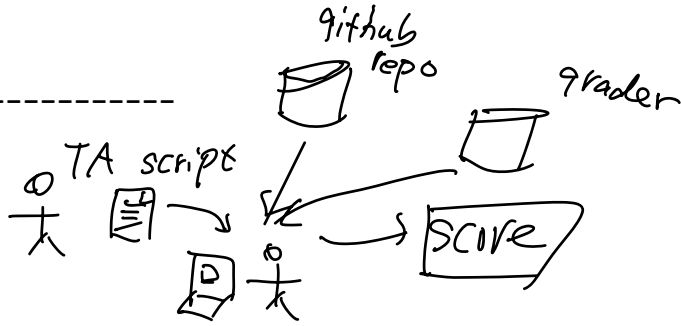


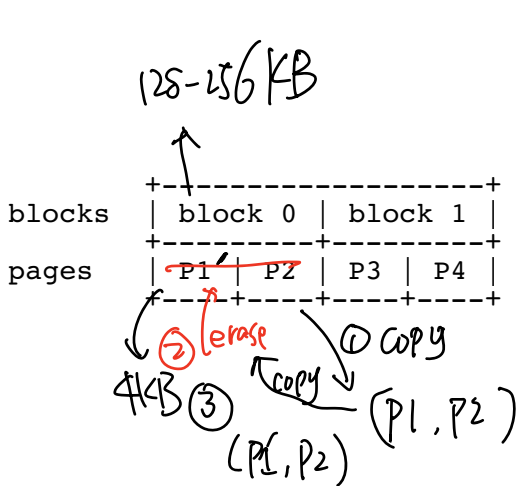
1. SSD continued
2. Intro to fs
3. Files
4. Directories

- Lab 3.
- Lab 5
- SSD.



- read: 1 page, $\sim 10s \mu s$
- erase: 1 block, $\sim 1s ms$ $\left. \begin{matrix} \nearrow 10x \\ \searrow 10x \end{matrix} \right\} 10x$
- program: 1 page, $\sim 100s \mu s$

• Wear-out:



flash bank

Q: a program updates a page free?

- FTL: flash translation layer

logical page $\xrightarrow{\text{FTL}}$ physical page

- log-structured FTL


* an example:

--Given a flash bank has two blocks; each has two pages.

--there are four writes to pages:

wirte(logic_page_1) ← LP1, LP10, LP99
wirte(logic_page_10)
wirte(logic_page_99)

--what will happen:



mapping:

LP1 => ~~P1~~, LP10 => ~~P2~~, LP99 => ~~P3~~
P4, P500, P999, LP1'

Question:

what will happen if the following op is write(logic_page_1')?
update

GC: garbage collection

- file system,
 - filetype
 - create/delete files
- Where are FSes implemented?
- persistency
- "name" a seq of bytes (file)
- human-understandable way to name (dirs)

• files

Q: $\begin{cases} \rightarrow \text{user: a seq of bytes.} \\ \rightarrow \text{FS: a set disk blocks} \end{cases}$

$\langle \text{file, offset} \rangle \longrightarrow \text{disk block addr}$

VMem, per-process:

$VA \xrightarrow{PT} PA$

Per-file:

$\text{offset} \xrightarrow{XYZ} \text{block addr (LBA)}$
 \swarrow inode \leftarrow imbalanced tree

How to find inode?

$\text{path/file name} \xrightarrow{\text{dirs}} \text{inode}$

Q: designing file-mapping.

- design parameters:



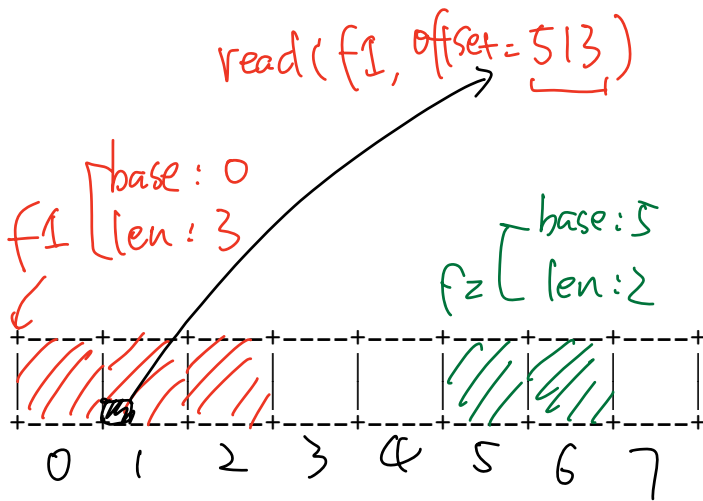
- * small files (most files are small)
vs.
large files (much of the disk is allocated to large files)
- * access patterns:
sequential access vs.
random accesses vs.
keyed accesses
- * disk utilization (metadata overhead and fragmentation)

- 3 Candidates

(A) Continuous allocation (extent-based file)

(B) linked files

(C) indexed files.

