

CS202: File Systems 1

◦ LAST TIME:

Disks

Interface: Read or write sector

block ← In this lec., notes, etc.

Performance characteristics

— sequential access is much faster than random access

Why?

FILE SYSTEM

PRIMARY USE: READ & WRITE DATA STORED IN
A DEVICE THAT PROVIDES PERSISTENT STORAGE

(PERSISTENT \equiv SURVIVES PROCESS & COMPUTER
RESTART)

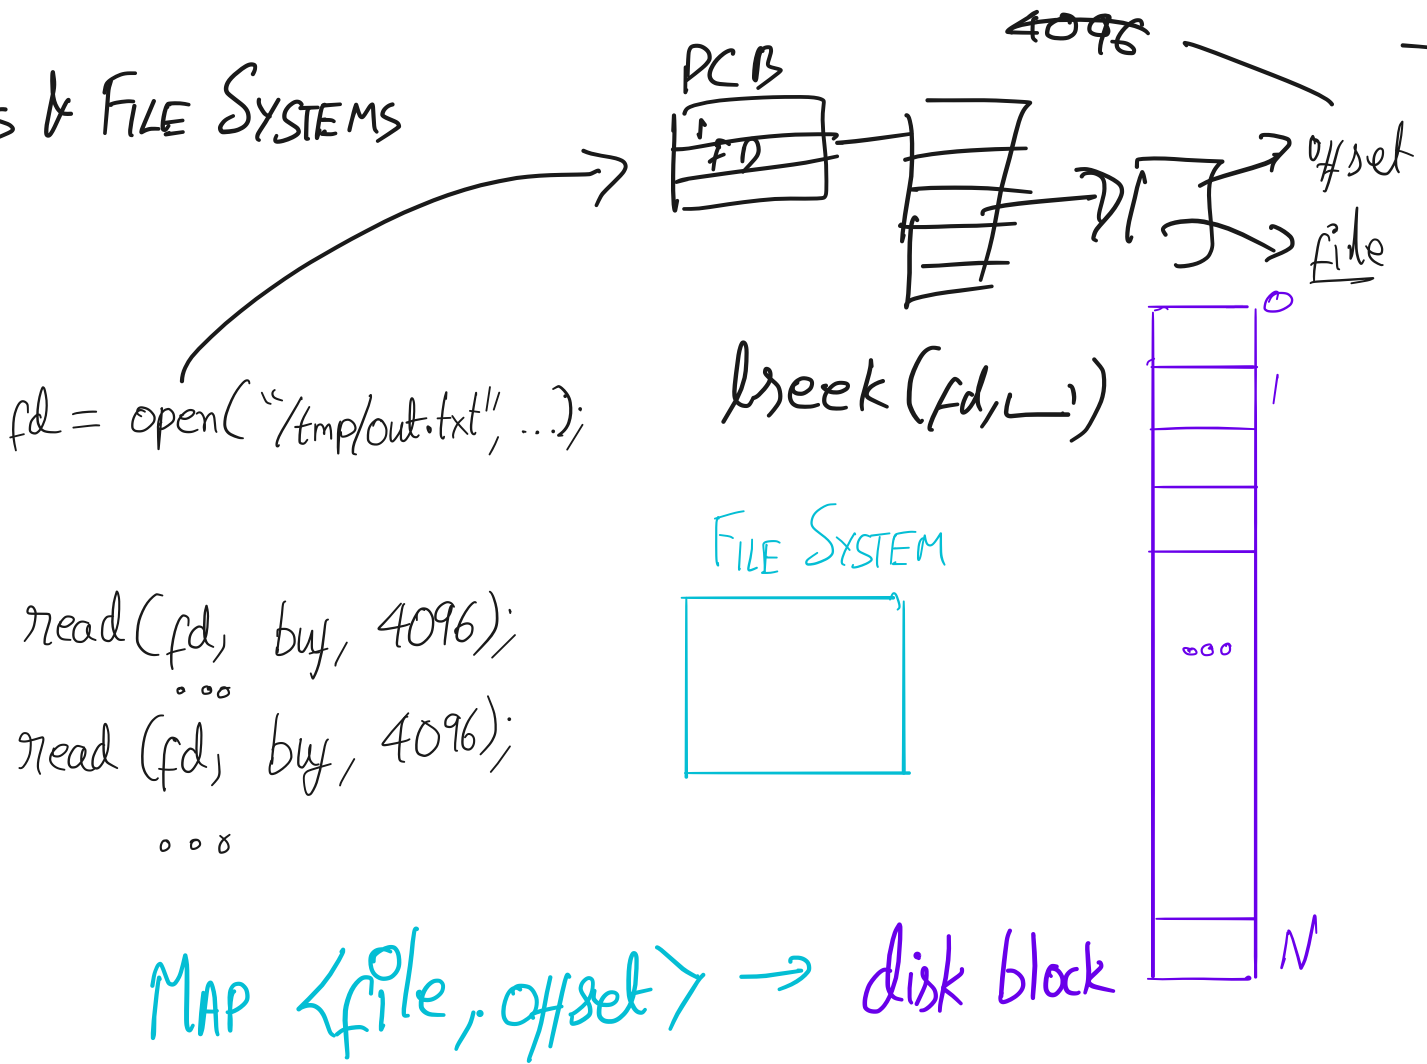
NEEDS TO PROVIDE A WAY TO

— MANAGE SPACE ON DEVICE

— NAME DATA STORED ON DEVICE.

FILE SYSTEMS HAVE TARGETTED MANY DEVICES,
WE WILL FOCUS ON DISKS.

FILES & FILE SYSTEMS



Interface

Open: Open a file

↳ O-CREAT: Create a new file

lseek: set offset

↳ Read from current offset

read: read from current

write: Write starting from current offset (might need to allocate blocks)

GOALS


- Minimize disk access time

↳ Sequential is better

- Minimize wasted space

EMPIRICAL OBSERVATIONS: Most files are small;
most of the disk is used by large files

Have both sequential & random access.



Q: How to Layout files

- Dictates how disk blocks are allocated
- Read and write performance.


