

# CS-245 Database System Principles - Winter 2001

- PLEASE write your serial number on the top of your first page.
- This assignment is due in class on Thursday, March 1<sup>st</sup>.
- State all assumptions.
- Email questions to [cs245ta-win01@lists.stanford.edu](mailto:cs245ta-win01@lists.stanford.edu).

## Assignment 6

### Problem 1 (25 points)

Consider a system that uses *undo logging* and produces the following sequence of log records (in exactly this order):

`<START>; <T, A, 10>; <T, B, 20>; <T, C, 30>; <T, D, 40>; <COMMIT T>`

One possible sequence of actions that leads to these log records is:

`Log(A); Log(B); Write(A); Log(C); Write(B); Log(D); Write(C); Write(D); Log(Commit);`

In this sequence, `Log(X)` represents the action that writes to disk the log record for action X, and `Write(X)` represents the action that updates object X on disk.

Give *all* possible legal sequences of actions (with undo logging) that satisfy all of the following three conditions:

- the sequence of actions can yield the given log sequence;
- `Write(A)` is before `Log(C)`;
- `Write(B)` is before `Log(D)`.

To give the sequences, use the same notation as for the example action sequence.

### Problem 2 (30 points)

The following is a sequence of undo-log records written by two transactions *T* and *U*:

`<START T>; <T, A, 10>; <START U>; <U, B, 20>; <T, C, 30>; <U, D, 40>; <COMMIT U>; <T, E, 50>; <COMMIT T>.`

Describe the actions of the recovery manager, including changes to both disk and the log, if there is a crash and the last log record that appears on disk is:

- a) `<START U>`
- b) `<T, E, 50>`
- c) `<COMMIT T>`

### Problem 3 (25 points)

Consider a system that uses *redo logging* and produces the following sequence of log records (in exactly this order):

```
<START>; <T, A, 10>; <T, B, 20>; <T, C, 30>; <T, D, 40>; <COMMIT T>
```

Give *all* possible legal sequences of actions (with redo logging) that satisfy all of the following three conditions:

- the sequence of actions can yield the given log sequence;
- Write(A) is before Write(C);
- Log(B) is before Log(D).

To give the sequences, use the notation of Problem 1.

### Problem 4 (20 points)

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

```
1) <START T1>
2) <T1, A, 50, 10>
3) <START T2>
4) <T1, B, 130, 10>
5) <T1, A, 70, 50>
6) <T2, C, 20, 10>
7) <T2, D, 30, 10>
8) <COMMIT T1>
9) <START T3>
10) <T3, E, 60, 10>
11) <T2, D, 40, 30>
12) <START CKPT (T2,T3)>
13) <T2, C, 70, 20>
14) <COMMIT T2>
15) <START T4>
16) <T4, F, 100, 10>
17) <T4, G, 110, 10>
18) <COMMIT T3>
19) <T4, F, 150, 100>
20) <END CKPT>
21) <T4, F, 140, 150>
22) <COMMIT T4>
```

Assume the log entries are in the format

```
<Tid, Variable, New value, Old value>
```

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

- a) if the system crashes just before line 10 is written to disk?
- b) if the system crashes just before line 15 is written to disk?
- c) if the system crashes just before line 16 is written to disk?
- d) if the system crashes just before line 19 is written to disk?
- e) if the system crashes just before line 22 is written to disk?
- f) if the system crashes just after line 22 is written to disk?