A excellent!

Dan Ports <drkp@mit.edu>
6.033 handson #7 - 2004/04/29
Karger TR1

1. The database doesn't show studentB because studentB was created by transaction 2, which was not committed.

2. When the database recovers, studentA should exist with a balance of 900 and studentC should exist with a balance of 3100, because these are the effects of transactions 1 and 3, which were committed; studentB should not exist.

3. The DB file does not contain the updates made by transaction 3 (creating studentC and debiting studentA) because this transaction has not been installed (via the END command)

4. We expect the DB to contain studentA with a balance of 900 and studentC with a balance of 3100, because these correspond to the effects of transaction 1 (which is committed and installed), and transaction 3 (which was committed before the crash and installed during the recovery procedure).

5. This is indeed the state of the database.

6. A "done" transaction is one that was committed and installed, and thus has an END record. A "winner" is one that was committed but not installed, and thus has an OUTCOME log record. A "loser" is a transaction that has not yet been committed and thus has no OUTCOME record.

7. The checkpoint command doesn't install into the database, it simply records which actions are winners, losers, and done at a given point.

8. The log only needed to be rolled back 6 lines. This is the advantage of the checkpoint: the log needs only be rolled back to the checkpoint.

9. The database is restored to the same state. Recovery is idempotent.

10. In the second recovery, transaction 3 is done, whereas it was a winner in the first. This is because the first recovery installs the results of transaction 3 (which was committed but not installed before the crash) and writes an END record for it; the second recovery finds it already installed.

$Id: handson7.txt,v 1.1 2004/04/29 09:33:10 dan Exp $