






Disk Scheduling

SSTF


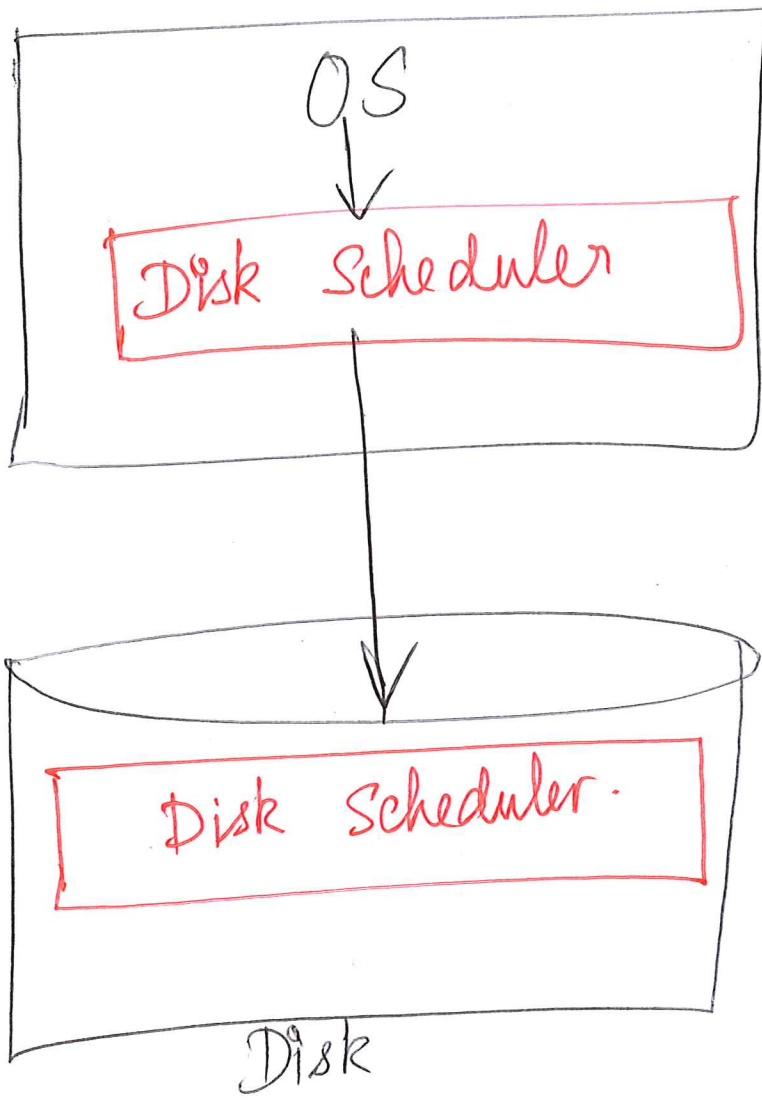


SCAN



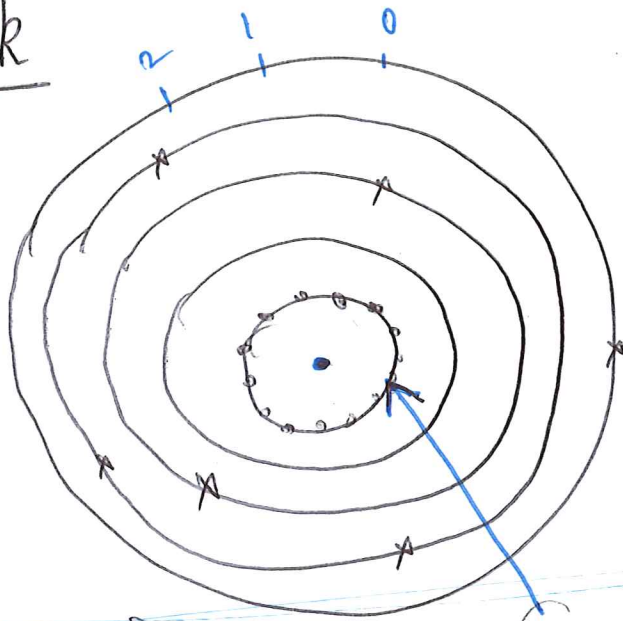
C-SCAN





I/O devices $\left\{ \begin{array}{l} \text{Persistence} \\ \text{interface} \\ \text{implementation.} \end{array} \right.$

Disk



$\text{Rate (Seq. Read)} \gg \text{Rate (Random Reads)}$

Disk Scheduling

1) FIFO.

2) SSTF
↓
Seek

starvation

ignores rotation

3) SPTF / SATF

4) BSATF (windows)


SCAN (Elevator)

1 2 3 1 2 3 1 2 3

C-SCAN

F-SCAN





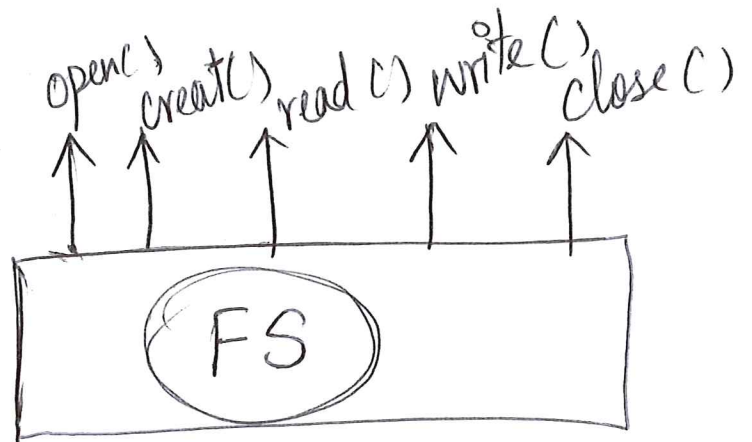
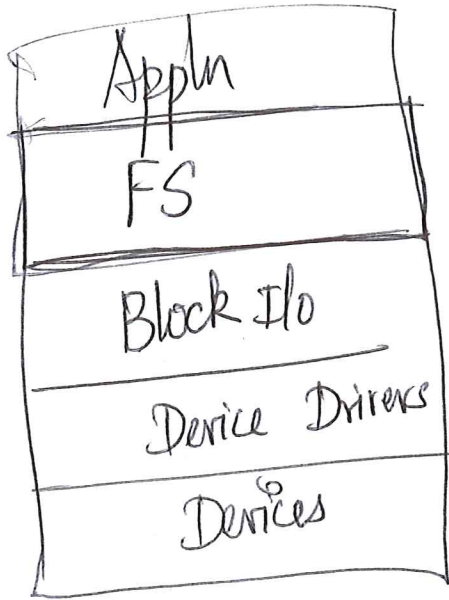
1. File System APIs.

2. File System Implementation

Data Structures

Access Methods

File System APIs



Basic Abstractions

Directory

- File (Regular File)**
- array of bytes
 - low-level name (inode #)
 - human-readable name

