

Review Session 5

- 0. Record
 - 1. Background Knowledge
 - 2. Lab 5 overview
 - 3. Final Exam study tips
 - 4. Q & A.
-

1. Background Knowledge.

cd ~/a/b.txt.

File:

User's perspective: named bytes/data on hardware storage.

File System's perspective: group of disk blocks.

Key Abstraction:

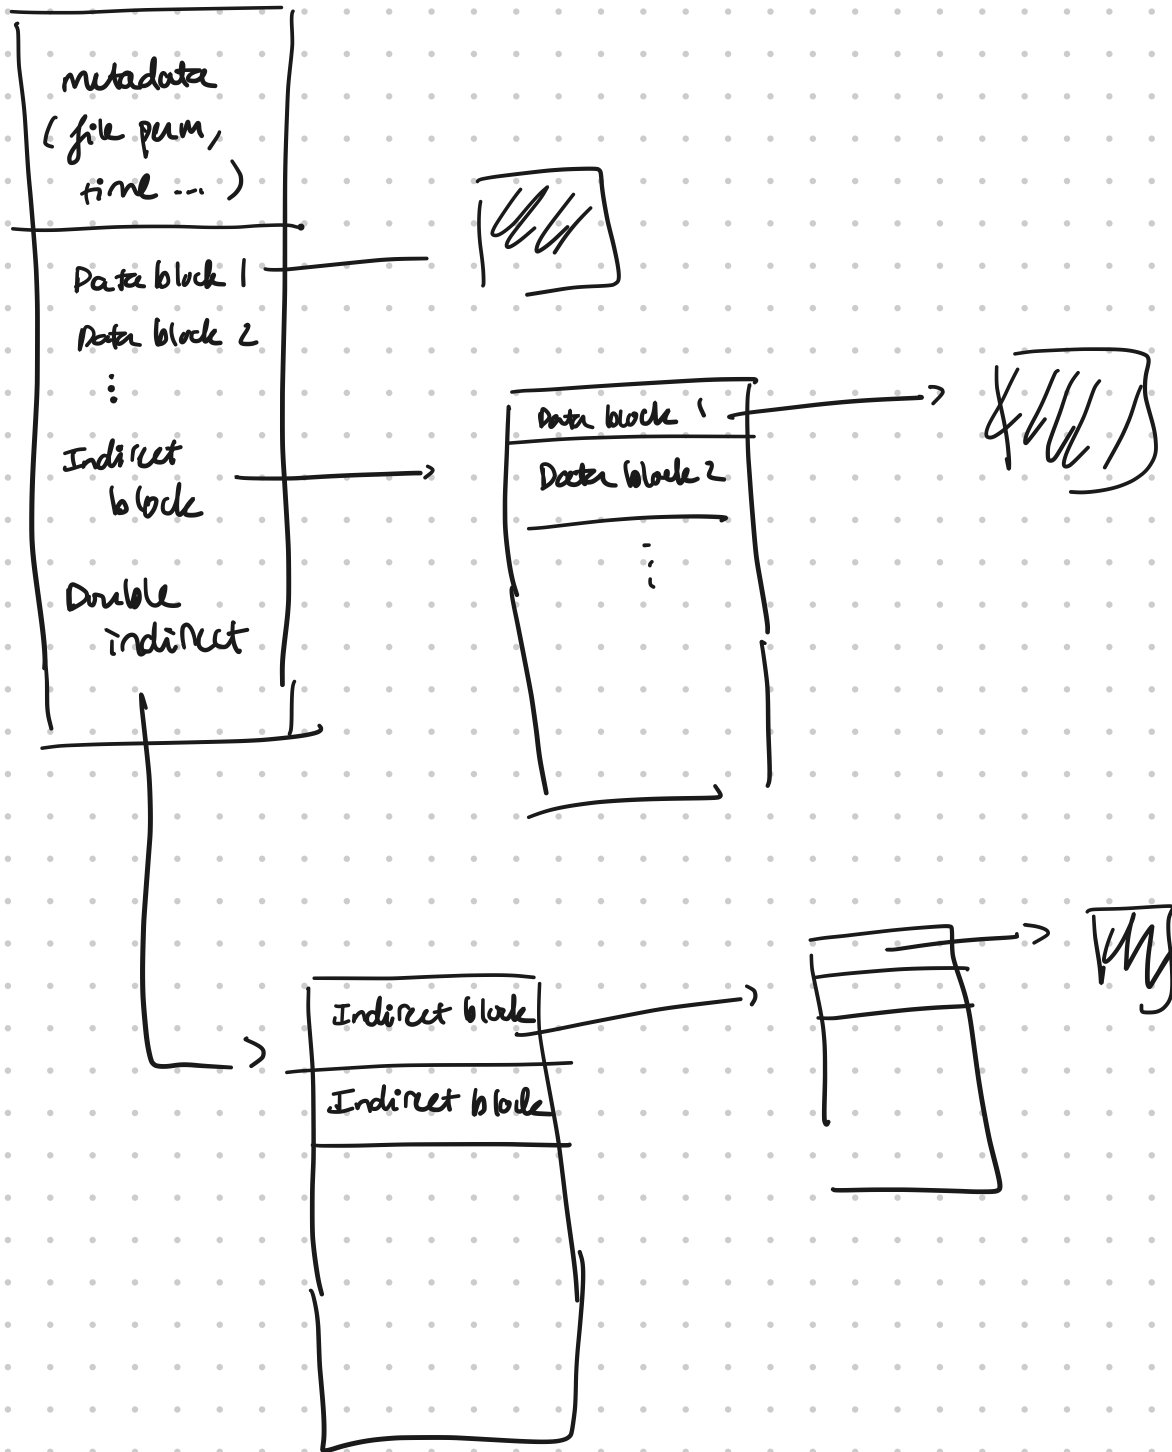
File: Data of 1 file might not be next to each other on the hardware.

