

Slide 1

Administrivia

- No class Wednesday.

Slide 2

Minute Essay From Last Lecture — Some Responses

- Almost unanimous preference for Windows scheme, as making it clearer what things are devices. It's a point . . .
- Which one is easier to program — I'm not sure! possibly equivalent.
- One clever thing allowed by UNIX naming conventions, though, is things such as Linux pseudo-filesystem `/proc`. Setting some kernel-level variables is done by — writing to “files” in this “directory”.

I/O Continued — Device Specifics

- Next, a tour of major classes of devices. For each, we look first at what the hardware can typically do, and then at what kinds of device-driver functionality we might want to provide.

Slide 3

Disks — Hardware

- Magnetic disks:
 - Cylinder/head/sector addressing may or may not reflect physical geometry — controller should handle this.
 - Controller may be able to manage multiple disks, perform overlapping seeks.
- RAID (Redundant Array of Inexpensive/Independent Disks):
 - Basic idea is to replace single disk and disk controller with “array” of disks and RAID controller.
 - Two possible payoffs — redundancy and performance (parallelism).
 - Six “levels” (configurations) defined. Read all about it in textbook if interested.
- Optical disks — CD, CD-R, CD-RW, DVD. Okay to skim details!

Slide 4

Disk Formatting

Slide 5

- Low-level formatting — each track filled with sectors (preamble, data, ECC bits).
- Higher-level formatting — master boot record, partitions (logical disks), partition table. Master boot record points to boot block in some partition. Partition table gives info about partitions (size, location, use).
- Partition formatting — boot block, blocks for file system.

Disk Arm Scheduling Algorithms

Slide 6

- A little more about hardware: Time to read a block from disk depends on seek time, rotational delay, and data transfer time. First two usually dominate.
- Earlier we said that typical device driver for disk maintains a queue of pending requests (one per disk, if controller is managing more than one). What order to process them in? several “disk arm scheduling algorithms”:
 - FCFS (first come, first served).
 - SSF (shortest seek first).
 - Elevator.How do they compare with regard to ease of implementation, efficiency?

Disk Error Handling

- Almost all disks have sectors with defects. Some controllers can recognize them (repeated failures) and avoid them; if not, o/s (device driver) must do this.
- Other kinds of errors also possible, e.g., failure to correctly position read/write head; also must be handled either by controller (if possible) or o/s.

Slide 7

Clocks — Hardware

- System clock — can be simple or programmable. Programmable clock can generate either one interrupt after specified interval or periodic interrupts (“clock ticks”).
- Backup clock — usually battery-powered, used at startup and perhaps periodically thereafter.

Slide 8

Clocks — Software

Slide 9

- Clock(s) can be treated as I/O devices, with device driver(s). Functions to provide:
 - Maintain time of day.
 - Enforce time limits on processes.
 - Provide timer / alarm-clock function.
 - Do accounting, profiling, monitoring, etc.
 - Do anything required by page replacement algorithm (turn off R bits in page table entries, e.g.).
- Provide this functionality in code to be called on clock-tick interrupts.

Character-Oriented Terminals — Hardware Overview

Slide 10

- Hardware consists of character-oriented display (fixed number of rows and columns) and keyboard, connected to CPU by serial line.
- Actual hardware no longer common (except in mainframe world), but emulated in software (e.g., UNIX xterm) so old programs still work. (Why does anyone care? some of those old programs are still useful — e.g., text editors — and usually very stable.)

Character-Oriented Terminals — Keyboard

Slide 11

- Hardware transmits individual ASCII characters.
- Device driver can pass them on one by one without processing, or can assemble them into lines and allow editing (erase, line kill, suspend, resume, etc.). Typically provide both modes.
- Device driver should also provide:
 - Buffering, so users can type ahead.
 - Optional echoing.

Character-Oriented Terminals — Display

Slide 12

- Hardware accepts regular characters to display, plus escape sequences (move cursor, turn on/off reverse video, etc.).
In the old days, escape sequences for different kinds of terminals were different — hence the need for a `termcap` database that allows calling programs to be less aware of device-specific details.
- Device driver should provide buffering.

GUIs — Hardware Overview

Slide 13

- PC keyboard — sends very low-level detailed info (keys pressed/released); contrast with keyboard for character-oriented terminal.
- Mouse — sends (delta-x, delta-y, button status) events.
- Display can be vector graphics device (rare now, works in terms of lines, points, text) or raster graphics device (works in terms of pixels). Raster graphics device uses graphics adapter, which includes:
 - Video RAM, mapped to part of memory.
 - Video controller that translates contents of video RAM to display. Typically has two modes, text and bitmap. High-end controllers may incorporate processor(s) and local memory.

GUI Software — Basic Concepts

Slide 14

- “WIMP” — windows, icons, menus, pointing device.
- Can be implemented as integral part of o/s (Windows) or as separate user-space software (UNIX).

GUIs — Keyboard

- Hardware delivers very low-level info (individual key press/release actions).
- Device driver translates these to character codes, typically using configurable keymap.

Slide 15

GUIs — Display (Windows Approach)

- Each window represented by an object, with methods to redraw it.
- Output to display performed by calls to GDI (graphics device interface) — mostly device-independent, vector-graphics oriented. A `.wmf` file (Windows metafile) represents a collection of calls to GDI procedures.

Slide 16

Slide 17

GUIs — Display (UNIX Approach)

- X Window System (its real name) designed to support both local input/output devices and network terminals, in terms of:
 - Programs that want to do GUI I/O.
 - Program that provides GUI services. Can run on the same system as applications, a different UNIX system, an X terminal (where it's the "o/s"), or under another o/s ("X emulators" for Windows — e.g., Exceed, XFree86).

Which is the "client" and which the "server"?

- Core system is client/server communication protocol (input, display events akin to those in Windows) and windowing system. "Window manager" and/or "desktop environment" is separate, as are "widget" libraries. Modularity makes for flexibility and portability, at a cost in performance.

Slide 18

GUI-Based Programming

- Input from keyboard and mouse captured by o/s and turned into messages to process owning appropriate window.
- Typical structure of GUI-based program is a loop to receive and dispatch these messages — "event-driven" style of programming.
- Details vary between Windows and X, but overall idea is similar. See example programs in textbook.

Network Terminals — Hardware

- Keyboard, mouse, and display as described previously, plus local processor; connected to remote system.
- Local processor can be very capable (X terminal, or even PC configured to run as one) or more primitive.

Slide 19

Other I/O-Related Topics

- “Stable storage” — use two disks to provide what appears to be a single more reliable one (i.e., write either succeeds or leaves old data in place).
- Power management significant — some devices have “sleeping” and “hibernating” states, o/s can try to determine when it would make sense to use them. Example — screen blanking.

Slide 20

I/O in UNIX/Linux

Slide 21

- Access to devices provided by special files (normally in `/dev/*`), to provide uniform interface for callers. Two categories, block and character. Each defines interface (set of functions) to device driver. Major device number used to locate specific function.
- For block devices, buffer cache contains blocks recently/frequently used.
- For character devices, optional line-discipline layer provides some of what we described for text-terminal keyboard driver.
- Streams provide additional layer of abstraction for callers — can interface to files, terminals, etc. (This is what you access with `*scanf`, `*printf`.)

I/O in Windows

Slide 22

- Hardware Abstraction Layer (HAL) attempts to insulate rest of o/s from some low-level details — e.g., I/O using ports versus memory-mapped I/O.
- Standard interface to device drivers — Windows Driver Model. Drivers are passed I/O Request Packet objects.

Minute Essay

- This wraps up what I have to say about I/O. Anything else you want to hear (more) about?
- (And best wishes for a good holiday!)

Slide 23