

Administrivia

Slide 1

- Reminder: Homework 4 due Friday. (Questions? and/or I have office hours this afternoon.)
- Courses next semester (shameless self-promotion?):
 - CSCI 3294 (Unix Power Tools) — similar to course two years ago, syllabus and notes on the Web, linked from my home page under “Old course materials”.
 - CSCI 3366 (Parallel Processing).

I/O Management

Slide 2

- Operating system as resource manager — share I/O devices among processes/users.
- Operating system as virtual machine — hide details of interaction with devices, present a nicer interface to application programs.

I/O Hardware, Revisited

Slide 3

- First, a review of I/O hardware — simplified and somewhat abstract view, mostly focusing on how low-level programs communicate with it.
- Many, many kinds of I/O devices — disks, tapes, mice, screens, etc., etc. Can be useful to try to classify as “block devices” versus “character devices”.
- Many/most devices are connected to CPU via a “device controller” that manages low-level details — so o/s talks to controller, not directly to device.
- Interaction between CPU and controllers is via registers in controller (write to tell controller to do something, read to inquire about status), plus (sometimes) data buffer.

Example — parallel port (connected to printers, etc.) has control register (example bit — linefeed), status register (example bit — busy), data register (one byte of data). These map onto the wires connecting the device to the CPU.

Accessing Device Controller Registers

Slide 4

- Two basic approaches:
 - Define “I/O ports” and access via special instructions.
 - “Memory-mapped I/O” — map some (real) addresses to device-controller registers.
- Some systems use hybrid approach.
- Making either one work requires some hardware complexity, and there are tradeoffs; memory-mapped I/O currently more common.

Direct Memory Access (DMA)

Slide 5

- When reading more than one byte (e.g., from disk), device controller typically reads into internal buffer, checking for errors. How to then transfer to memory?
- One way — CPU makes transfer, byte by byte.
- Another way — DMA controller makes transfer, having been given a target memory location and a count.
- Which is better? consider speed of DMA versus speed of CPU, potential for overlapping data transfer and computation. DMA is extra hardware and could be slower than CPU, but would appear to offer potential to overlap transfer and computation.

Interrupts, Revisited

Slide 6

- When I/O device finishes its work, it generates interrupt, typically actually signalling interrupt controller.
Interrupt controller signals CPU, with indication of which device caused interrupt, or ignores interrupt (so device controller keeps trying) if interrupt can't be processed right now.
- Processing is now similar to what happens on traps (interrupts generated by system calls, page faults, other errors):
Hardware locates proper interrupt handler (probably using interrupt vector), saves critical info such as program counter, and transfers control (probably switching into supervisor mode).
Interrupt handler saves other info needed to restart interrupted process, tells interrupt controller when another interrupt can be handled, and performs minimal processing of interrupt.

Interrupts, Revisited, A Bit More

- Notice that pipelining complicates things — restarting is much easier with precise interrupts (all instructions before interrupted one complete, none past interrupted one complete, etc.), but these are difficult to get with pipelined processor.

Slide 7

Mechanics of I/O — Polling Versus Interrupts

- Programmed I/O: Program tells controller what to do and busy-waits until it says it's done.
Simple but potentially inefficient.
- Interrupt-driven I/O: Program tells controller what to do and then blocks. While it's blocked, other processes run. When requested operation is done, controller generates interrupt, interrupt handler unblocks original program,
- I/O using DMA: Similar to interrupt-driven I/O, but transfer of data to memory done by DMA controller, only one interrupt per block of data.

Slide 8

Minute Essay

- We talked about two approaches to communicating with I/O devices — special I/O instructions, and “memory-mapped I/O” (reading/writing particular memory locations). What implications do you think the two choices have for programmers’ ability to write device drivers in a (moderately) high-level language such as C?

Slide 9

Minute Essay Answer

- With memory-mapped I/O it should be possible to write device drivers entirely in C; with special I/O instructions this would not be possible without compiler modifications or some amount of assembly-language code.

Slide 10