4/24 welcome to CS 537!

should we hold class outside?

☐ Yes

☒ No

Do you believe that this is

a RULE

QUESTION ?

☒ No

Today

Wrong Thing:

Enjoy Life

Later

Today: File Systems (Implementation)

-> Locality } (2)

-> [Crash Consistency] } (1)

Review:

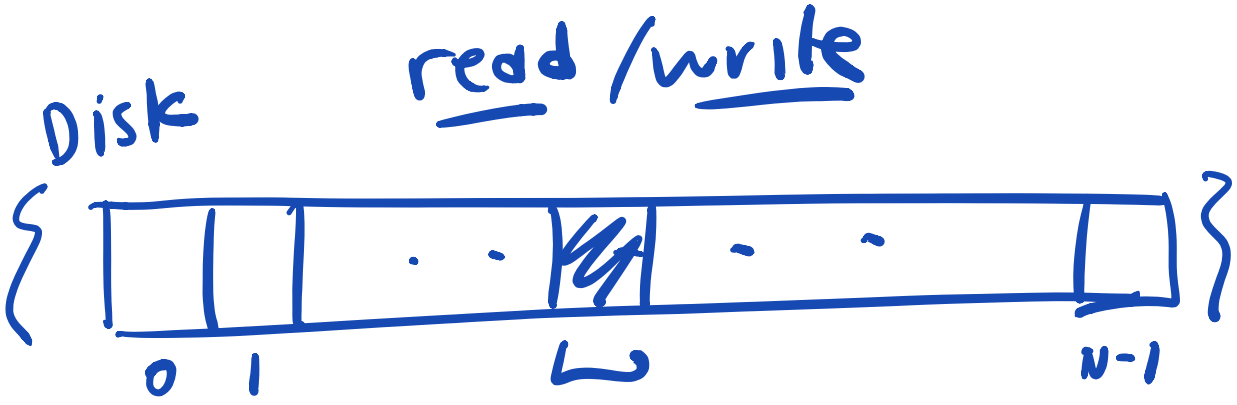

File System: uses API (open, close, read etc)

main memory / cache / write buffer

ops on disk
=> low-level reads/writes

FS:

=> On-disk structures

=> Access Methods

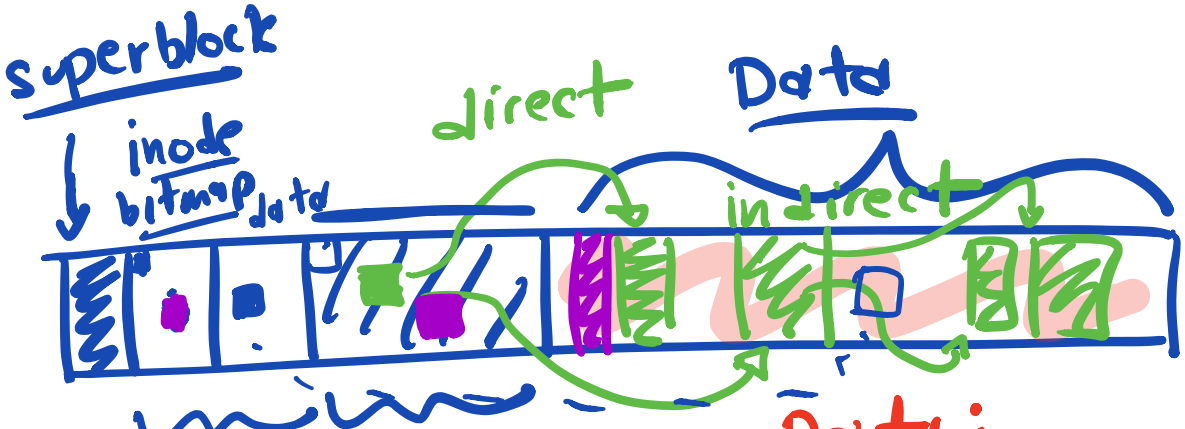

alloc. structures

per-file metadata
=> inode table

Data:
-> user data
-> dir data
-> indirect blocks

create: "/foo" (empty file)

Do

Sett on Thurs:

not erase

3:30 Drinks on him!

Crash: what is a "crash"?

- ⇒ power loss ← UPS
- ⇒ BSOD: kernel panic / bug (0 ⇏ 0)
- ⇒ user restarts

why important for FS?

⇒ in the middle of an update

Examples: updates

⇒ File Creation:

{what data structures are updated?}

/ ← root ⇒ creat("/foo");

e.g. ⇒ inode bitmap (find free, mark it used)

10 ⇒ inode: all its various contents

⇒ root directory: data

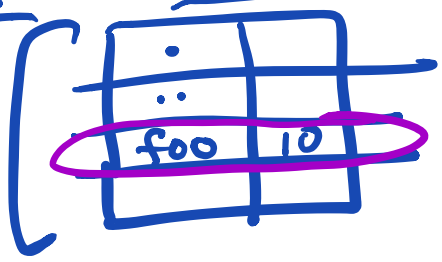

name ⇒ low-level name (inode #)

⇒ root inode: update times

[could also (if dir grows)
⇒ data bitmap,

new data block

existing file, goal => append new data block

[=> how API? (system calls)
=> which blocks written?]

