

Administrivia

- Homework 4 (problems about I/O) will be on the Web soon, to be due a week from Monday.

Slide 1

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 2

Disk Formatting

Slide 3

- 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 (more about that in next chapter).

Disk Arm Scheduling Algorithms

Slide 4

- 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 5

Other I/O-Related Topics

- 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 6

I/O in Unix/Linux

Slide 7

- Access to devices provided by special files (`/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. (See figure on p. 729.)
- For character devices, optional line-discipline layer provides some of what we described for text-terminal keyboard driver. (See figure on p. 729.)
- Streams provide additional layer of abstraction for callers — can interface to files, terminals, etc. (This is what you access with `*scanf`, `*printf`.)
(Aside: How do you get the man page for the `printf` function? (`man printf` gives you something else.) Can be several man pages for given name, in different “sections”. Get all of them with `man -a`.)

I/O in Windows

Slide 8

- 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. (See figure p. 779.)
- Standard interface to device drivers — Windows Driver Model. Drivers are passed I/O Request Packet objects. (See figure on p. 829.)
- Interesting comparison of o/s sizes on p. 771.

Minute Essay

- This wraps up what I plan to say about I/O (though not about filesystems, which we'll talk about next week). Anything else you'd like to hear about?

Slide 9