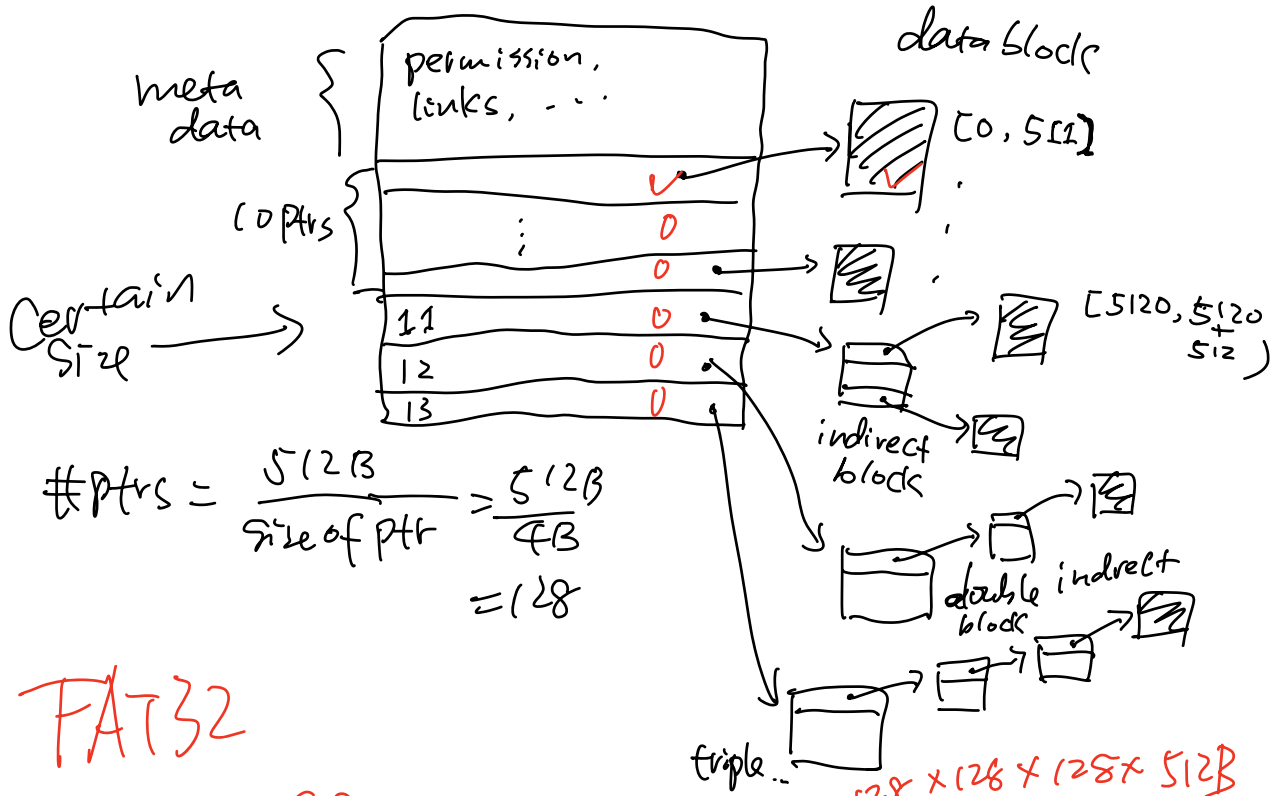


Q: How large is the table?

→ LARGE NUMBER

$$\#entries = \frac{\text{Size of a file}}{\text{block size}}$$

. Unix inode




$$\#ptrs = \frac{512B}{\text{size of ptr}} = \frac{512B}{4B} = 128$$

FAT32

↳ 4GB

$$128 \times 128 \times 128 \times 512B = 1GB$$

disk



df -i ~