File name: give the name to the data.
Don't have remember exact locations.

Directories: container of files to help organize.

Q: How do we implement this in Unix?

Inode



Q: What is a similar data structure we have seen?

Page Table.

The sparse / imbalanced tree allows us to handle both small & large files. For small files, we can fit it into the direct pointers to the data block. Else, we can use an indirect block or double indirect block.

FS contains a fix-sized array of inodes. Each inode is indexed by a number, referred to as i-number.


Mental Model:

$\langle i\text{-number} \rangle \longrightarrow \text{inode.}$

How about directories?

- Directories is also implemented using inode. Its content in each entry is a mapping of name to inode number.



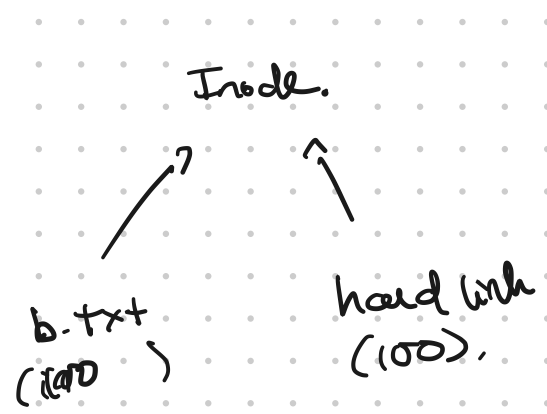
name \longrightarrow inode #.



(Symbolic)


Hard link vs Soft link:

- Hard link: have the same inode # as the file
 - cannot have hard link to directories or across file system.
- Softlink: allocated its own inode. Its content is the path to the file name.



b.txt -> 100

b.txt.link -> 102



2. Lab 5 overview.

FS in lab 5 is slightly different.

- 1 region in which both inode & data block reside. There are usually inode regions & data block regions
- Each inode is allocated its own disk block instead

of being packed together in a single disk block.

- A sector performs a read of 512B (Hardware)
FS reads in terms of block size: 8 sectors $2^3 \cdot 2^9 = 2^{12}$ (4KB)

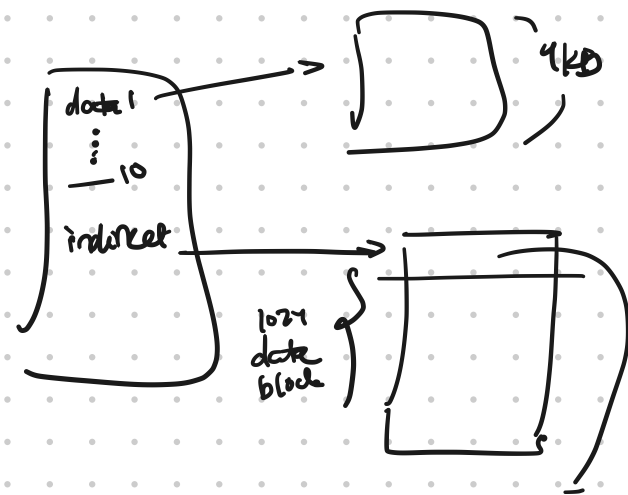
- Superblock: block 0, holding metadata about the FS & pointer to root dir.