(a) CONTIGUOUS ALLOCATION

Put all of a file's blocks next to each other

+ Great for seq & random access

- Adding a new block to a file is challenging


- FRAGMENTATION



(b) LINKED FILES


+ Easy to extend file

~~F~~ ^{on avg reasonable}




+ ~~fast~~ sequential access

- BAD RANDOM ACCESS PERFORMANCE

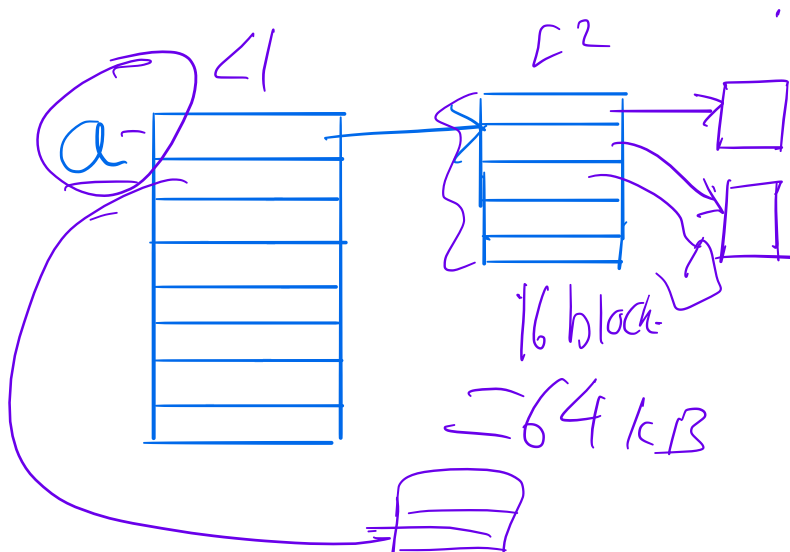


© INDEXED