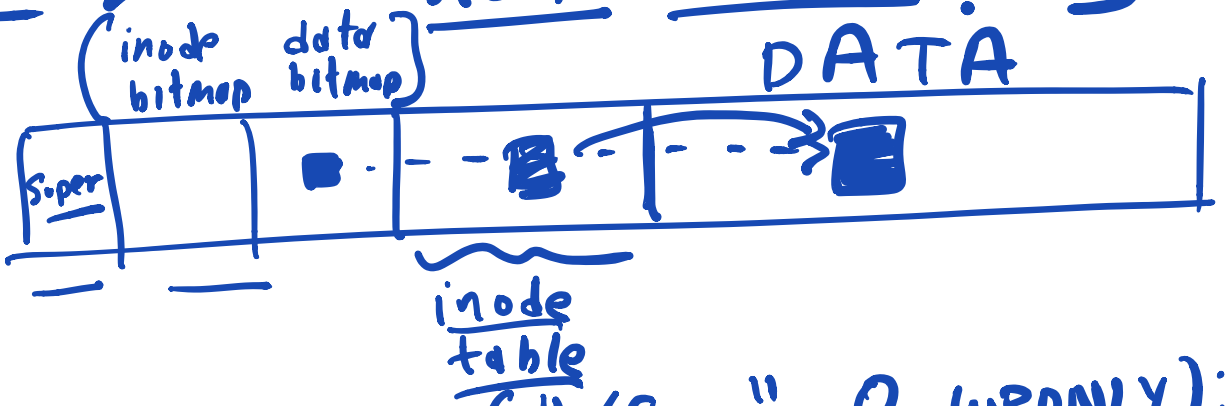


```
int fd = open ("/foo", O_WRONLY);
lseek (fd, 0, SEEK_END);
[ write (fd, buffer, size); ]
close (fd);
```

size: 4KB

Data/Metadata:

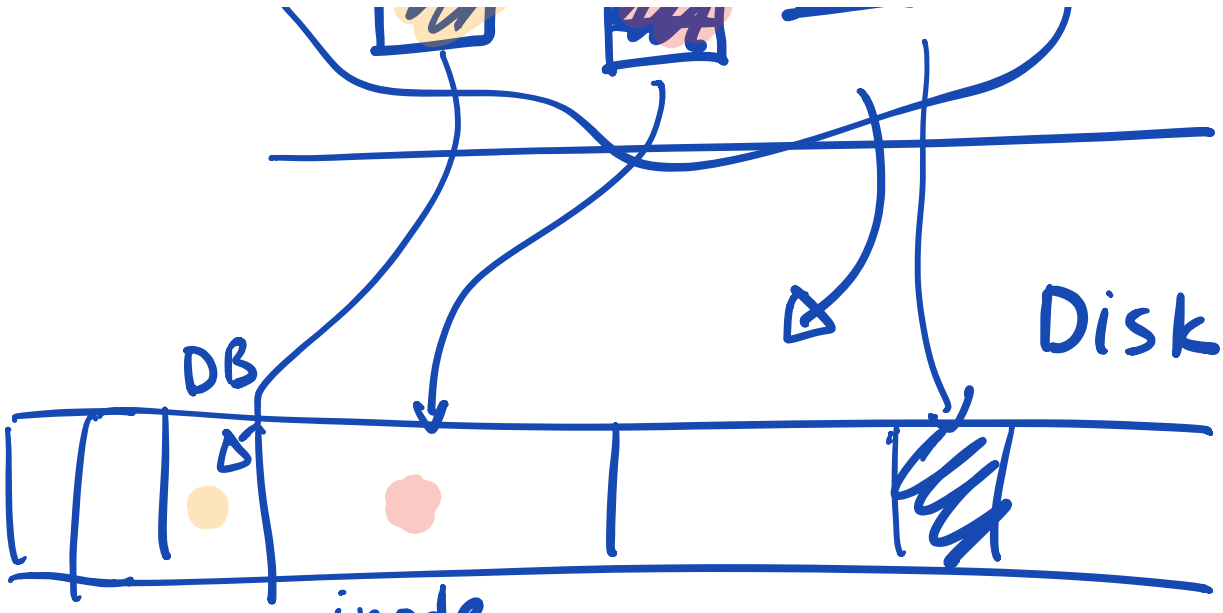

(D) (I) (DB)

Data, Inode, Data Bitmap:

orders of writing:

D	D	I	I	DB	DB
I	DB	D	DB	I	D
DB	I	DB	D	D	I

crash here or here

[what happens?]

