

Recitation 6: Filesystems

Kai Mast

Filesystem Abstraction

ext4
(mounted to /)

btrfs
(mounted to /home/kai)

Partition 1
(/dev/ubuntu-vg/
disk1)

Partition 2
(e.g. /dev/ubuntu-vg
/disk2)

Virtual Disk
(e.g. /dev/ubuntu-vg)

Disk 1
(e.g. /dev/sda)

Disk 2
(e.g. /dev/sdb)

One size fits it all?

Assuming a spinning disk, which one performs well for...

- Doing many “random” writes?
- Reading many small-files?
- Being space-efficient
- Unpredictable workloads?

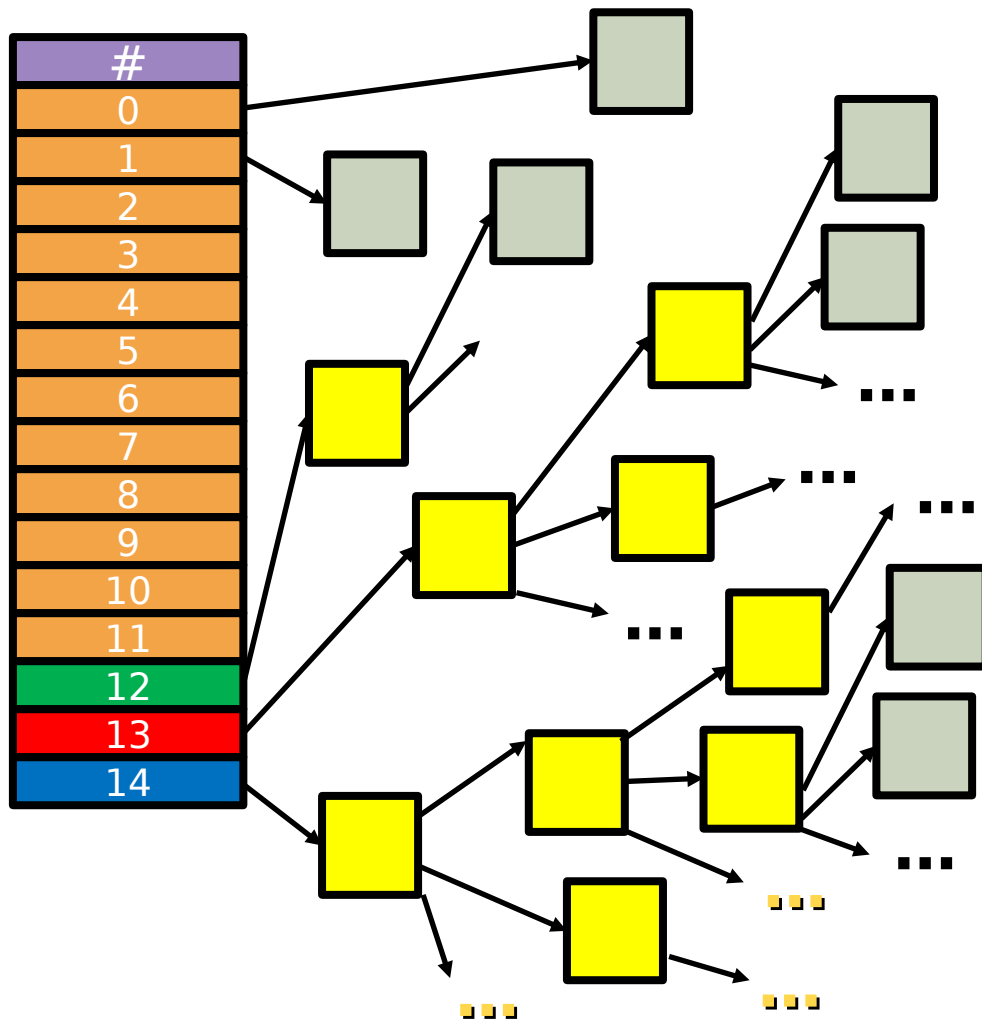
Some options: NTFS, FAT, ZFS...