- special type of file
- human-readable name
- low-level name
- array of records



human name	low-level name
a.txt	1231271
dir1	1470712

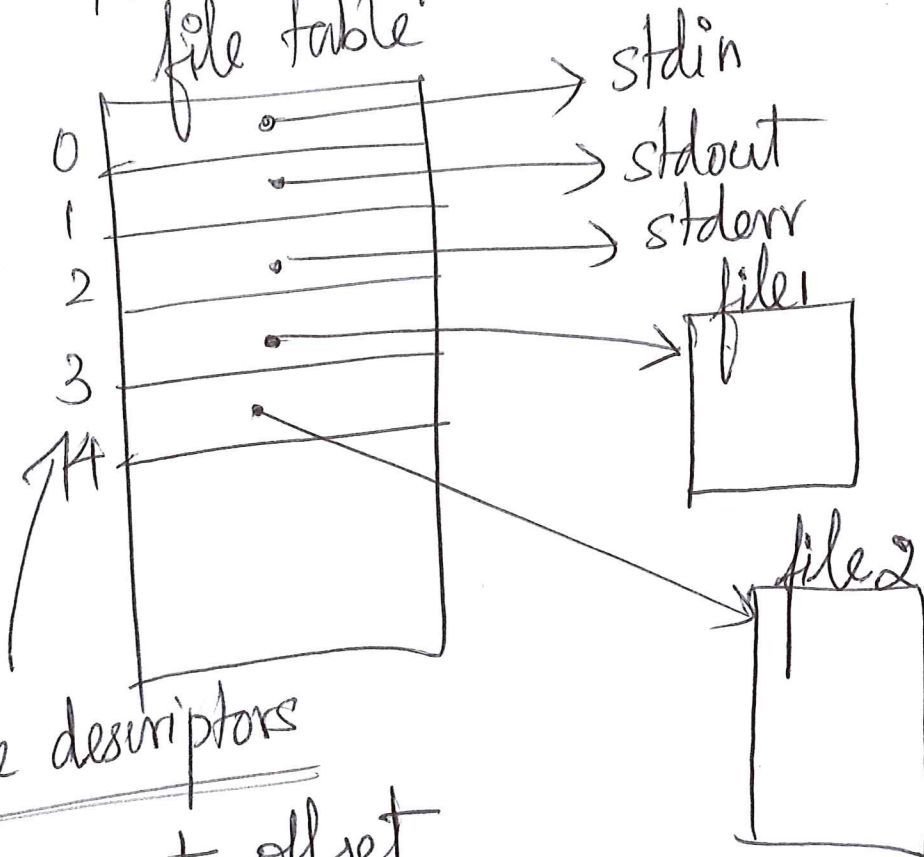
Files

1) Creation

open ()
creat ()

OS


Process - opens a file
file table




file descriptors

current offset

fd = 3.






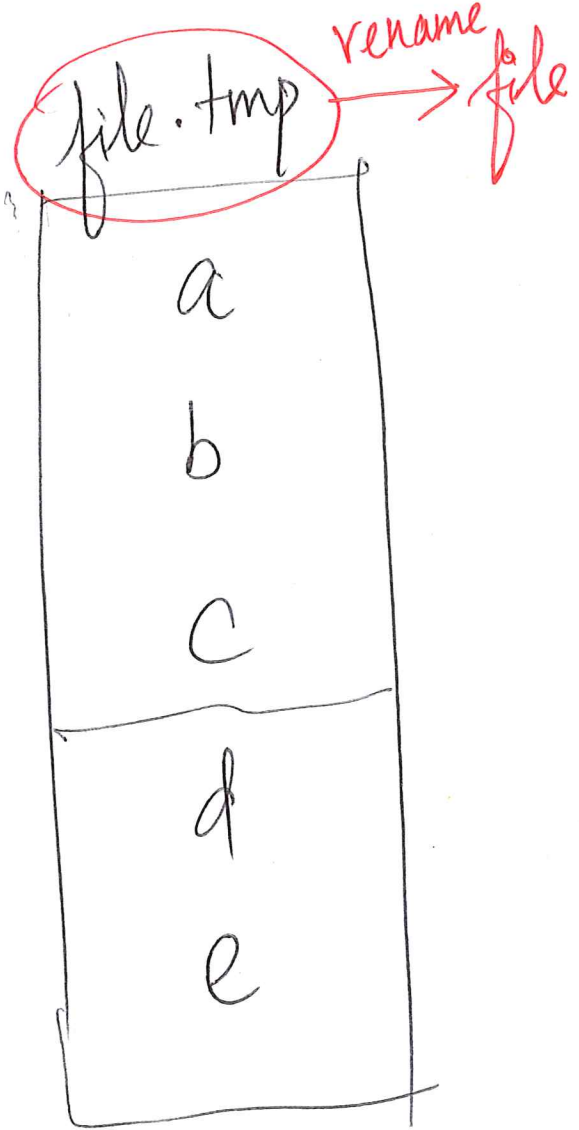
lseek()

