

1. Last time
  2. Intro to file systems
  3. Files
  4. Implementing files
  5. Directories
  6. Performance
- 

## 2. Intro to file systems

What does a FS do?

- provide persistence
- create a way to name data on the disk

FS's can be implemented in lots of places

- We focus on the disk, generalize later

Note: disk is the 1<sup>st</sup> thing we've seen that is both modifiable and persistent.

---

### 3. Files

What is a file?

From user's view: a named, contiguous run of bytes

From FS's view: collection of disk blocks

Job of a FS:

map {file, offset in file}  $\xrightarrow{\text{FS}}$  disk address

operations:

create(file), delete(file), read(), write()

Goal: operations have as few disk accesses as possible and minimal space overhead

---

### 4. Implementing files

A. Contiguous

B. Linked files

C. Indexed files

When file's metadata is known

Assume for now that a given file's metadata is stored in the system.

Access patterns to support:

- Sequential
- Random access

Ideal is good sequential + good random access performance.

Candidate designs:

read (fib, 1496)

IBM os/360

A. contiguous allocation

user pre-specifies length

metadata contains location, size

[ <free> a1 a2 a3 <free> b1 b2 <free> ]

+ Simple: easy, low overhead

+ Fast access for both seq + random

- Fragmentation

B. linked files


point to (disk address) of file's first block

metadata is pointer (array)




- + Fragmentation gone
  - + Seq access pretty good + easy
  - Random access is a disaster
  - Pointers take up room in blocks; messes up data alignment
- read (f.l.b, 2500)  
↓  
10<sup>6</sup>)

### C. Indexed files



- + Seq, RA are both easy
  - Metadata enormous
- attempt 2

disk addr of data



attempt 3 (final)  
 imbalanced tree

classic Unix file system

"inode"

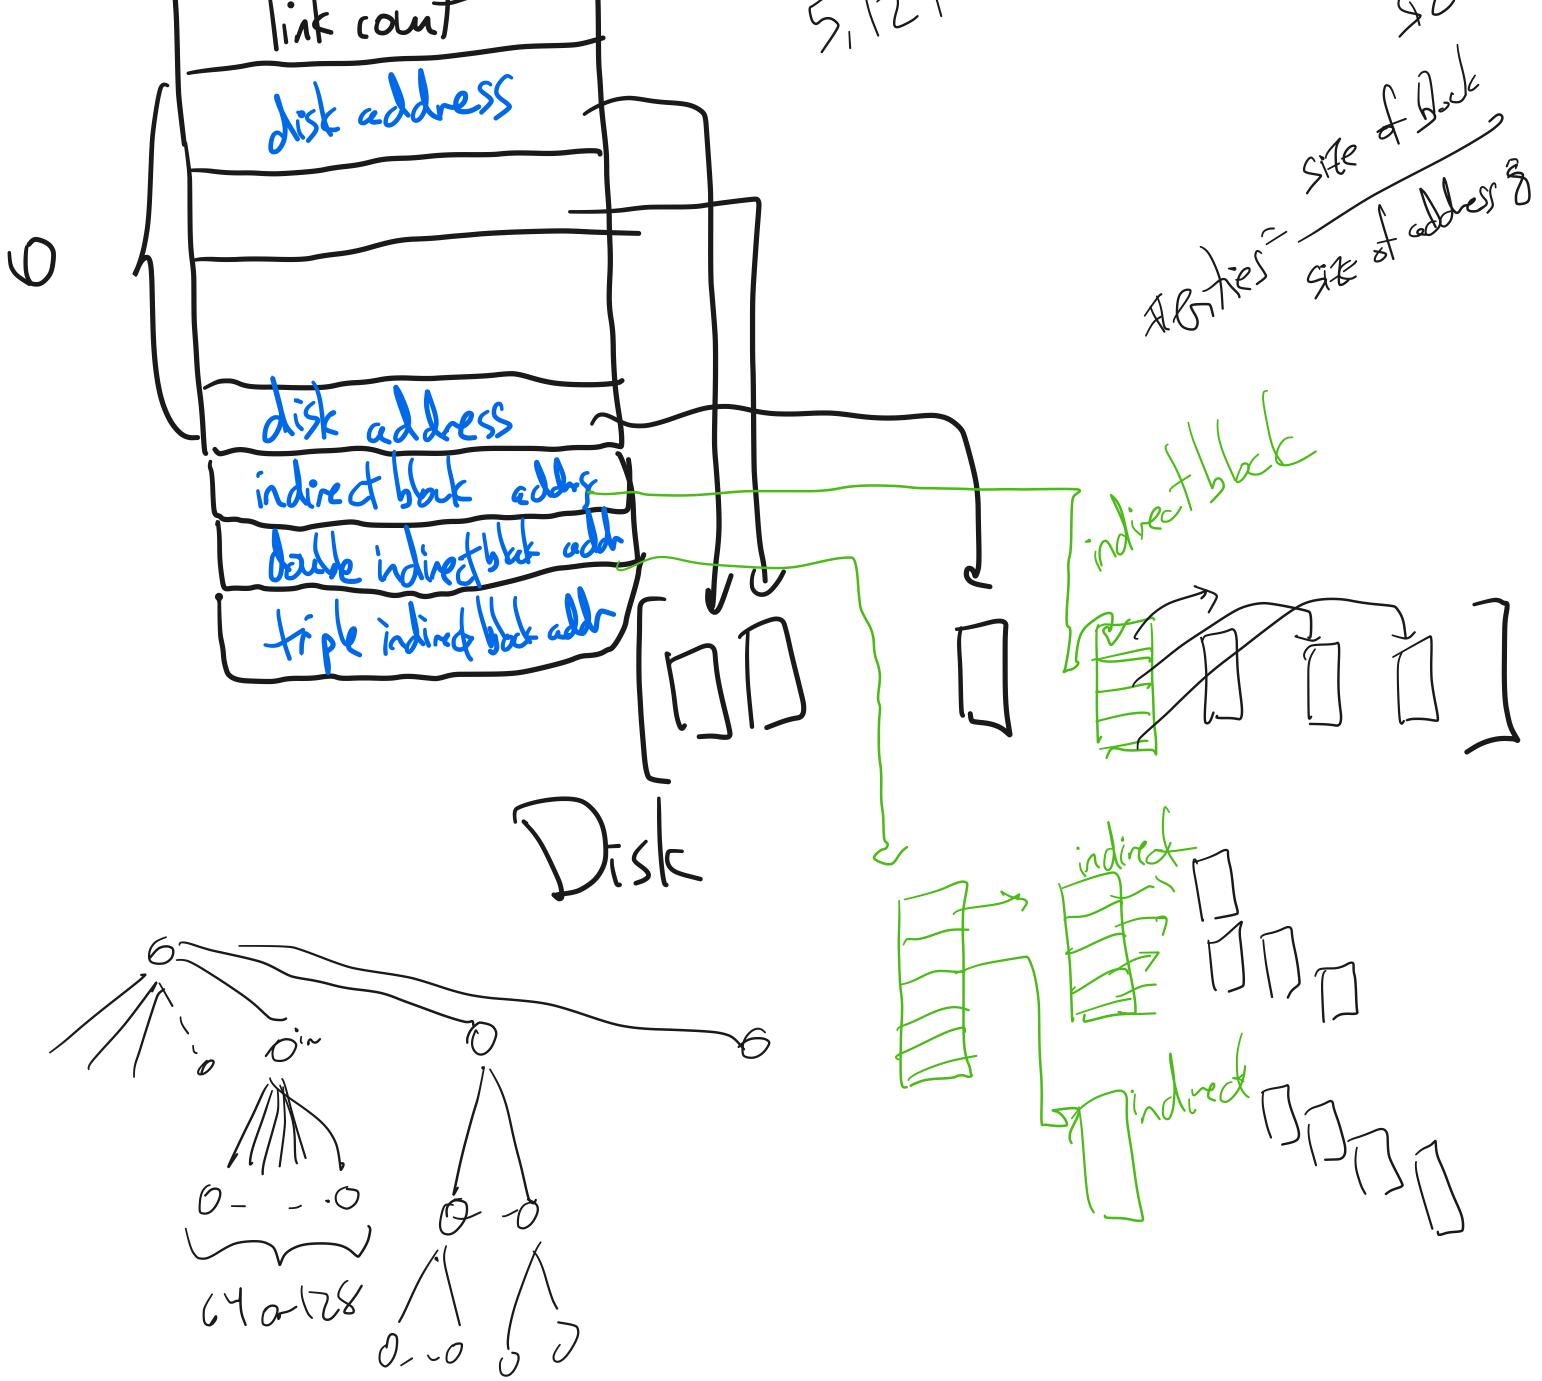
Metadata: inode

inode:

perm S
mtime
atime
ctime
...


5,120

137



- + Simple, easy to build, fast access to small files
- + Max. file length can be enormous
- Worst case seeks/accesses not great
- Worst case overhead (e.g., 11-block file) not ideal
- metadata + data stream throughout the disk

inodes  
 stored in a fixed-size array, known location  
 vocabs: "inumber"



stat (&\_);

5. Directories multics → Unix

Approach 0: Users remember their dir location

Short history of directories

Approach 1: Single dir. for whole system

<u>Dir</u>	
name <sup>1</sup>	inum
name <sup>2</sup>	inum
name <sup>3</sup>	inum
foo,	<u>576</u>
bar,	110578
baz,	205208

answers.txt

Approach 2: Single dir. for each user

- clumsy

- 1,000  $\Rightarrow$  10,000


Approach 3: Hierarchical name space

foo, 573 000

bar, 645 000

mod, dir 2

zoo, dir 3



bin

devcam

sbin

tmp

usr

tcpdump

alice bob



ls, grep, vim

ping

|||

|||

dir:

<name, inode#>

<bin, 1011>


~~bin~~

<sbin, 2048>

:

could be the inum of a directory

inode



bootstrapping:

" / " is inode 2

example:


/a /foo.c



/b/c/essay.txt

/a /b /a/foo.c

12 14 15 16 17

data blocks





ln old-file synonym

delete:

rm  
↓  
unlink