

How does it look on the disk?

Disks are Sector addressable (not byte addressable)

a set of sectors \rightarrow Block

Simple Disk



each inode = 256 byte

each block = 4KB $\Rightarrow 5 \times 4KB = 20KB$ inode data $\Rightarrow 80$ inodes

- we want to access inode #32 $\Rightarrow 32 \times 256 = 8KB$

8KB + 12KB = 20KB. considering sectors of size 512

$$\text{Disk read} \Rightarrow \frac{20 \times 1024}{512} = \text{Sector } 40$$

$$5 \times 8 = 40 \text{ sectors}$$

Read sector 40

Remzi What is a directory? what is an inode?

How are files accessed on a FS? pointer to the blocks of file.

Directory Organization : entry name, inode num

where is it? usually treated as a special type of file!

⇒ directory has an inode with the type set as "directory"

How to implement it?

1) Linear list: - simple to program
- not efficient

- directories are searched quite often → keeping a sorted version is helpful, but maintaining a sorted list is not cheap.
Linear search: costly

- Some systems cache directories to speed up the process.
(in RAM rather than disk)

- Balanced tree is a better option.

2) Hash Table: - lower directory search time.

- collisions
- fixed size

Example: assuming that the file system has been mounted in the superblock is already in memory (everything else, inodes, directories on ^{disk})

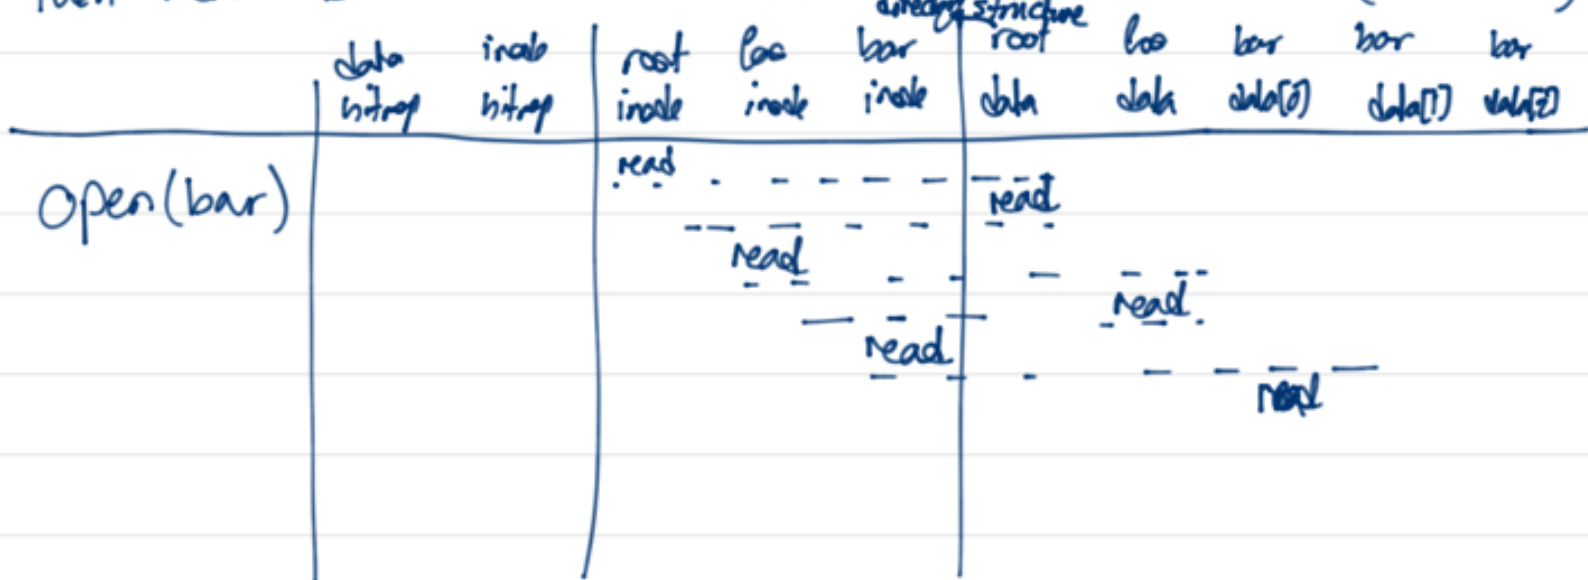
- we want to open a file /foo/bar read it & then close it

Open (" /foo/bar", O_RDONLY)

each directory has an inode itself, usually the root directory has a well-known inode # (2 in Linux) \Rightarrow after accessing inode 2

we have to load the directory of foo to read its inode val

then read bar file inode from data within foo (data block)



\Rightarrow Linear vs Hashed \Rightarrow faster access

Important challenge: Allocation Methods:

Compact

1) Contiguous allocation:
 ↗ Sequential
 ↘ direct.

Drawbacks: -waste of space (external fragmentation)

Extent-based (segmentation) - waste of space (internal " ")? user has to estimate size of the file or for a dynamic size file ⇒ overestimate allocate disk blocks in extents. ⇒ waste.

Assuming sequential access to file blocks and to avoid disk head movements



direct entry:



2) Linked allocation (only sequential)

unless (FAT) → cached

overhead: address of next block

sequential ✓

direct access X

→ store indexes in a different DS:

File Alloc. Table



- how to add a block to a file?

- how does sequential access is affected? - more disk head movements - unless cached