lseek(fd, SEEK_END)

rm file → unlink()



write →



Directories

- special type of file
- create, read, delete
- mkdir ls rmdir

- Write - NOT possible

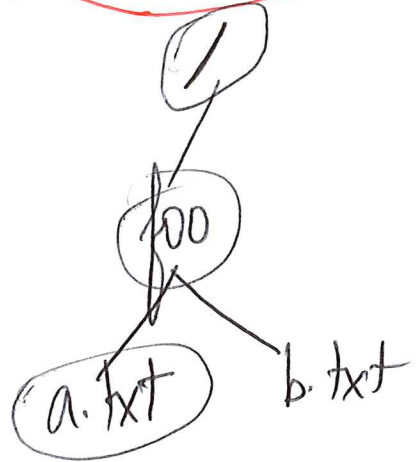

/foo

HN.	inode #
.	1731
..	2
a.txt	1812
b.txt	1012

. → current directory
 .. → parent directory

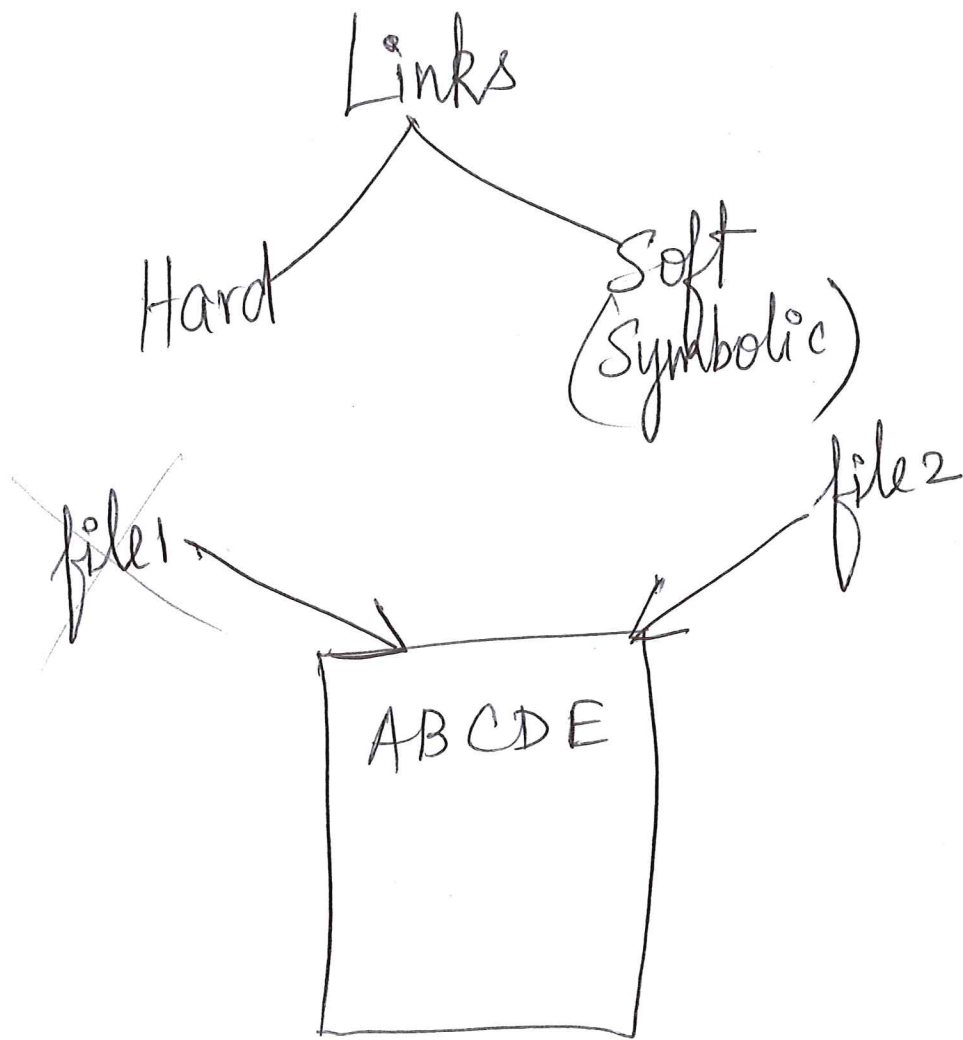
Absolute pathnames: /foo/a.txt

Relative pathname: a.txt



open(/foo/a.txt)

- ① / (root)
- ② foo
- ③ a.txt




① Hard Links



> ln ~~file1~~ file2

A. N.	inode#
file1	123
file2	123

→ No hardlinks for dir.



Symbolic links (soft links)




> ln -s file1 soft

File System Implementation

FS \rightarrow array of block
(Disk)

VFS



1 block = 4 KB (8 sectors)