- Bit map: an array of bits

001...0111...0

1: free to use

0: already allocated.



- Each inode contains

- 10 data blocks. Each 4KB
- 1 indirect block. 1024 direct blocks
- 1 double indirect blocks: 1024 indirect blocks.

Pointers to Pointers mental model:

int a = 10;


int * b = &a;

int ** c = &b;



*b = 1;

**c = 10;



Q, which table are where the file block numbered

• 5? Data block #5.

• 1032? Indirect block, 1022 entry in the table/mapping/array.

• 2060? Dindirect,
2nd indirect table/mapping, ...
2nd entry.


$$\begin{array}{r} 2060 \\ - 10 \\ \hline \end{array}$$

2050

$$\begin{array}{r} \text{disk address} - 1024 \\ \hline \end{array}$$

1026

$$\begin{array}{r} 1026 \\ - 1024 \quad (1 \text{ indirect in dindirect}) \\ \hline 2 \end{array}$$

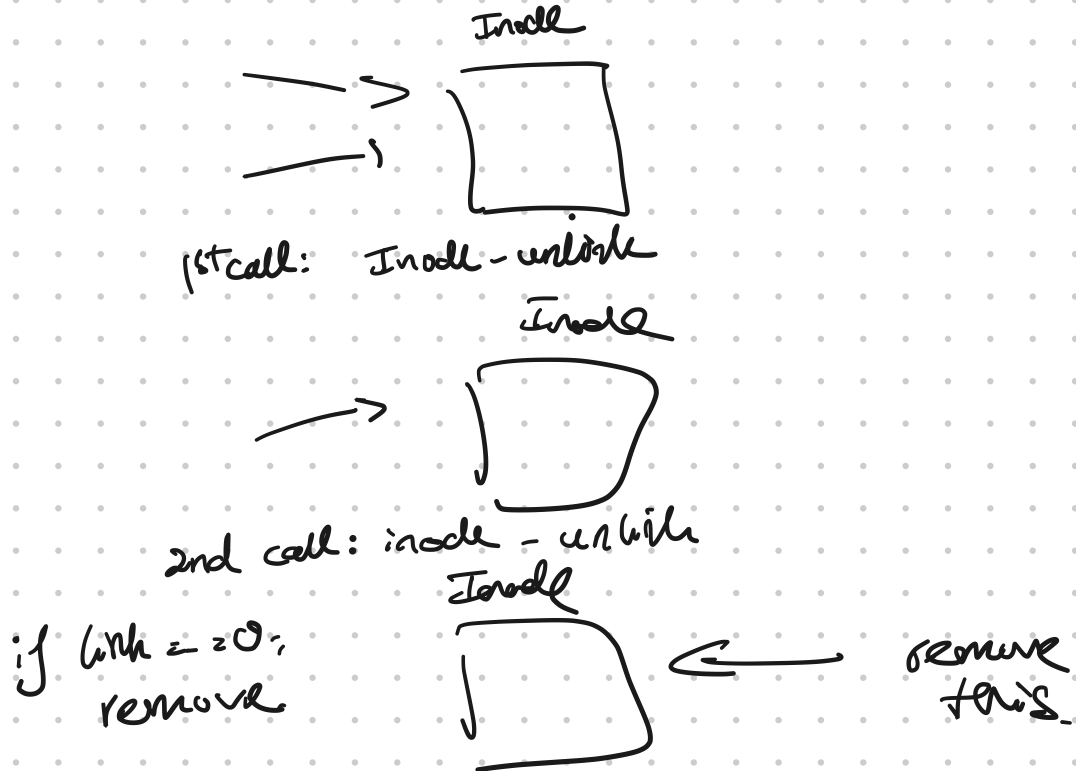


inode - block walk: find the fileblock & allocate

inode - get - block:

Set to start of the block

inode - unlink : decrement the link count.



Final Exam Study Tip:

- Double check your lab. Correct the mistakes.
- Review all the materials.
- Keep up w/ the readings:
 - Cross reference between Mike's & OSTEP.
- Start study early. Go to OHP, ask questions.
- Practice. Do the prev exams in a set of time.
 - Set a timer. Do all probs → Check the answers.
- Concurrency. I have no tips. Practice . . .
- Try not to overthink.