address: 100

DB inode

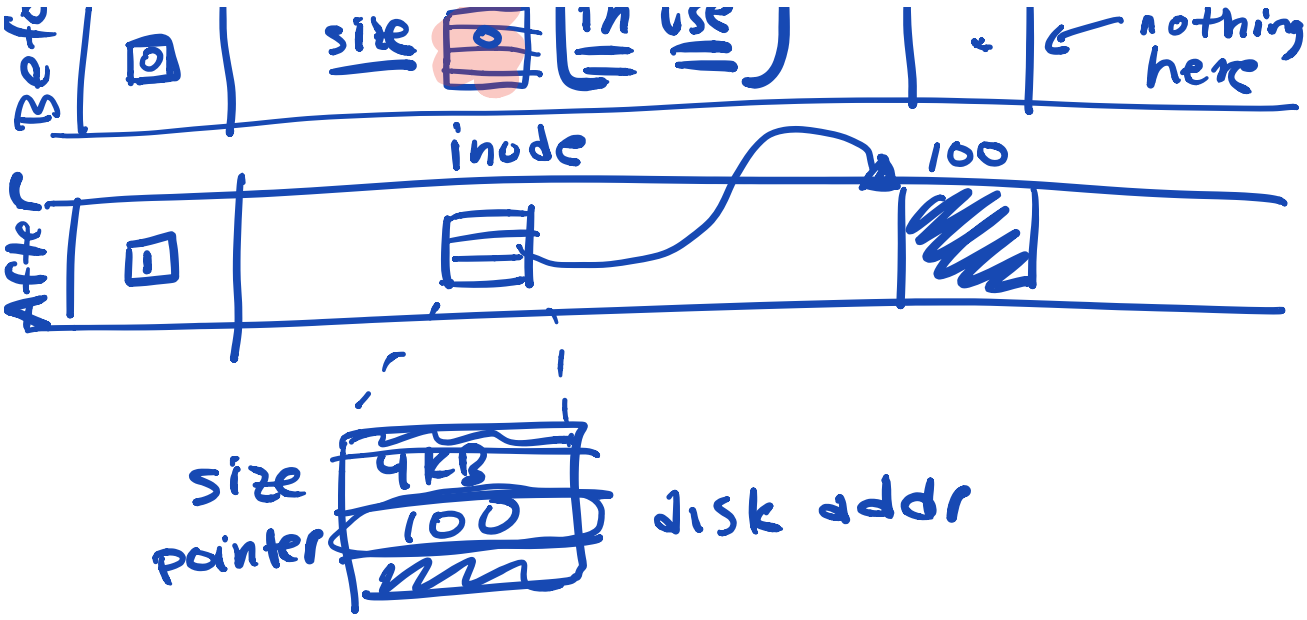

(D) | (I) | (DB)

Data, Inode, Data Bitmap:

orders of writing:

D ①	D ③	I ⑤	I	DB	DB
I ②	DB ④	D ⑥	DB	I	D
DB	I	DB	D	D	I

①③ looks good! file system metadata "consistent"

② inode ↔ data bitmap don't agree! ⇒ inconsistent

④ inconsistent : spare leak

⑤ inconsistent :

could read garbage (or zeroes) (or privacy problem)

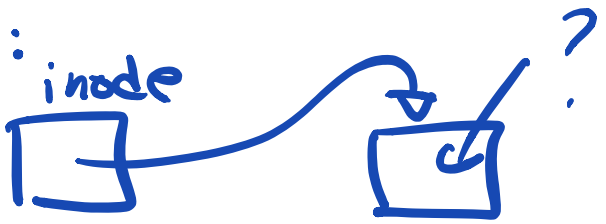

Solutions ?

Check : is the FS consistent ?

⇒ before usage (before mount)

⇒ file system check (fsck)

⇒ [focus of PS (good news)]

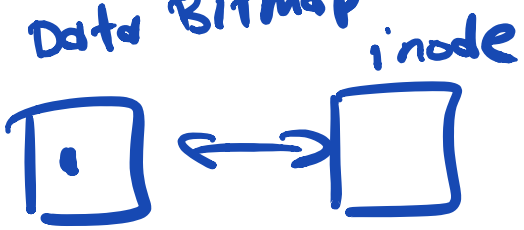

But: Disks got large (RAID too)

=) too slow?

Something else: Journaling (or write-ahead logging)

=) Eager:

→ some work w/ every update

→ recovery (after crash):

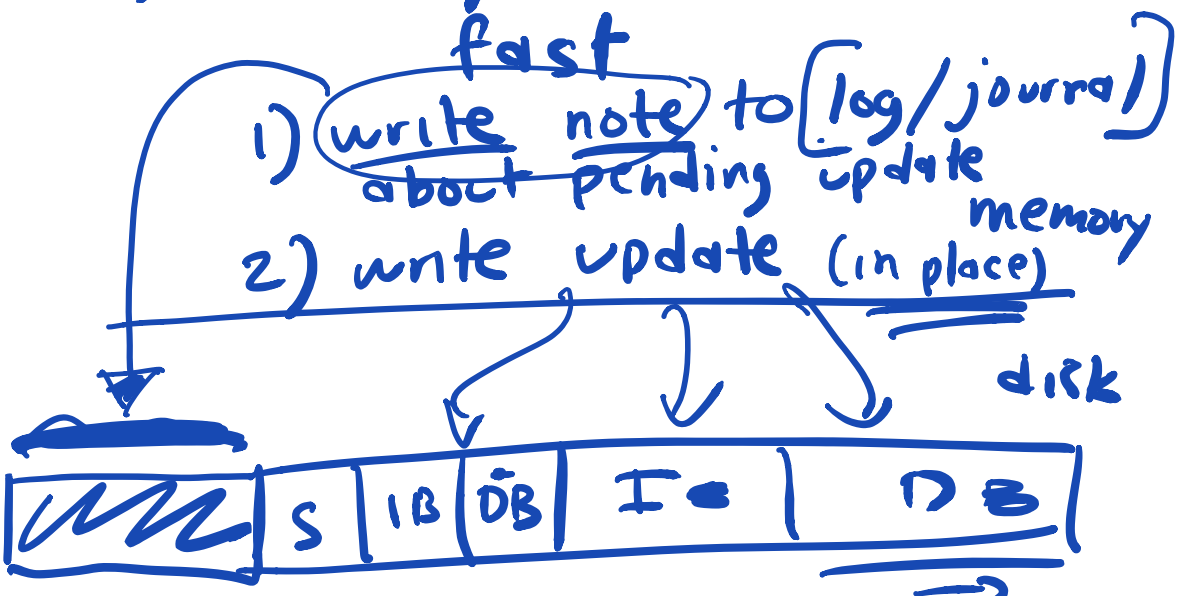

Journal (or Log)

Details: example -> append to existing file

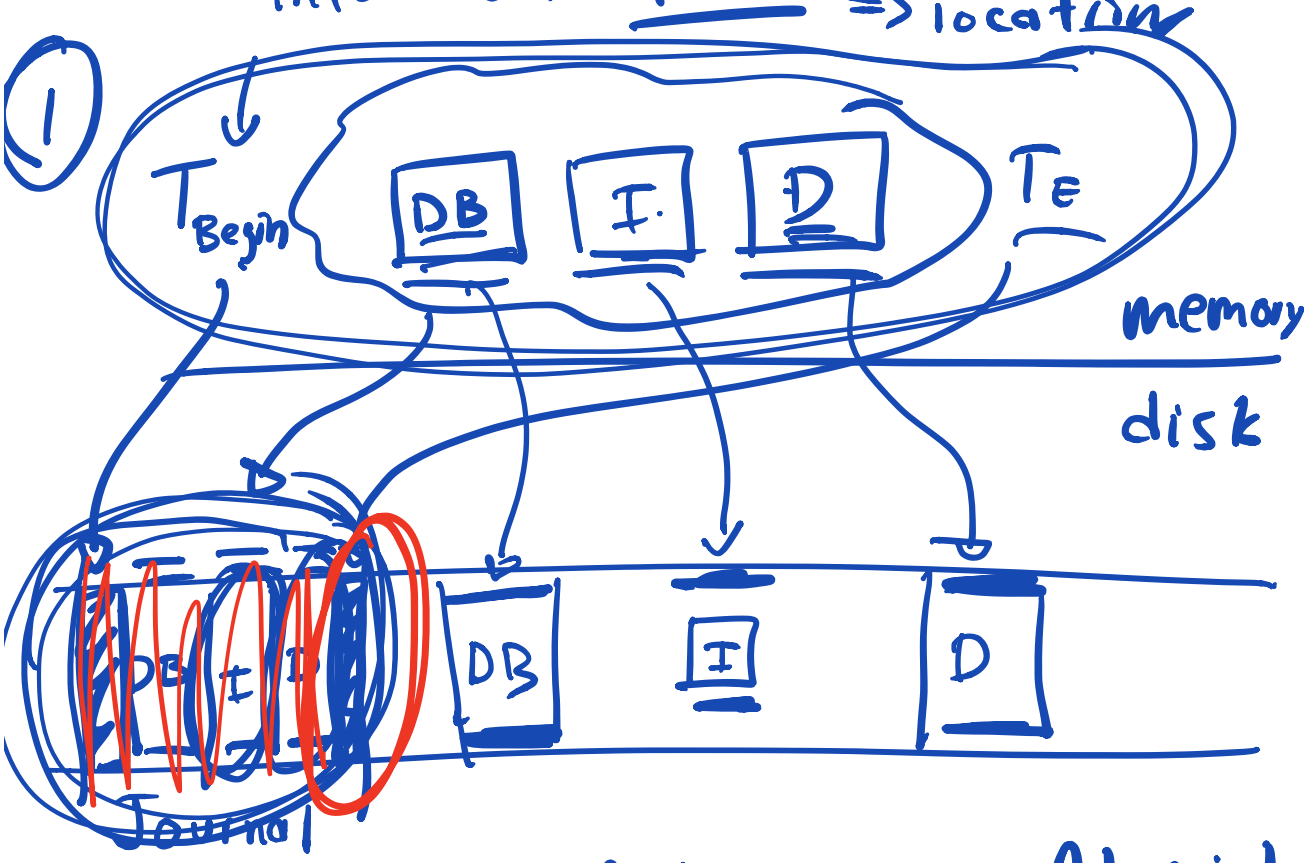

All or none: Atomicity

2 writes "in place"

=> PB, I, D

Crash:

during step 2 (after step 1)
recover from journal

during step 1:
or (issue all writes)
to journal

Aside: disk guarantees
that any 512 byte write
is atomic
but larger write may
partially complete

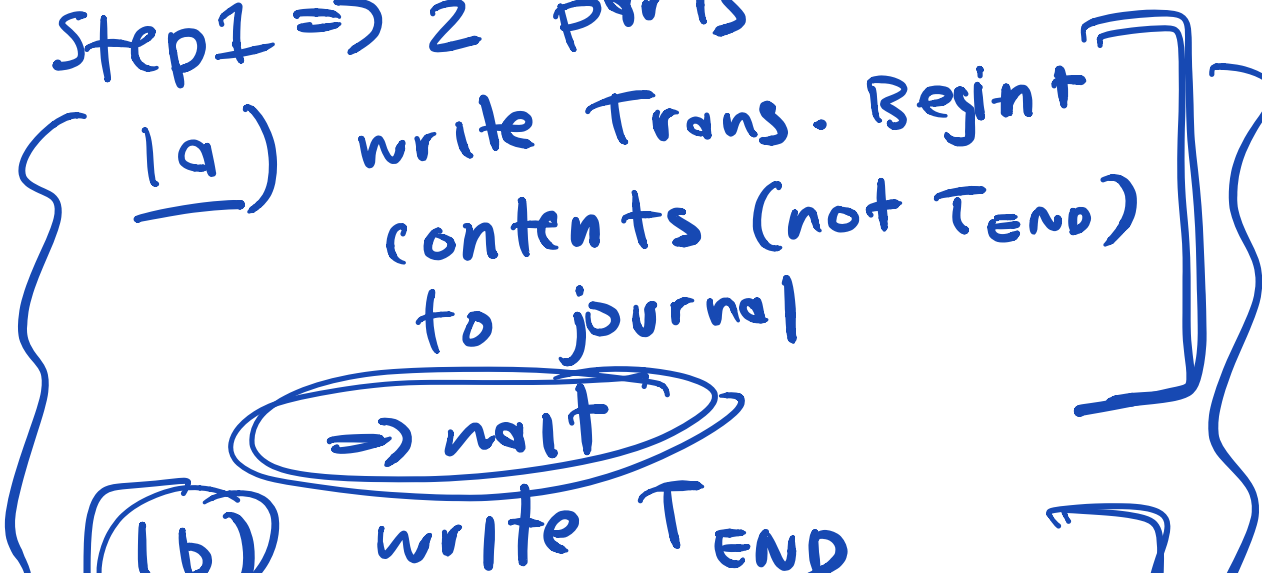

=> transaction commit

wait

2) in place updates

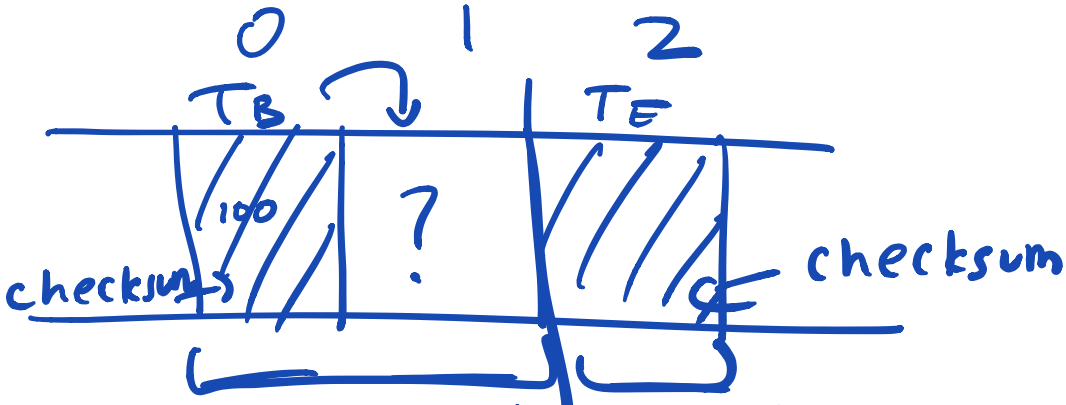

2005: transaction checksum => {Linux ext4}

Crash Consistency:

