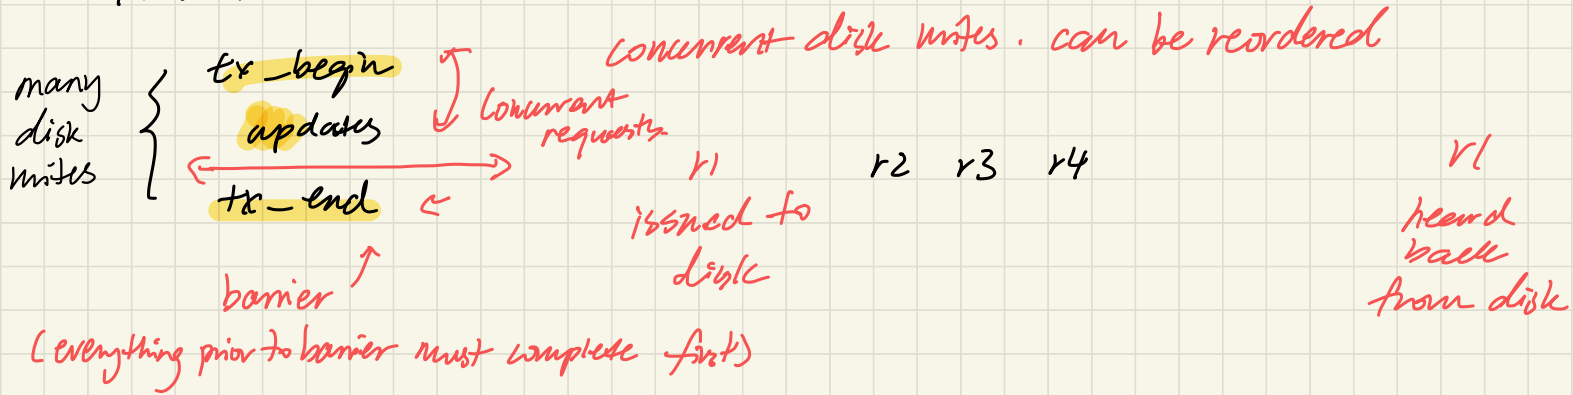


5/20/24

Transactions Wrapup & Copy on Write FS

Transactions



→ resolve the incomplete txn

→ compute a checksum (write as part of tx_end).


→ on recovery, validate committed txns' checksum before applying them

★ Logging isn't the only way to achieve crash safety!


Copy-on-write Filesys

→ **in place update** :

→ can't overwrite
multiple disk blocks
atomically



inode table reserved at
fixed location



(txn & journaling)

→ achieve crash consistency by
writing changes elsewhere first

→ **copy-on-write** :

→ write changed blocks to new locations

→ once all new blocks are written, use a **single**

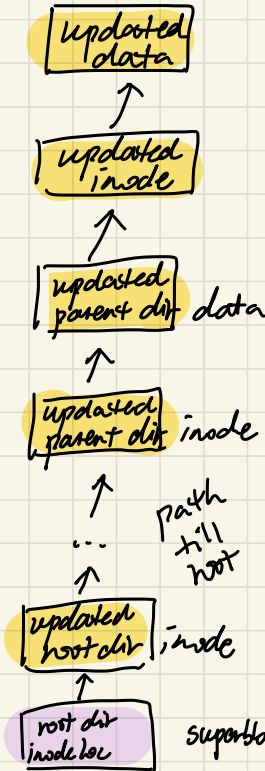
disk write to make all the new blocks visible at once

★ no fixed location for almost all structures & data blocks,
data & metadata locations change after every update