Filesystem Caches

Simple idea: Keep data in memory to avoid expensive reads from disk and batch writes.

- Why batching writes?
- What could possibly go wrong?


Unix Filesystem Inode



Legend



:Data block



:Indirect block



:Indirect pointer



:Triple Indirect pointer



...

...

Practice Question 1

A disk has a partition on it, subdivided into blocks of 2^{13} bytes.

You want to put a Unix-like file system on the partition, with one superblock in position 0, followed by a sequence of blocks filled with i-nodes. Each i-node is $128 = 2^7$ bytes. You want to have enough i-nodes to store 2^{20} files.

How many i-node blocks are needed?

$$(2^{20} \times 2^7) / 2^{13} = 2^{14}$$

Practice Question 2

A block pointer identifies a block on the partition, and is 4 bytes long (enough to identify 2^{32} blocks). An “indirect block” (a block filled with block pointers) can have $8192 / 4 = 2048$ (2^{11}) block pointers.

Suppose now that an i-node contains 13 block pointers. The first 10 point to the first 10 data blocks. The next three point to an indirect block, a double indirect block, and a triple indirect block. The maximum file size can be approximated by just the number of data blocks reachable from the triple indirect block pointer (the rest is negligible).

In theory, how much data (in bytes) could be accessed from the triple indirect block pointer in the i-node? For this question, assume the size of the disk is unbounded.

$$(2^{11})^3 * 2^{13} = 2^{46}$$

Practice Question 3

Assume now that the file system cache is empty except for the superblock. Assume the file with i-node #2015 has the string “Hello World” in it (that is, the file is just 11 bytes long).

How many disk accesses would be necessary to read the contents of this file, given that you (and the kernel) know the i-node number?

- 1) read inode-block containing inode #2015 to find the data block number
- 2) read the data block itself

Practice Question 4

In reality this same file's i-node number has to be retrieved first. Suppose the name of the file is /etc/test.txt. Assume that the contents of each directory fits in a single block. The root directory / is described in i-node #2 by convention. Assume /etc is in i-node #5 (you know this, but the kernel doesn't).

Again, assuming only the superblock is in the cache and a cache large enough so the same block never has to be read more than once, how many disk accesses are required to read the file?

- 1)read block with inode of root directory (#2) to find location of root directory
- 2)read root directory to find inode # of etc. (#5)
(block with /etc inode already in cache)
- 3)read etc directory to find inode # of test.txt
- 4)read inode block with inode #2015
- 5)read the content of the file itself

Practice Question 5

File /etc/shakespeare.txt (which you know to be in i-node #7, but the kernel doesn't) contains the complete works of Shakespeare (2^{22} bytes or about 4Megabytes).

Assuming only the superblock is in the cache, how many disk accesses are required to retrieve the whole thing?

- read block with inode of root directory (#2) to find location of root directory
- read root directory to find inode # of etc
- *(inode block with etc inode (#5) to find location of etc directory already in cache)*
- read etc directory to find inode # of shakespeare.txt
- *(inode block with shakespeare inode (#7) already in cache)*
- read the first 10 direct blocks
- read the indirect block
- read the remaining $512 - 10 = 502$ blocks

=> 516 in total

Practice Question 6

Suppose somebody wants to add the text "All's Well That Ends Well." to the end of the complete works of Shakespeare (the new text will be contained in a new data block at the very end).

Suppose that the file system has only the superblock in its cache and can allocate free blocks without going to the disk.

How many disk reads and how many disk writes are necessary (assuming a "write-through" cache? (Assume that among the i-nodes only i-node #7 has to be updated for this operation.)

Reads:

- block with i-node of root directory (#2) to find location of root directory
- root directory to find inode # of etc (#5)
- inode block with etc inode (#5) to find location of
- etc directory to find inode # of shakespeare.txt
- the indirect block

Writes:

- allocate and write the new data block
- the updated indirect block
- update size and last modified time in inode #7 and write the inode block