

CS162  
Operating Systems and  
Systems Programming  
Lecture 17

Disk Management and  
File Systems

October 28, 2009  
Prof. John Kubiawicz  
<http://inst.eecs.berkeley.edu/~cs162>

Review: Want Standard Interfaces to Devices

- **Block Devices:** e.g. disk drives, tape drives, Cdrom
  - Access blocks of data
  - Commands include `open()`, `read()`, `write()`, `seek()`
  - Raw I/O or file-system access
  - Memory-mapped file access possible
- **Character Devices:** e.g. keyboards, mice, serial ports, some USB devices
  - Single characters at a time
  - Commands include `get()`, `put()`
  - Libraries layered on top allow line editing
- **Network Devices:** e.g. Ethernet, Wireless, Bluetooth
  - Different enough from block/character to have own interface
  - Unix and Windows include `socket` interface
    - \* Separates network protocol from network operation
    - \* Includes `select()` functionality
  - Usage: pipes, FIFOs, streams, queues, mailboxes

10/28/09

Kubiawicz CS162 @UCB Fall 2009

Lec 17.2

Review: How Does User Deal with Timing?

- **Blocking Interface:** "Wait"
  - When request data (e.g. `read()` system call), put process to sleep until data is ready
  - When write data (e.g. `write()` system call), put process to sleep until device is ready for data
- **Non-blocking Interface:** "Don't Wait"
  - Returns quickly from read or write request with count of bytes successfully transferred
  - Read may return nothing, write may write nothing
- **Asynchronous Interface:** "Tell Me Later"
  - When request data, take pointer to user's buffer, return immediately; later kernel fills buffer and notifies user
  - When send data, take pointer to user's buffer, return immediately; later kernel takes data and notifies user

10/28/09

Kubiawicz CS162 @UCB Fall 2009

Lec 17.3

Goals for Today

- Finish Discussing I/O Systems
  - Hardware Access
  - Device Drivers
- Disk Performance
  - Hardware performance parameters
  - Queuing Theory
- File Systems
  - Structure, Naming, Directories, and Caching

Note: Some slides and/or pictures in the following are adapted from slides ©2005 Silberschatz, Galvin, and Gagne. Many slides generated from my lecture notes by Kubiawicz.

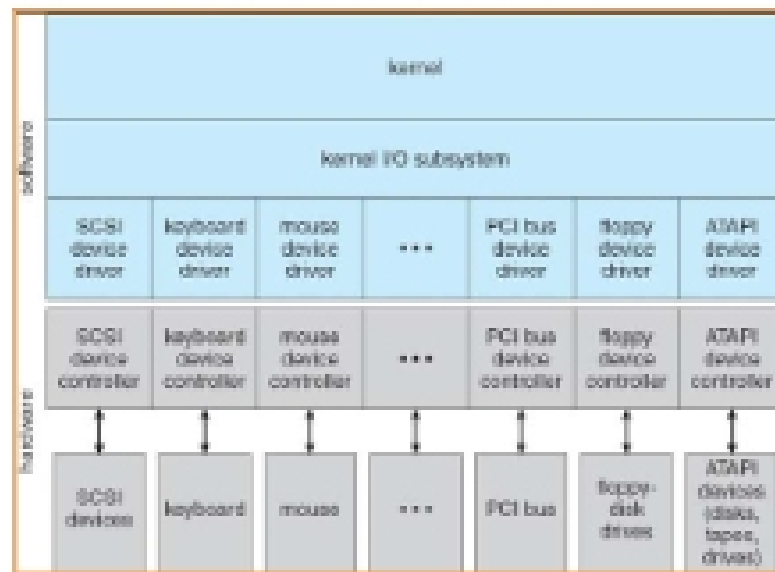
10/28/09

Kubiawicz CS162 @UCB Fall 2009

Lec 17.4



## A Kernel I/O Structure




10/28/00

Kubiatowicz CS162 @UCB Fall 2000

Lec 17.0

## Administrivia

- Midterm I results:
  - Mean: 65.6, STD: 15.0
  - Min: 33.5, Max: 98.0
- Gloukop Grades:
  - Will try to update sooner
- Group Evaluations (Both Projects 1 and 2)
  - These **MUST** be done: you will get ZERO if you don't fill them out
  - Fill them out honestly (you will be potentially asked about them)
- Next Week's Sections
  - Fill out a survey form to see how class is going
  - Give you an opportunity to give feedback
- Other things
  - Group problems? Don't wait.
  - Talk to TA/talk to me
  - Let's get things fixed!



10/28/00

Kubiatowicz CS162 @UCB Fall 2000

Lec 17.10

## Device Drivers

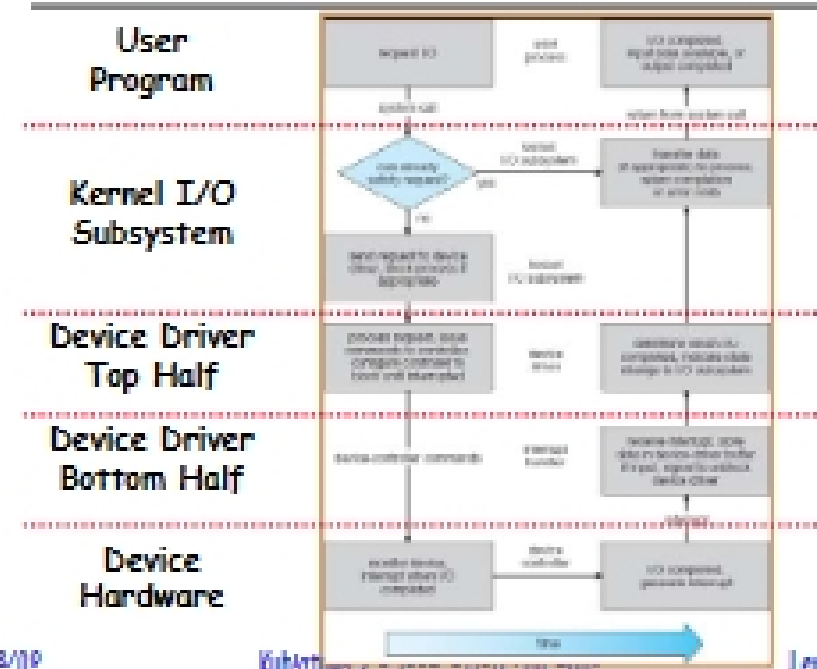
- **Device Driver:** Device-specific code in the kernel that interacts directly with the device hardware
  - Supports a standard, internal interface
  - Same kernel I/O system can interact easily with different device drivers
  - Special device-specific configuration supported with the `ioctl()` system call
- Device Drivers typically divided into two pieces:
  - Top half: accessed in call path from system calls
    - implements a set of **standard, cross-device** calls like `open()`, `close()`, `read()`, `write()`, `ioctl()`, `strategy()`
    - This is the kernel's interface to the device driver
    - Top half will *start* I/O to device, may put thread to sleep until finished
  - Bottom half: run as interrupt routine
    - Gets input or transfers next block of output
    - May wake sleeping threads if I/O now complete

10/28/00

Kubiatowicz CS162 @UCB Fall 2000

Lec 17.11

## Life Cycle of An I/O Request



10/28/00

Kubiatowicz CS162 @UCB Fall 2000

Lec 17.12