1 sector = 512 bytes

Size of HDD = 64 x 4 KB

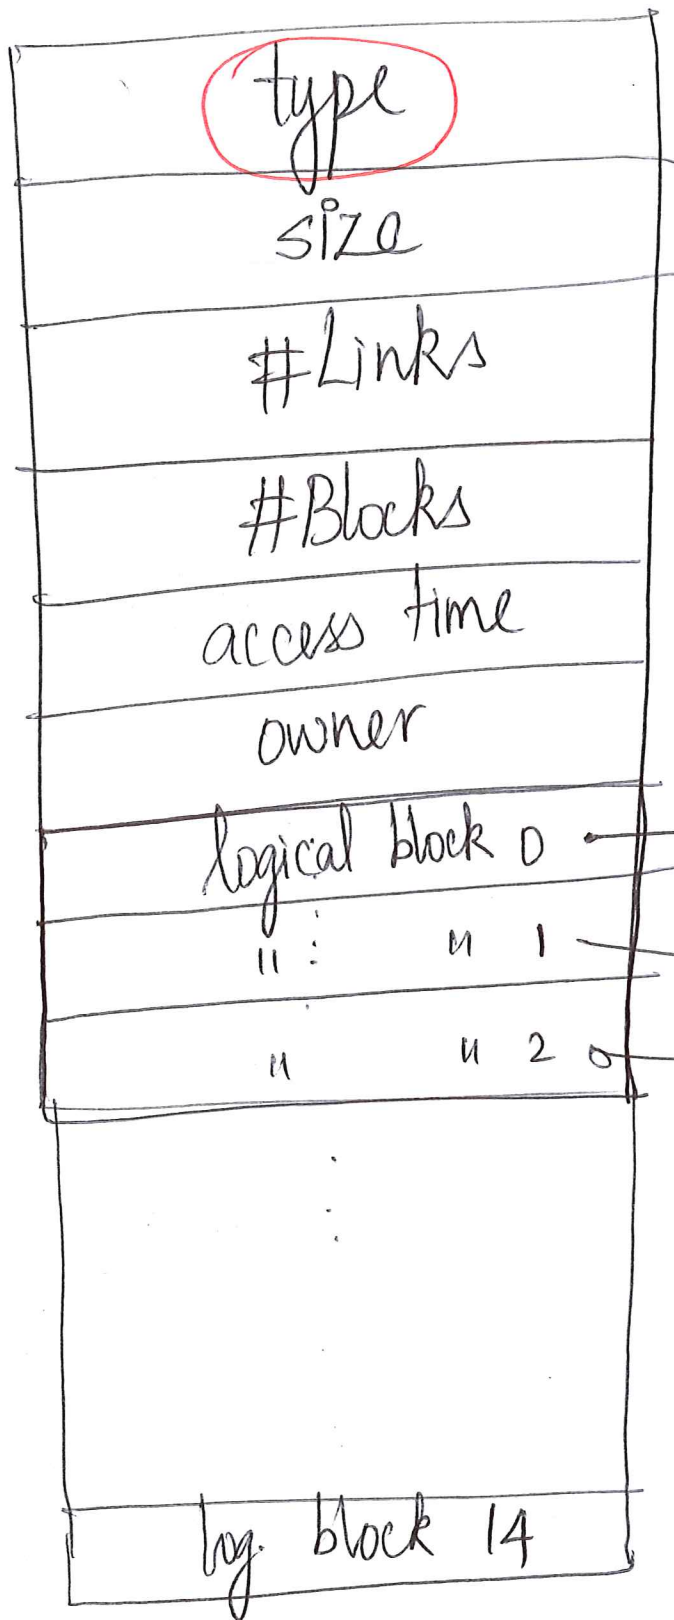
= $2^6 \times 2^2$ KB = 256 KB

① How to store user data?

"DATA blocks"

2

② How to store per-file metadata?



Inode

size = 128 or

256 bytes ✓

max # of files
= 5

1 block $\rightarrow \frac{4KB}{256} = \frac{2^{12}}{2^8} = 2^4 = 16$

\rightarrow Data block 9

\rightarrow Data block 21

\rightarrow Data block 30

1 block \rightarrow 16 inodes

5 blocks \rightarrow 80 inodes

\downarrow
80 files

max. file size = 15 X 4 KB
= 60 KB

Inode



Triple indirection.

max file size
 $= 14 \times 4 \text{ KB} +$
 $\underbrace{1024 \times 1024 \times 4 \text{ KB}}_{\approx \underline{4 \text{ GB}}}$

double indirect ptr.

4 bytes \rightarrow ptr to data block.

of ptrs in indirect block
 $= \frac{4 \text{ KB}}{4} = 1 \text{ KB} = \textcircled{1024}$

max file size = $14 \times 4 \text{ KB} + \underbrace{1024 \times 4 \text{ KB}}_{\textcircled{4 \text{ MB}}}$

FS \rightarrow read block 8?

$$8 \times 4 \text{ KB} = 32 \text{ KB}$$

$$\text{sector \#} = \frac{32 \text{ KB}}{512} = \frac{2^{15}}{2^9} = 2^6 = \textcircled{64}$$

sector #

FS \rightarrow read inode #32?

offset from the
inode start

$$= 32 \times 256 = 2^5 \times 2^8 = \underline{8 \text{ KB}}$$

inode start
from the beg
of disk

$$= 12 \text{ KB} + 8 \text{ KB}$$
$$= \underline{\underline{20 \text{ KB}}}$$

$$\text{sector \#} = \frac{20 \text{ KB}}{512} = \frac{20 \times 2^{10}}{2^9}$$

$$= \textcircled{40} \text{ (sector \#)}$$

③ Directories?



reuse inodes — type = d

Data block

H.N.	inode #
.	123
..	2
a.txt	75

data of dir.

④ How to track free inodes & data blocks?



Read

/foo/bar.txt → ^{data} 12 blocks

- ① Read inode of / (root)
- ② Read data of / (root)
- ③ Read inode of foo.
- ④ Read data of foo.
- ⑤ Read inode of bar.txt
- ⑥ Read data[0] of bar.txt.
- ⑦ Write inode of bar.txt

HN	9#
foo	1234

...	
...	
bar.txt	789

create and write to /foo/bar.txt

root
data

bar.txt
inode

foo
inode

root
inode

db

ib

foo
data

bar
data [0]

R

R

R

R

R

W

R

W

W

R

R

W

W

W

create
foo/bar.txt

x
write()

