



Welcome Back!

Abstraction



Abstraction



H/W

Where we were:

→ FS API

open, read/write/
close, ^{link}unlink,
mkdir, rmdir, etc.

→ Implementation

simple implementation:

{ → On-disk data structures }


Σ → Access methods



close up

~256 bytes

k



Bitmaps (allocation)

- Type
- size, # of blks
- owner, permissions
- addresses of data blocks

info about each file / directory

(direct pointer)

Super block (meta-info about FS)

Questions:

Access?

Large Files?

Performance:

Use Memory as a cache (DRAM)

pathname:

/ a / b / c / d / main.c

e.g.

(a, 10)

root inode

root data

→

a's inode #

a's inode

a's data

→

b's inode #

b's inode

b's data

c's inode

c's data

d's inode

d's data

→

main.c inode

main.c inode

Large Files:

inode:

type

size

blocks

owner

permissions

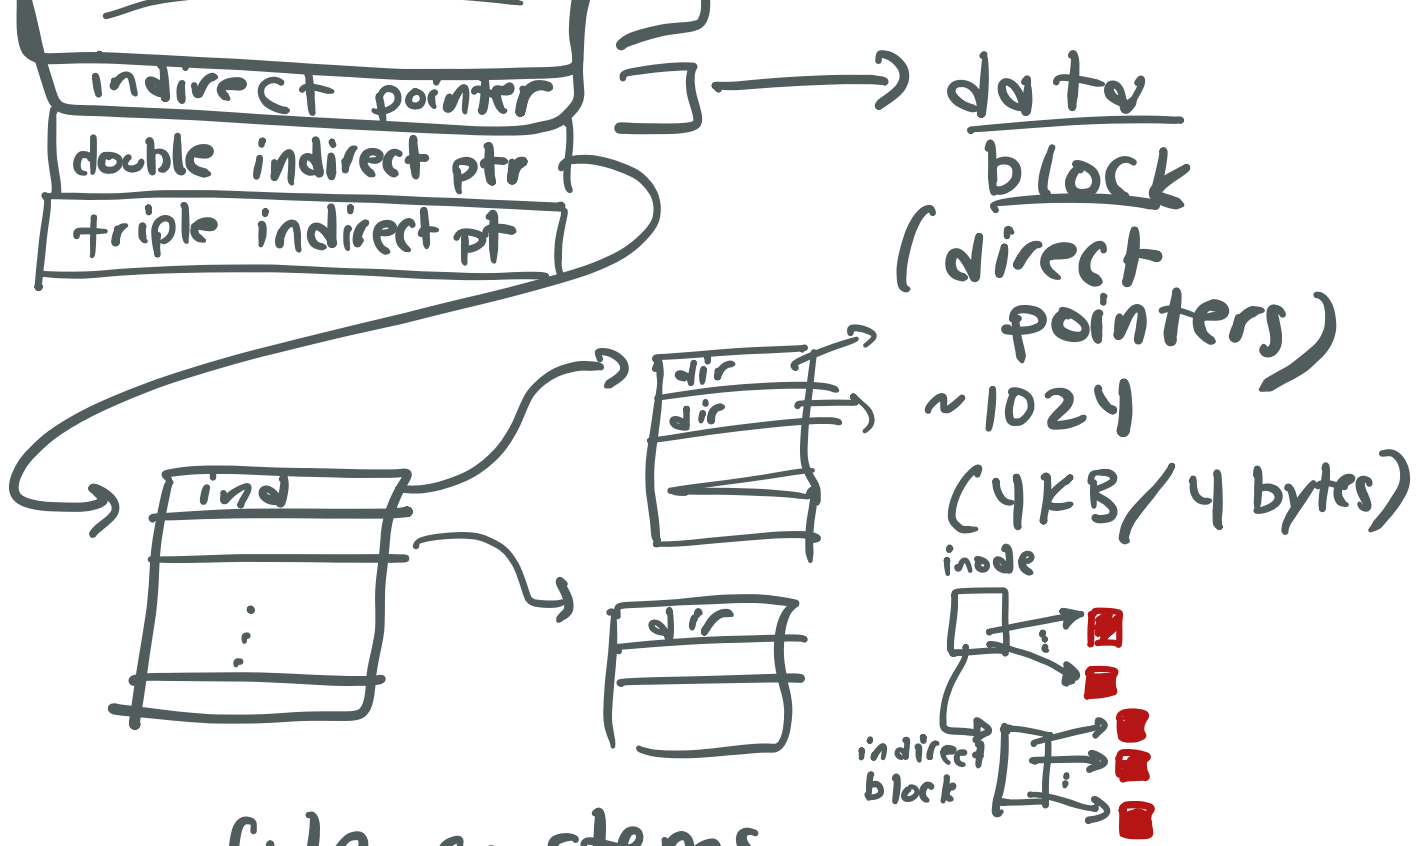
direct pointers

INODE


addresses
of first

N blocks

(N=12)



=> some file systems use extents:
[not just pointers]



=> (Last Project)
coming soon

=> Caching / Buffering ^{write}

↓
frequently accessed
↓
in memory

→ delayed write benefits
+ batch
+ latency
+ avoid I/O (overwrite, delete)

↓
write() → fast
buffers in memory
+ returns
(not immediately written to disk)

(use LRU-like replacement)

=> crash =>
(lost data)

to ensure no data loss:

→ write() (slow)

→ fsync()
forces writes to disk

Locality:

Case study of
early file system

(Berkeley Fast File System)
[FFS]

FFS: main idea

"treat disk like a disk"

(technology aware)

=> simple ideas



"block group":

→ contiguous part of disk

→ allocate "related" items
in a group

(spread out "unrelated" items
into different groups)

block



Allocation Policies:

- mkdir () Directory
- creat () File

spread across disk
(e.g. pick group w/ low # of directories)

put files in same group as parent directory

Large File exception: