obtained solving point 2 of Problem 1.

Suppose there is a crash. Using the primitive actions: `WRITE`, `OUTPUT`, `ABORT(T)`, give the actions that need to be done to recover from the crash in each of the following situations:

1. The last record in the log file is $\langle T_4, F, 0 \rangle$.
2. The last record in the log file is $\langle \text{COMMIT } T_1 \rangle$.

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

Recovery With Nonquiescent Checkpoints

To recover a crash using an Undo Nonquiescent Checkpoints: Scan the log from the end:

- If we first meet $\langle \text{END CKPT} \rangle$ then we can restrict to transactions that began after the $\langle \text{START CKPT}(T_1; \dots; T_k) \rangle$. The log before $\langle \text{START CKPT}(T_1; \dots; T_k) \rangle$ is useless and can be deleted.
- If we first meet $\langle \text{START CKPT}(T_1, \dots, T_k) \rangle$, then the crash occurred during the checkpoint. We need to undo:
 1. All those transactions T with $\langle \text{START } T \rangle$ after the $\langle \text{START CKPT} \rangle$ but no $\langle \text{COMMIT } T \rangle$;
 2. All transactions T_i on the list associated with $\langle \text{START CKPT} \rangle$ with no $\langle \text{COMMIT } T_i \rangle$. The log before the start of the earliest of these incomplete transactions can be deleted.

Problem 3

The following is a sequence of redo-log records written by three transactions T, U, V :

```
< START T >< T, A, 10 >< START U >< U, B, 20 >  
< T, C, 30 >< START CKPT(T, U) >< U, D, 40 >< COMMIT U >  
< T, E, 50 >< START V >< V, C, 45 >< END CKPT >  
< COMMIT V >< T, D, 45 >
```

Suppose there is a crash and the log file is as above. Using the primitive actions: `WRITE`, `OUTPUT`, `ABORT(T)`, give the actions that need to be done to recover from the crash using a REDO-Logging technique.

Recovery for Redo Logging

- While incomplete transactions did not change the DB, committed transactions cause problems since we don't know which of their changes have been written to disk.
- **To recover a crash using a Redo-Logging we do:**
 1. Find the set of committed transactions from the log;
 2. Scan the log *forward* from the beginning and for each $\langle T, X, v \rangle$ do:
 - If T is committed, write value v for X to disk, i.e., do:
`WRITE(X, v); OUTPUT(X);`
 - Otherwise, do nothing.
 3. For each incomplete transaction T add $\langle \text{ABORT } T \rangle$ to the log, and flush the log.