↳ STORE THEM IN A TREE



Q. How many entries?



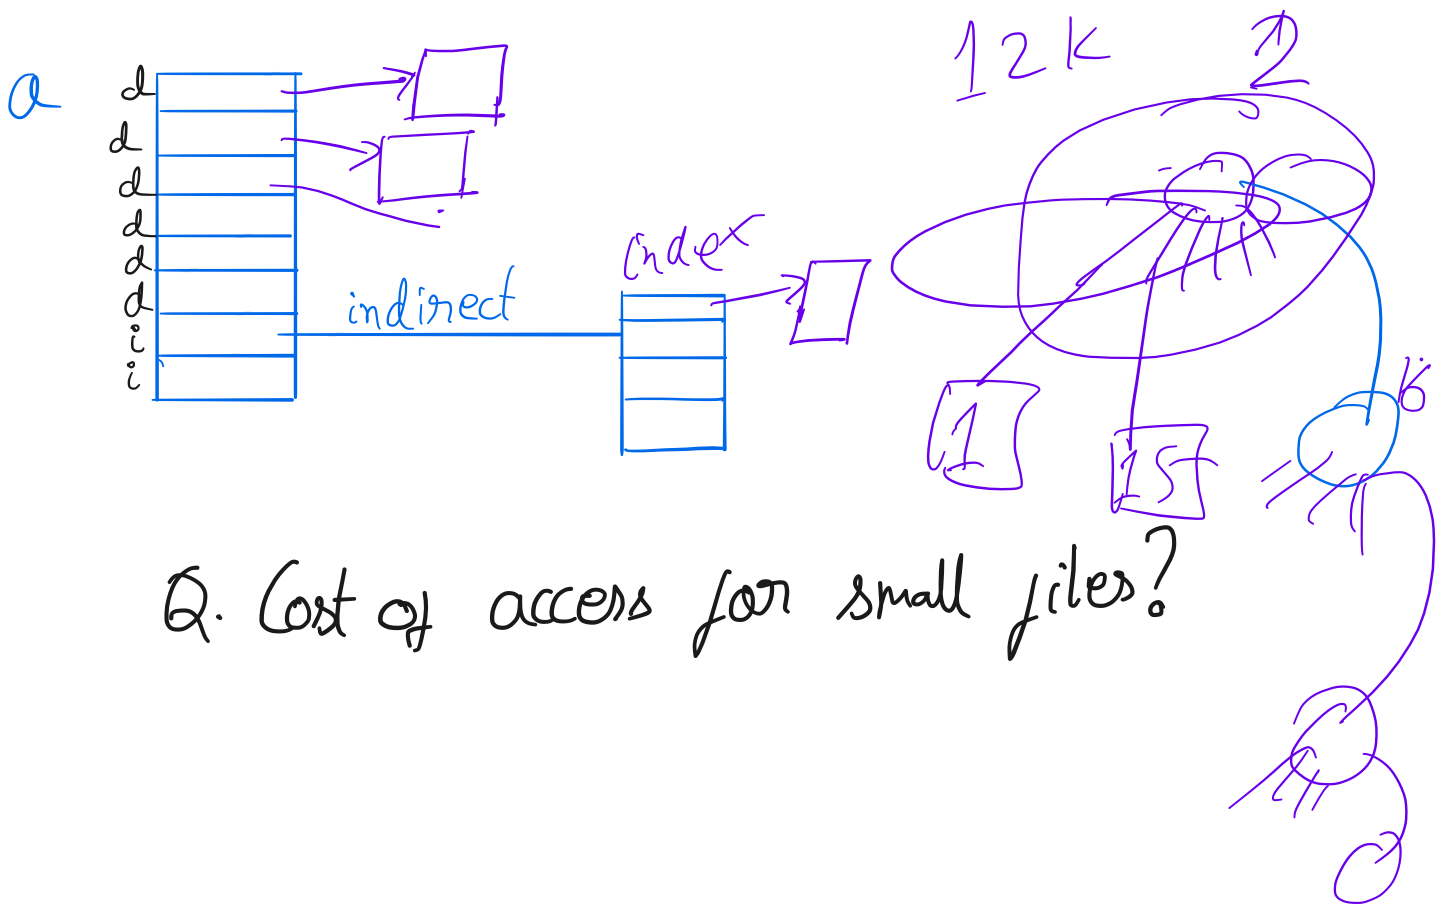
`fd = open(a)`

`seek(fd, 3 * 4 * 210)`

`read(fd, c, 1)`

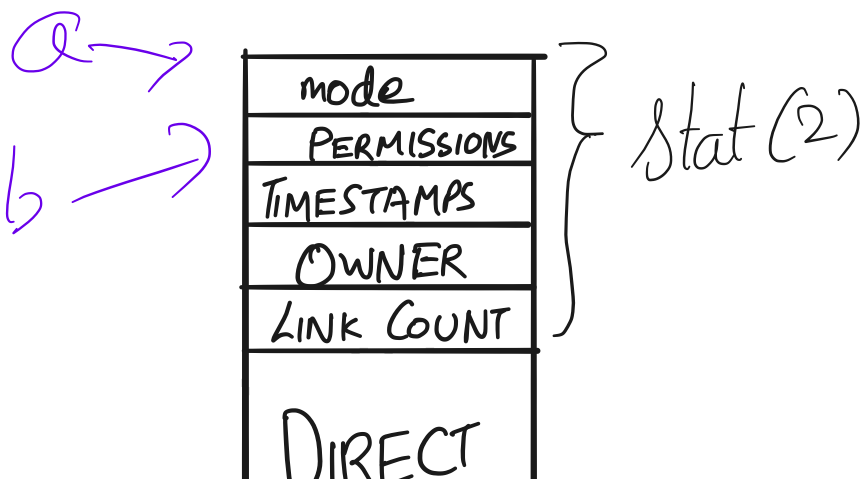
Q: Cost of reading 1 byte?

Q: Cost of reading file with 1 block?




Q. Cost of access for small files?


Index Nodes — inode



Jim a.



DIRECTORIES



70,000