Cow Filesys

→ recursive update problem:

if directory entry stores actual location of the metadata



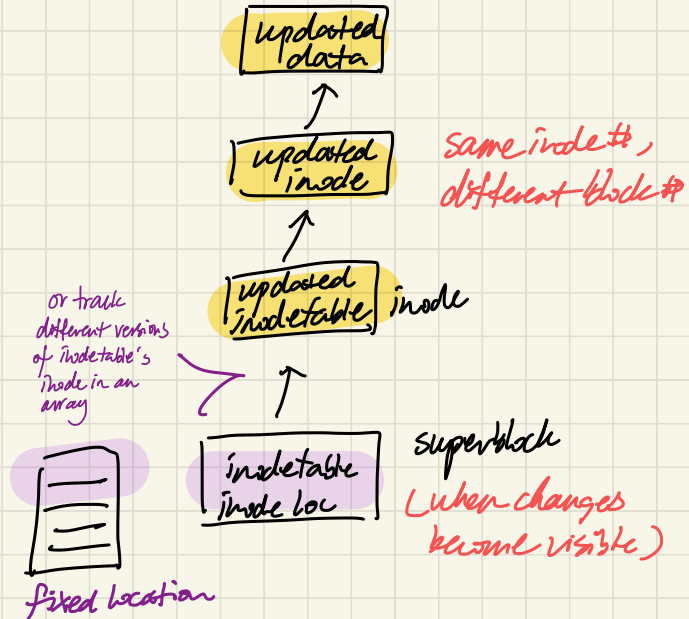
name / block #

* tracking physical pointer in a cow system causes recursive updates!

⇓
use logical pointer (inode #)

* requires inode # to blk # translation!

