

# **CSE 451: Operating Systems**

## **Spring 2010**

### **Module 17**

## **Journaling File Systems**

**John Zahorjan**  
**zahorjan@cs.washington.edu**  
**Allen Center 534**

# File System Consistency (Not Performance)

- Buffering is necessary for performance
- Suppose a crash occurs during a file creation:
  1. Allocate a free inode
  2. Point directory entry at the new inode
- In general, after a crash the disk data structures may be in an inconsistent state
  - metadata updated but data not
  - data updated but metadata not
  - either or both partially updated
- fsck (i-check, d-check) are *very* slow
  - must touch every block
  - worse as disks get larger!

# Journaling file systems

- Became popular ~2002
- There are several options that differ in their details
  - Ext3, Ext4, ReiserFS, XFS, JFS, ntfs
- Basic idea
  - update metadata (and possibly all data), *transactionally*
    - "all or nothing"
  - if a crash occurs, you may lose a bit of work, but the disk will be in a consistent state
    - more precisely, you will be able to quickly get it to a consistent state by using the transaction log/journal – rather than scanning every disk block and checking sanity conditions

# Where is the Data?

- In the file systems we have seen already, the data is in two places:
  - On disk
  - In in-memory caches
- The caches are crucial to performance, but also the source of the potential “corruption on crash” problem
- The basic idea of the solution:
  - Always leave “home copy” of data on disk in a consistent state
  - Make updates persistent by writing them to a sequential (chronological) journal partition/file
  - At your leisure, push the updates (in order) to the home copies and reclaim the journal space

# Redo log

- Log: an append-only file containing log records
  - $\langle \text{start } t \rangle$ 
    - transaction  $t$  has begun
  - $\langle t, x, v \rangle$ 
    - transaction  $t$  has updated block  $x$  and its new value is  $v$ 
      - Can log block “diffs” instead of full blocks
  - $\langle \text{commit } t \rangle$ 
    - transaction  $t$  has committed – updates will survive a crash
- Comments
  - Committing involves writing the redo records – the home data needn't be updated at this time
  - No guarantees about the consistency of the file data, as seen by the application. (This is just how things were to begin with...)

# If a crash occurs

- Recover the log
- Redo committed transactions
  - Walk the log in order and re-execute updates from all committed transactions
  - Aside: note that update (write) is *idempotent*: can be done any positive number of times with the same result.
- Uncommitted transactions
  - Ignore them. It's as though the crash occurred a tiny bit earlier...
    - Yes, you still lose some recent updates

# Managing the Log Space

- A cleaner thread walks the log in order, updating the home locations of updates in each transaction
- Once a transaction has been reflected to the home blocks, it can be deleted from the log
  - What if we crash between updating the home blocks and deleting the log?

# Impact on performance

- The log is a big contiguous write
  - very efficient
- And you do fewer synchronous writes
  - very costly in terms of performance
- So journaling file systems can actually improve performance (immensely)
- As well as making recovery very efficient



# What About Deadlock?

- Why might it arise?
- What would you do about it?

# Want to know more?

- CSE 444! This is a direct ripoff of database system techniques
- “New-Value Logging in the Echo Replicated File System”, Andy Hisgen, Andrew Birrell, Charles Jerian, Timothy Mann, Garret Swart
  - <http://citeseer.ist.psu.edu/hisgen93newvalue.html>