

Administrivia

- mid term solution
 - lab 2 discussion (?) } will publish as soon as possible
 - lab 3 - delayed → } extension
 - concurrency:
 - if you are reading a shared variable, you must protect the read unless you can prove that no other {thread/process} will write to that variable
 - read may not be atomic
 - if you are writing a shared variable, you must protect the write unless you can prove that no other {thread/process} will write or read that variable
 - writes & read may not be atomic
 - race conditions are hard to test for because they usually happen rarely (usually, long after the system has gone into production)
 - need analytic tools
 - careful inspection
 - another pair of eyes
 - message passing: easier to reason about
 - study Erlang
- (can't prove absence of race/deadlock because you need to show true at all possible execution paths
 → can show presence thought it was obvious: (can't prove a negative)
 can prove: "no deadlock here"

Previously on CSS 503A...

- userland vs kernel
 - system calls from userland processes to request kernel services
- concurrency: processes & threads
- filesystems:
 - abstract data structure on top of mass storage devices
- filesystem API: uniform interface
 - regular files
 - devices

- filesystem layers

- application

open/close/read/write/seek

- logical filesystem

filename → file# / file handle / protection

- file organization

inode (file control block)

NTFS: master file table entry

- basic filesystem

block# → drive / cylinder / track

buffer/cache

- I/O control

device access

- device

- electronic signals

File System Implementation

- disk: data structure

boot sector

partition table

super block

magic #

mount count / max mount count

- run disk check utility

block size / blocks / group

free blocks

free inodes

1st inode (/ or mount point)

↳ root file system

* loopback device

- kernel driver:

lets you mount a file as if

it's a separate file system

file control block - aka inode

- multiple filesystem types

- ISO 9660

- Ext 2 - 3 - 4

- NTFS

- FAT 16 / FAT 32

- FUSE

- userland filesystem

" File System in User Space

- Virtual Filesystem

- directory: ^{→ "Special file"} map name → inode #

- inode: data structure

- file metadata

- size

- time stamps

- owner/group

- permission bits

- # links

- ptr to data blocks

- name is not an intrinsic property of a file

- hard link: 2 file names → single node

- Symbolic (soft) link: directory entry that points to inode that points to directory entry

- cannot do hard links to directories

{ - early unix: directory is file

{ - now: not so much

- need inode # + device # to uniquely identify file

File Structure

Super block

- where are the inodes?
- free/allocated block list

inode: file metadata

directory

- "special" file

- bit set in inode

- used to be able to ~~over~~read/close

→ now requires separate ^{set of} system calls

Disk Allocation & Indexing

- memory allocation problem
 - like malloc/new

⊕ Latency

- disk rotation
- read/write head seek

- clever placement can make a world of difference

- 2 basic strategies

- contiguous (Array)
- linked list (sequential)

- or: some combination

Disk Allocation & Indexing (cont.)

- contiguous allocation

OS - 360 / VM / CMS

Allocate "cylinders"

+ minimize disk seek time

- requires pre-allocation (inflexible)

- fragmentation

- "external" fragmentation

Disk Allocation & Indexing (cont.)

- linked list

+ inode (File Control Block) only
requires pointer to 1st block

- Smaller inode

+ Fragmentation is not an issue
(file is pre-fragmented)

- only suitable for sequential access

- link pointers reduce available
block data

- broken link → unrecoverable data

(e.g. write failure due to
OS bug or H/W fault)

→ really bad thing when it
happens

- allocation by large blocks:

best of both worlds performance

- as disk fragments, performance degrades

Windows: defragmentation utility