=> [FS] $\in$ initial deadline

xv6 file system image

C/Linux

look for inconsistencies

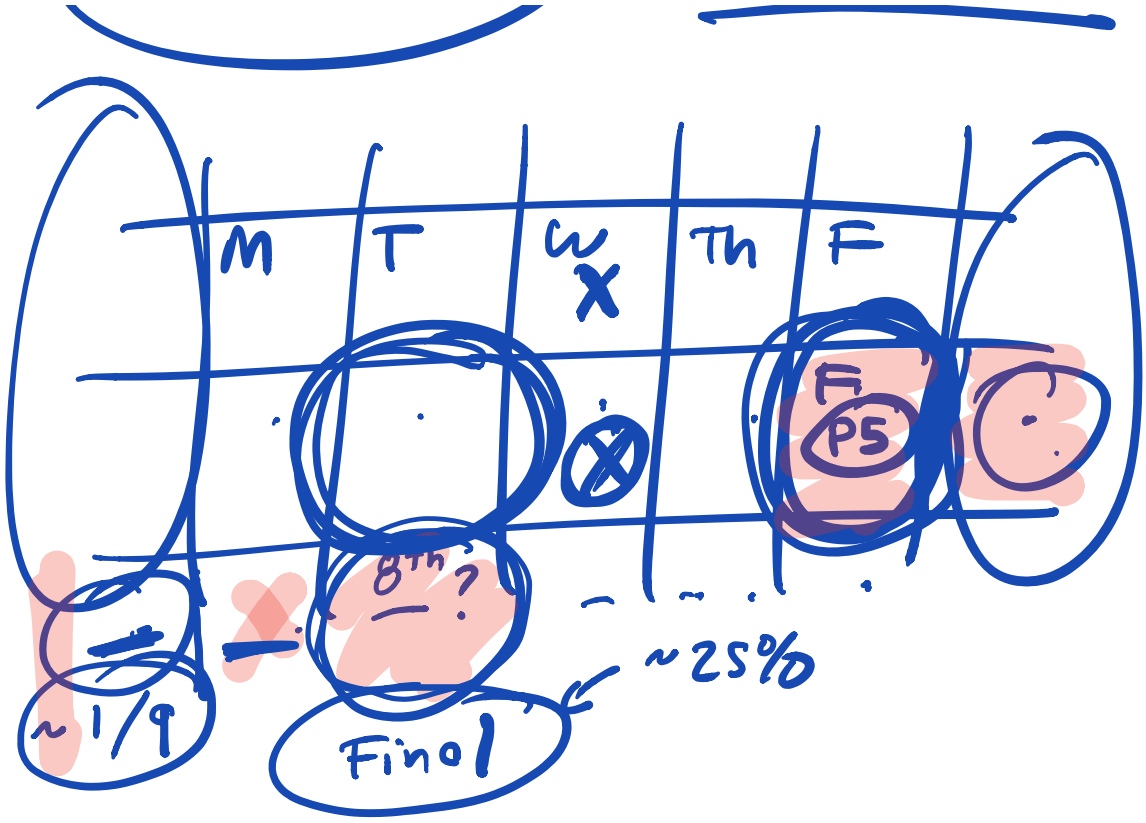
M
T
W
Th
F
F
P5
8th ?
~25%
~1/9
Final

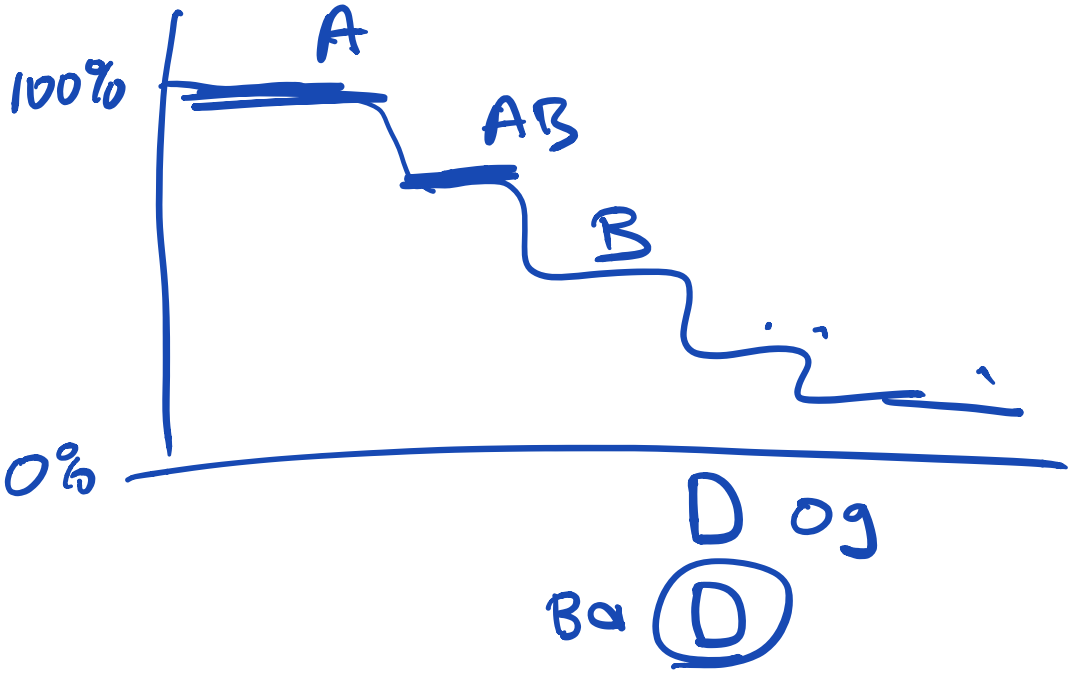
A
100%
AB
B
0%
D
Og
Ba
D

Locality

Hard Disks: long seeks costly

SSDs: locality important too

goal:
put files/etc. on disk => accessed together should be near one another

Fast File System (FFS)

=> Performance

Problem: old unix was slow

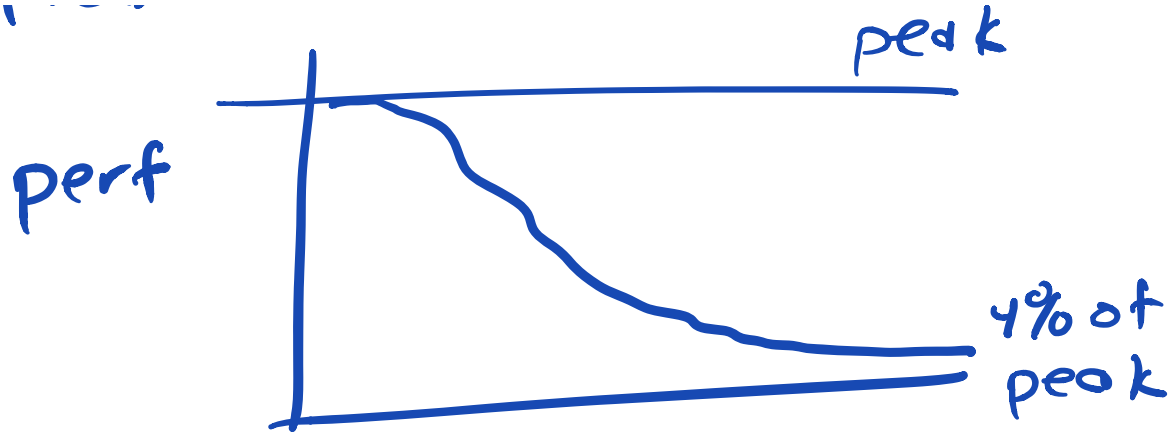

=> data blocks, inodes scattered through disk

like:

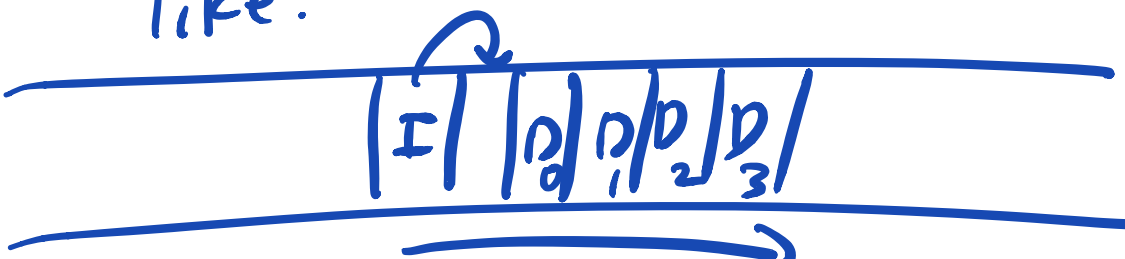

actually:

