

6.S081: File systems

User-Friendly names / pathnames

Share files between users/processes

Persistence / durability

Why interesting?

two tabs

Abstraction is useful

Crash safety ← lib.

Disk layout

storage devices are slow

Performance

buffer cache concurrency

API example / file systems calls

fd = open("X/Y", 0);

write(fd, "abc", 3)

link("X/Y", "X/Z") & multiple names

unlink("X/Y")

write(fd, "def", 3)

fd

path name

human readable

offset is implicit

File system structures

Inode = file info, independent of name
inode #

→ link count } both are zero

open file count

fs must maintain an offset

FS layers names fct's

inode — read
write

icache → synchronization

logging

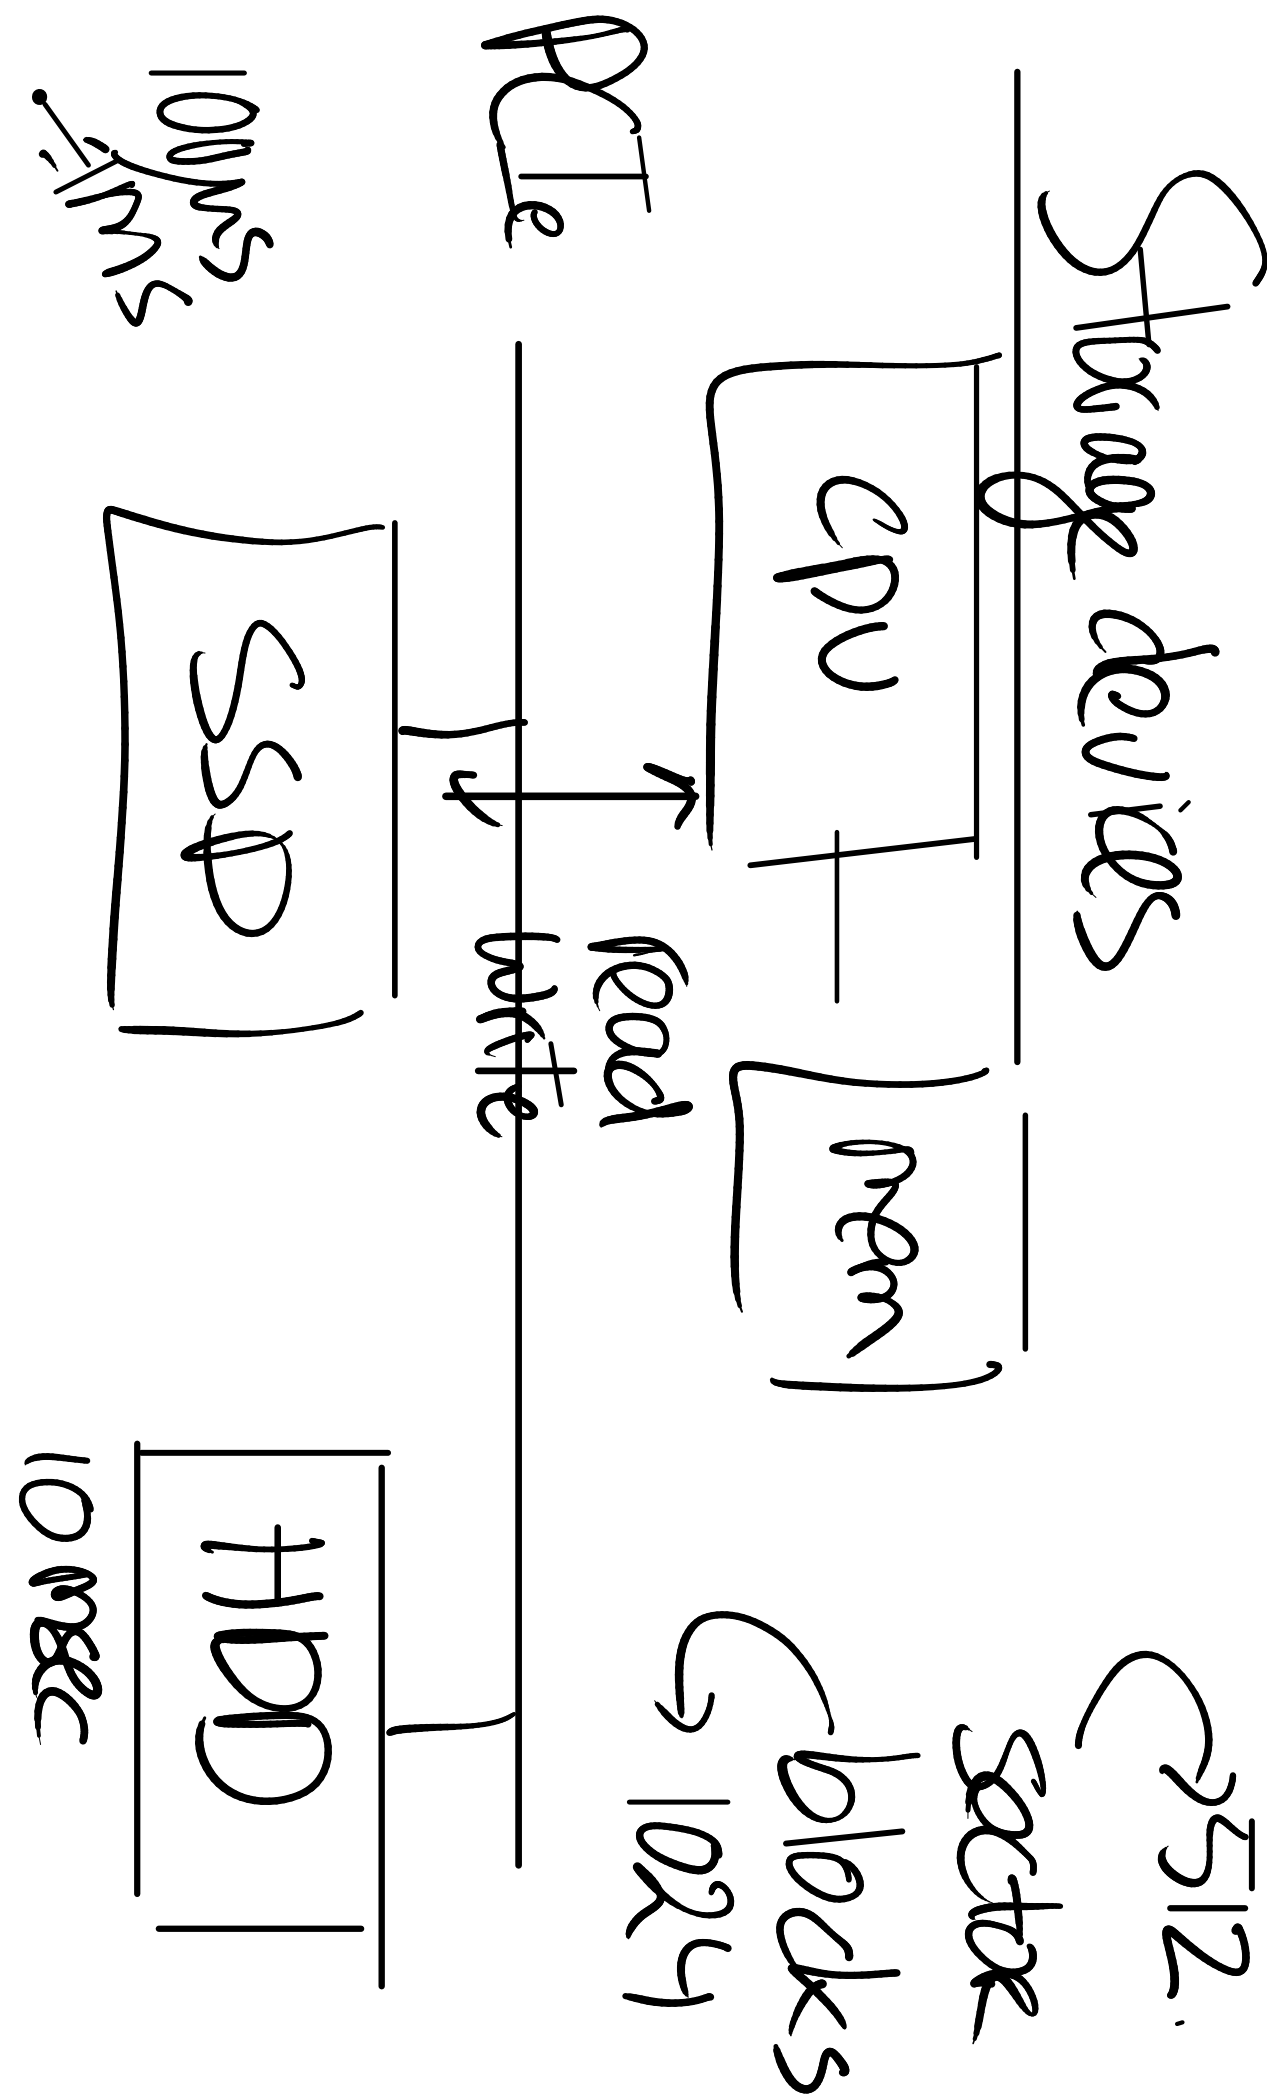
buf cache

disk

memory

cache


Storage devices



Disk layout

768 bytes

data



superblocks

metadata


kernel

inode 10 →

$$32 + \text{inode} \times 64$$

$$\frac{1024}{1024}$$


On-disk mode



max file size:

$$(256 + 12) \times 1024 \text{ bytes}$$

268 \Rightarrow read w/e



8000

8000/1024

8000 \approx 7 1024 \approx 82

Directories (XV6)

dir = file w. some structure

pathname
lookup

X/X

num# ₂₅₁	file name
---------------------	-----------

entry

2
k₁

→ root node (7)
scans block for name

→ read node 251
Scan its block for X

Brache block cave

One copy of block in men

Sheep locks

LRD

two levels of binding

Summary

$f_s =$ on-disk data structure
↳ XUB very simple

block cache

WEd: crash safety