6.033 Computer System Engineering

Spring 2009

Recoverability

Recoverable sector

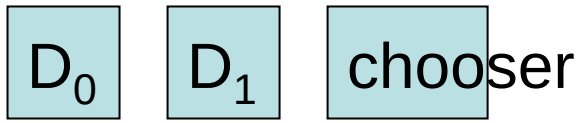

Larger actions?

Commit point

```
id ← begin_ra()
                                ABORT()

     COMMIT()

     end_ra()
```

Version histories

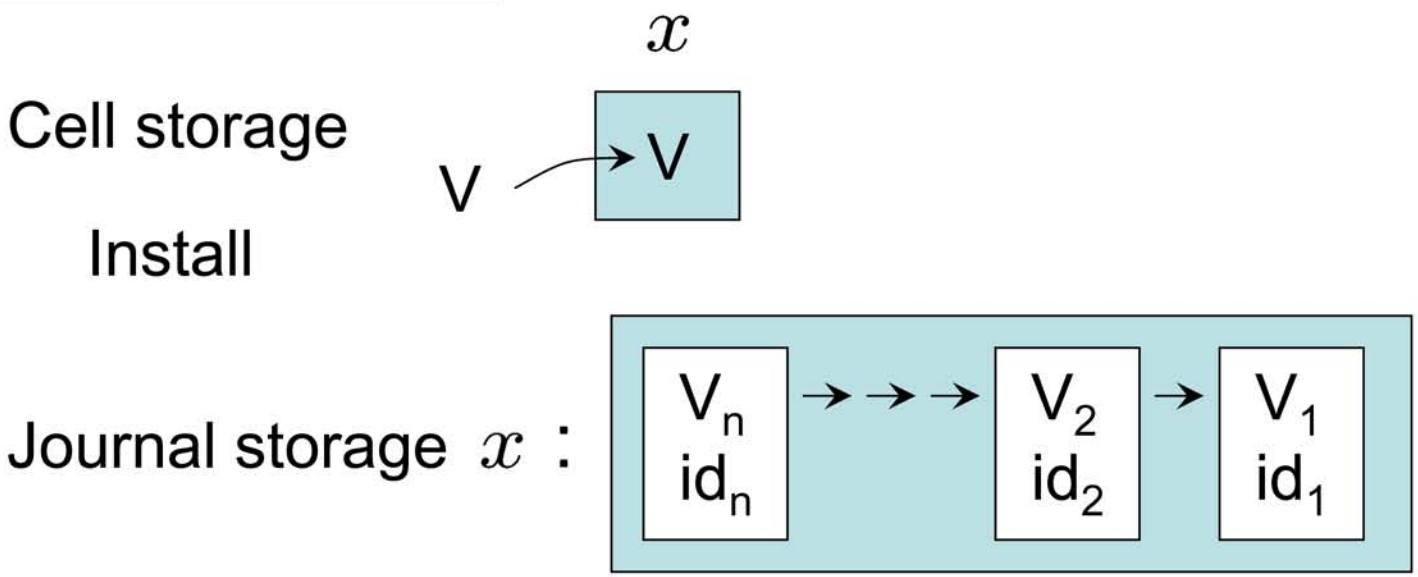


```
WriteJournal(item x, v, id)
```

`ReadJournal(item x, id)`

Commit record table

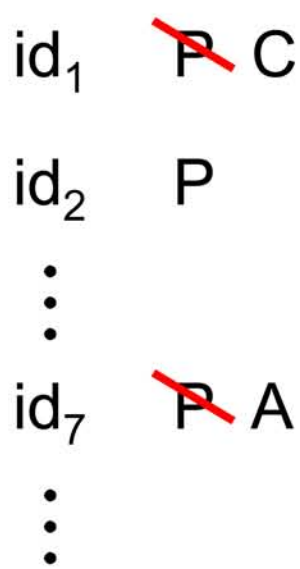

1) Bootstrap
2) Slow

Logging:

Cell st. : read / write

Log : non-volatile + sequential

Plan:

1) Fail ➔ Recover from log

→ uncommitted ➔ back out
(undo)

→ committed ➔ install
(redo)

2) **`ABORT()`** ➔ undo cur. action (redo)

Append-only:

```
type: UPDATE
id: 172
undo: x ← old
redo: x ← new
```

```
type: OUTCOME
id: 174
status:
    COMMITTED
```

1) When to write log?

2) How to recover?

Disk-bound DB:

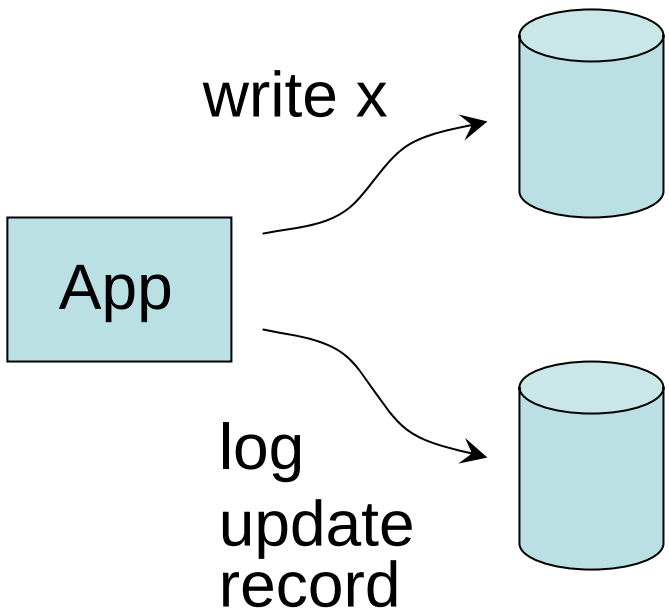

1) WAL protocol:
   Write ahead logging

2) Log COMMIT record
   before returning from
   `commit()`

Recovery:

1) Scan log backwards
2) Winners: COMMITTED or ABORTED
   Losers: Everything else
3) Redo COMM. Winners
   Undo Losers