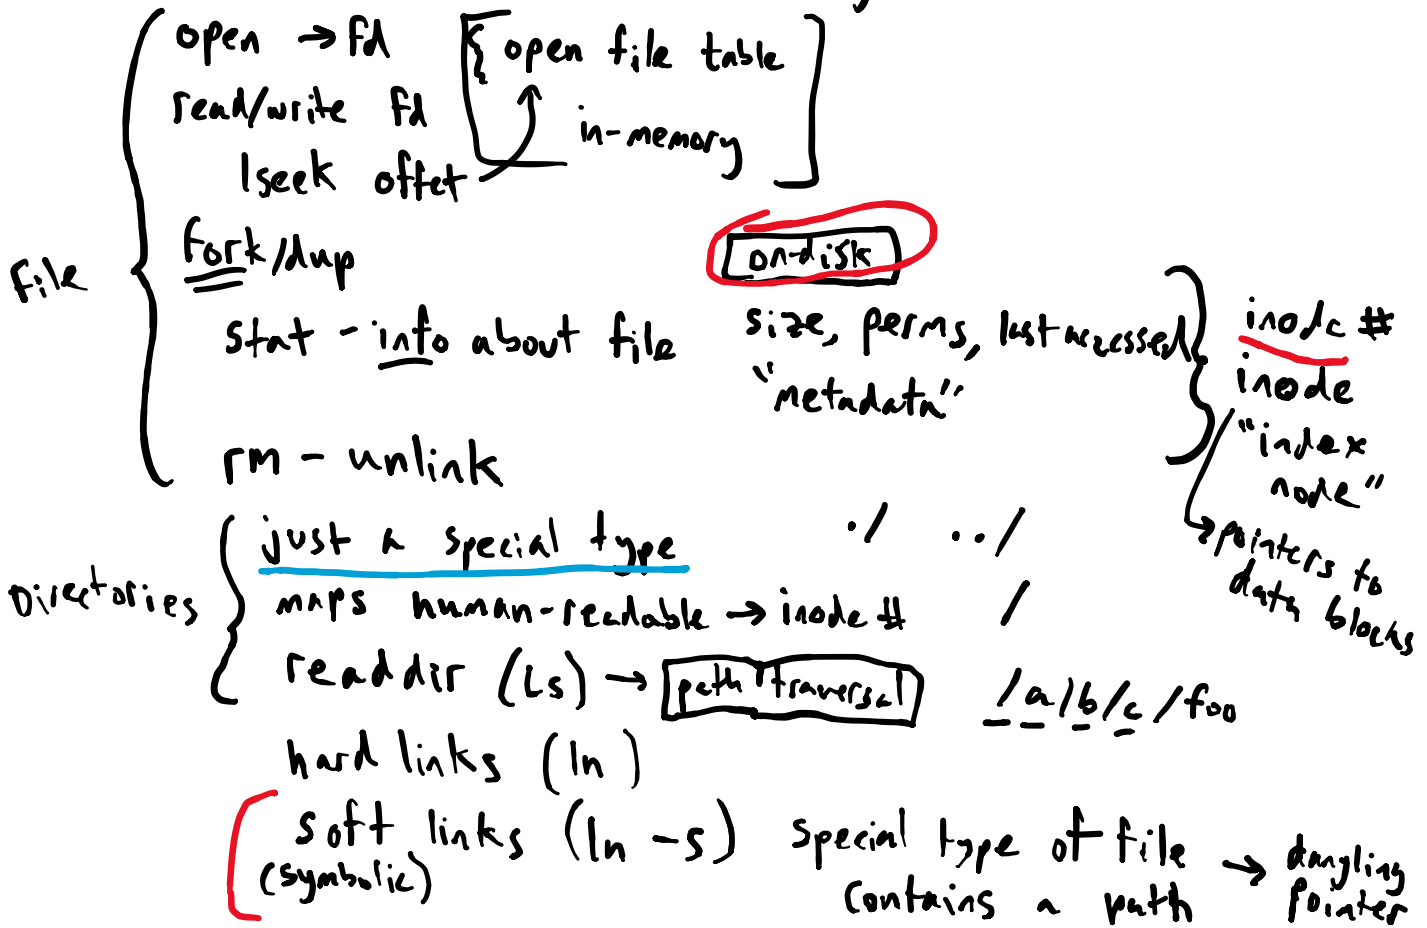


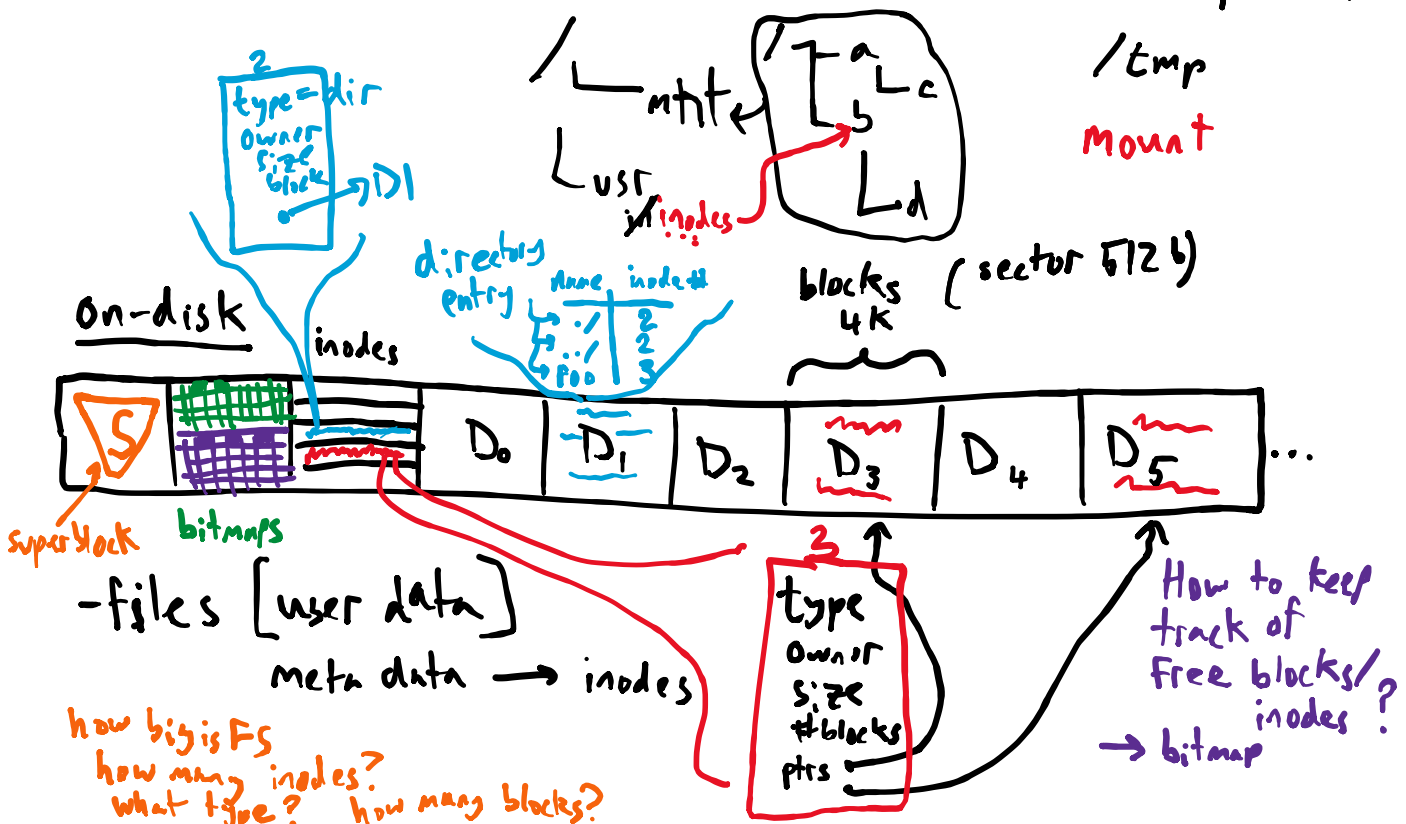
Persistence - I/O device (disk)

FS API files / directory



Mounting different FS

each filesystem is a "volume" → disk partition



metadata about FS
SUPER BLOCK

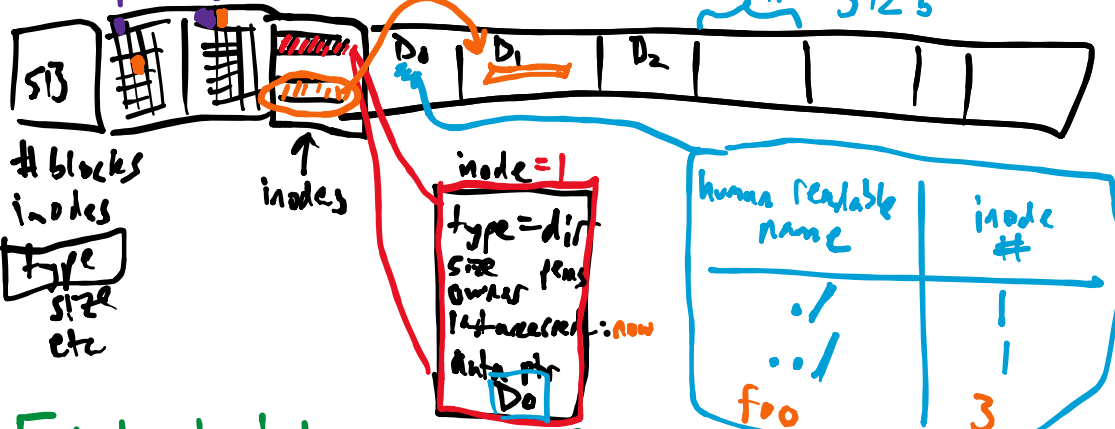
dd

mkfs: create an empty filesystem

mkfs <disk partition>



↳ create on-disk structs so OS can mount



blocks
inodes
type
size
etc

[what about large files?]

Create a file `fd = open("/foo", O_WRONLY, O_CREAT)`

- 1. read / 1. check permissions read inode
- 2. check uniqueness read D0

- 2. create empty file "foo"
 - allocate inode {
 - 1. read inode bitmap
 - 2. write inode bitmap #3
 - 3. write inode type = file
owner = ...
size = 0

- 3. link this to directory
 - 1. update dentry write to data block of /
 - 2. update inode of / write inode of / access time, ...

- 4. create in-memory structures
 - fd → open file table "in-memory inode"
 - return fd

write to file write (fd, buffer, 4k);

1. allocate a data block

read data bitmap

write data bitmap

2. update inode change-size

read inode

write inode

- blk ptr

- modify time, etc...

3. write data block

Close file

free in-memory descriptor

free file struct from open file table

} no disk activity

what about long pathnames

open: what is read from disk " /a/b/c/d/x.txt "

root inode (perms, loc of data)

root data (inode # of a)

a's inode

a's data (look for b's inode #)

b's inode

path traversal

b's data (look for c's inode #)
c's i
c's d
d's i
d's data

Next time xtxt inode (final perm checks)

efficiency : page cache

: what about large files?