with logical pointer

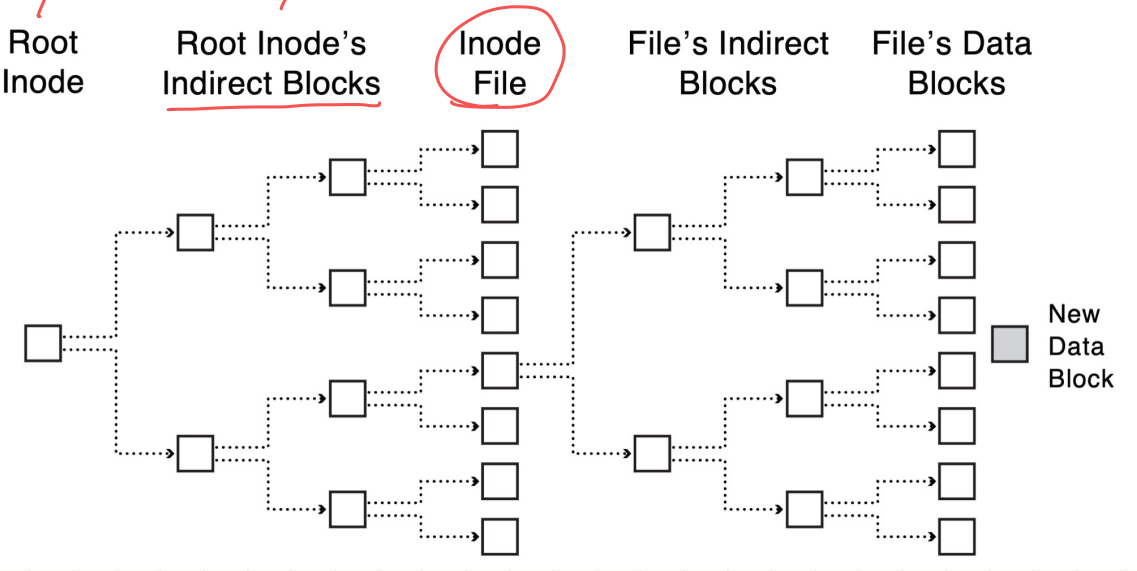


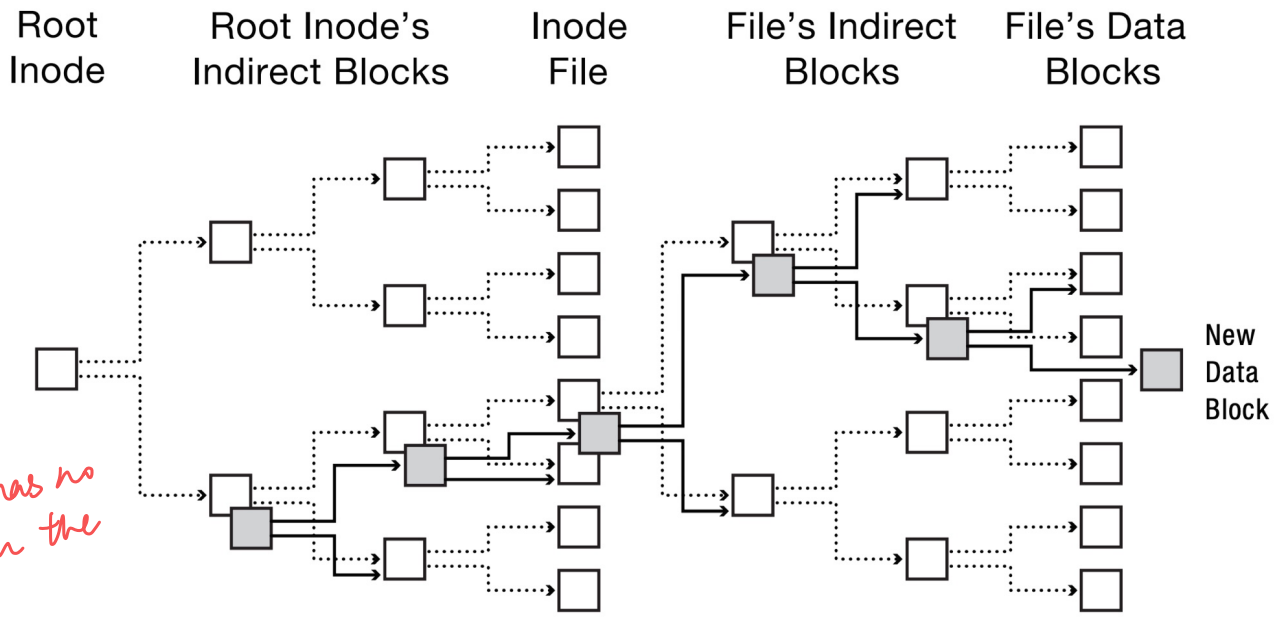
★ data layout via indexed blocks

metadata for inodetable

data pointers

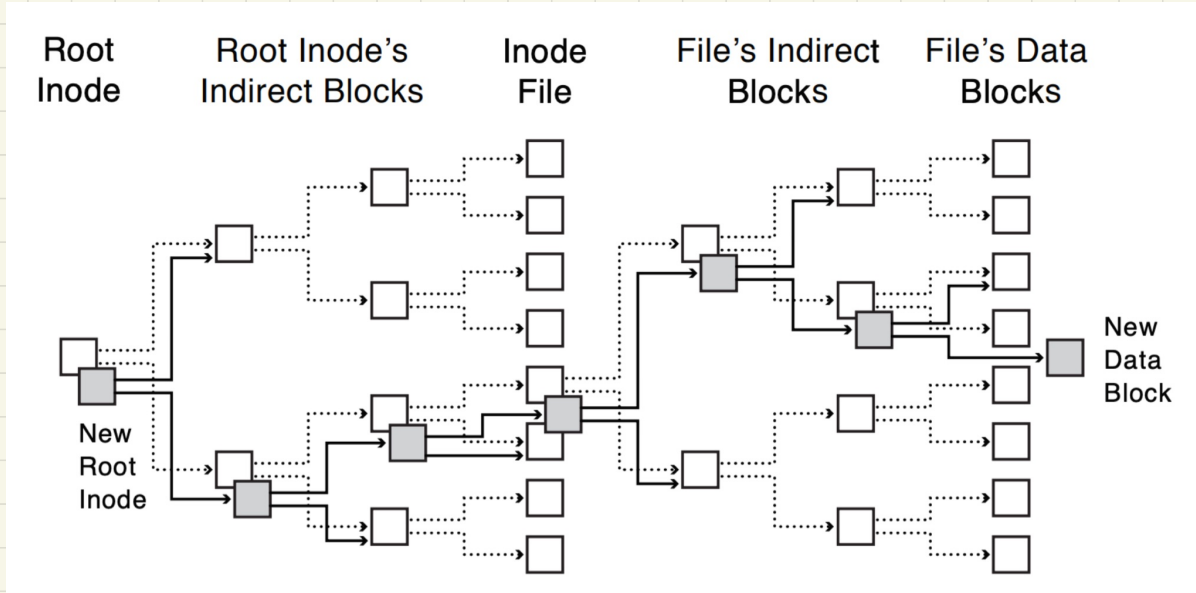
inodetable data

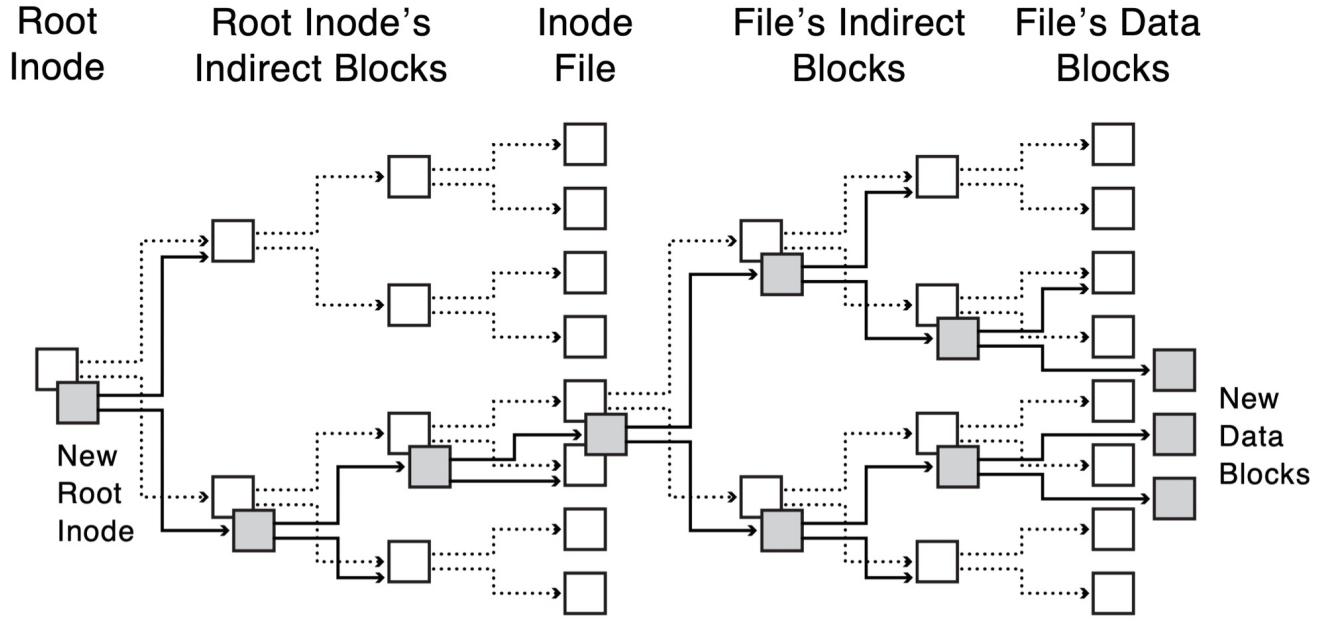




crash has no effect on the file sys

→ lots of overhead (many blocks to change for a single write)
→ batch updates!





☆ provides crash safety w/out logging, but often times a log is still used to speed up fsync.