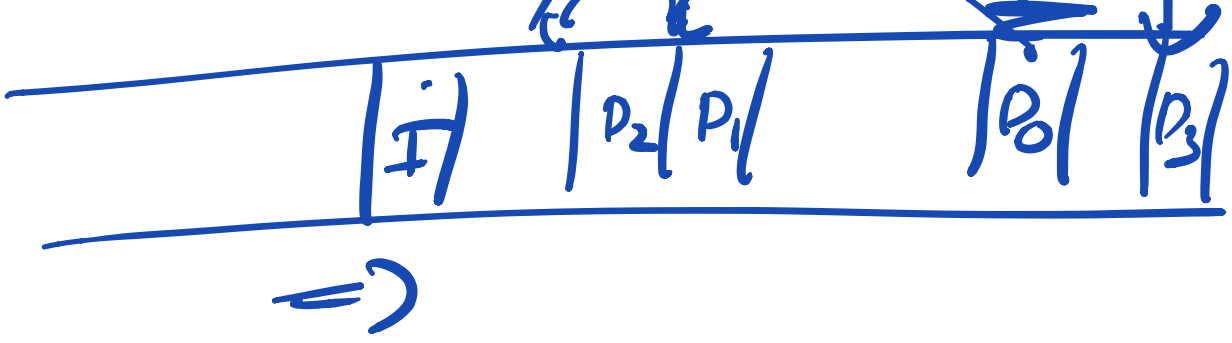

=>

FFS: Treat Disk like a Disk!

On-disk structure: change

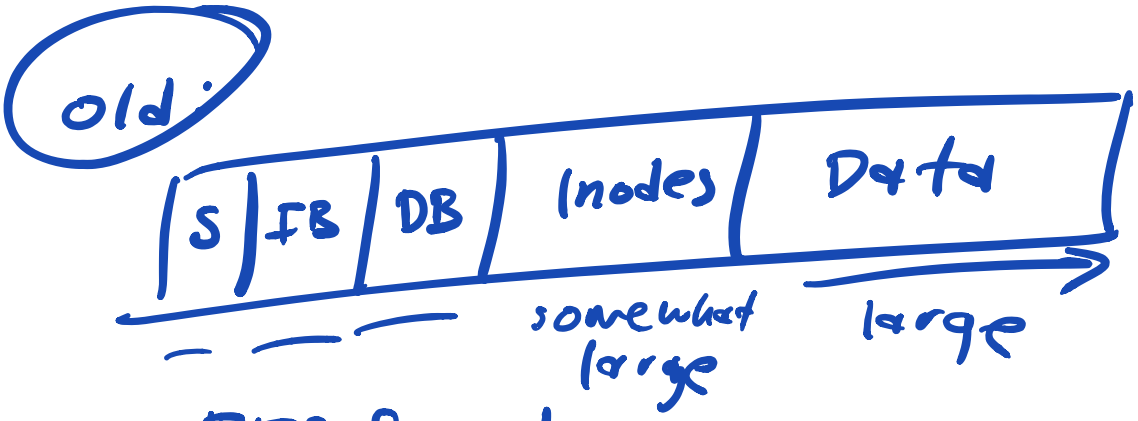

new FFS format:

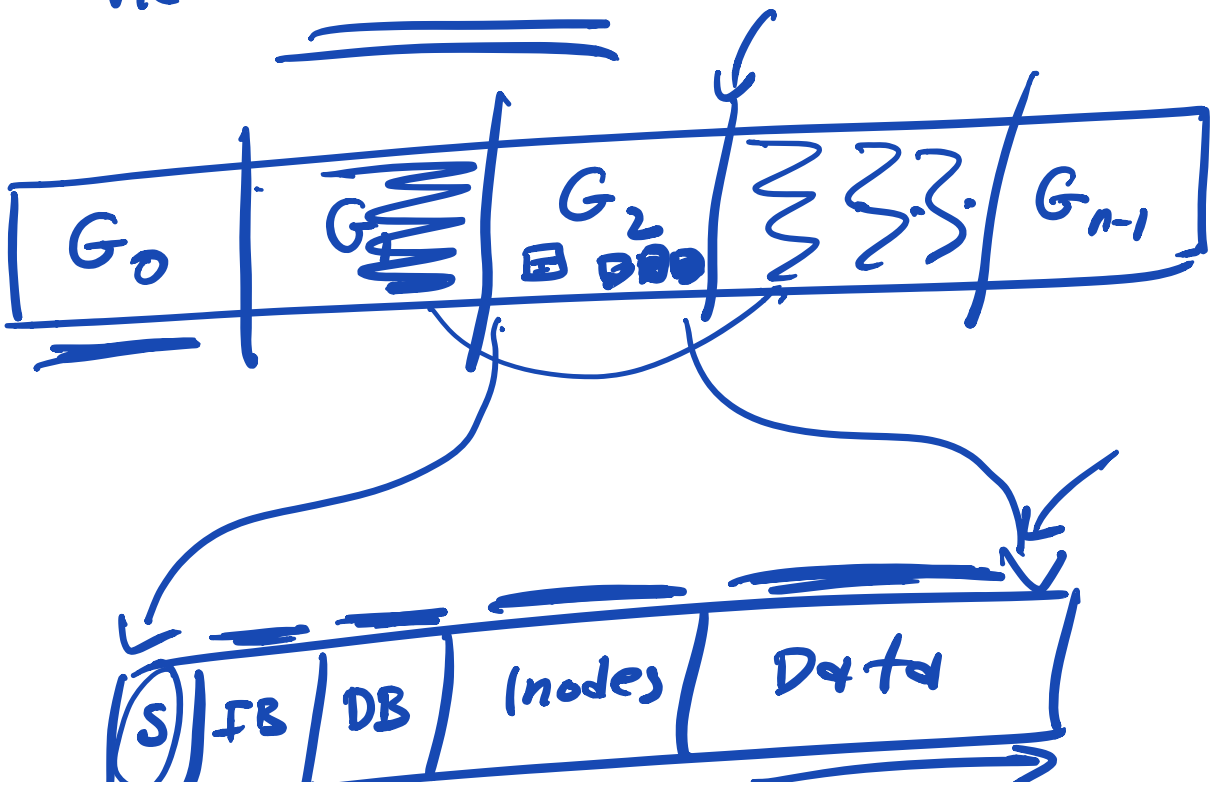

somewhat large

Goal: Put related stuff into same group (unrelated → different)

what is "related"?

- → data same file
- → files in same dir.
- → inodes/data

make new dir:

pick some group X (how?) → (free space, inodes)

all files in dir (inodes/data) ⇒ group X

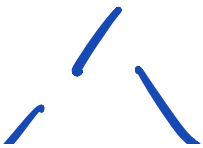

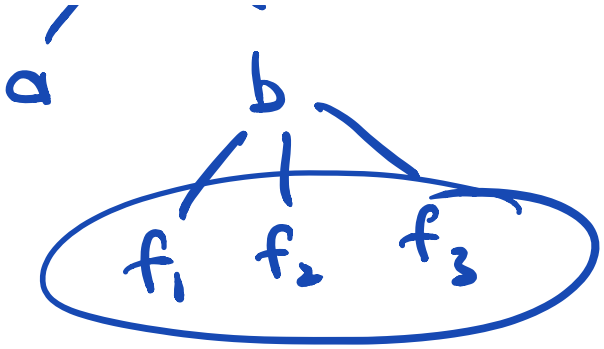

Large Files: exception

only put first N blocks of file => desired group

put next chunk of file in each other groups

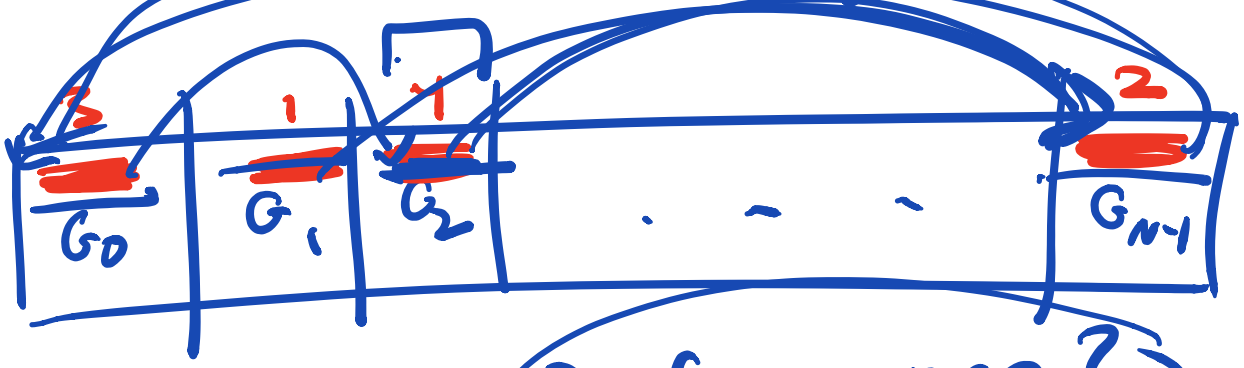

Problem: Performance?

[how big is each chunk?]

=> need to know:
  => how big file is ✓
  => group size ✓
  => memory ✓

=> cost of seek
=> transfer speed

read:
read chunk
seek
tolerable
10%

most transfer:

transfer: 100 MB/s
seek (avg): 10 ms

10^{x} MB $\cdot \frac{\text{sec}}{100\ \text{MB}} =$

$\Rightarrow$ [Final] $\Leftarrow$ Study this

FFS : [first time: same API, new implementation]

$\Rightarrow$ symbolic link

$\Rightarrow$ long file names

$\Rightarrow$ other perf. opts (not as relevant)

Treat Disk like Disk

$\Rightarrow$ [peak, sustained

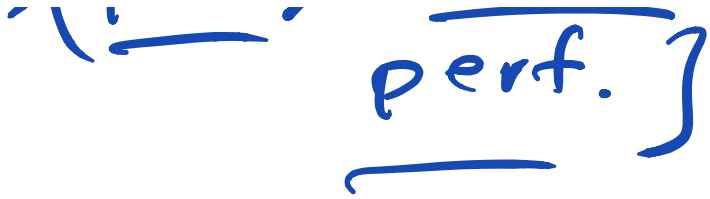
perf.