CS110 Discussion 14
OS, I/O, DMA, Networks

Tianyuan Wu

ShanghaiTech University

Short Intro to OS

- One of the first things that runs when your computer starts
- Loads, runs and manages programs
- Services: File System, Network stack, printer, etc.
- Finds and controls all the devices in the machine in a general way

What does OS do?

- As referee: Allow fair sharing of resources among applications:
  - e.g., Scheduler: Fair share of CPU, disk & network
- As illusionist: Provide the application with “infinite” resources:
  - e.g., VM & scheduler: illusion of having dedicated CPU and all memory;
- As glue: Provide the application with standard service interface:
  - e.g., System Calls & File system: standard interface for disk access

IO – Program interact with outside world

- What must the processor do for I/O?
  - Input: reads a sequence of bytes
  - Output: writes a sequence of bytes
- Interface options
  - Some processors have special input/output instructions (e.g., Intel X86)
  - Memory Mapped Input/Output (e.g.,RISC-V)
  - RISC-V: https://riscv.org/wp-content/uploads/2017/05/riscvprivileged-v1.10.pdf section 3.5

I/O - Polling

- Consistently check the device for the data to read or write
- Control register and data register

- Input: Read from keyboard into a0

```
                li    t0, 0xffff0000 #ffff0000
Waitloop:       lw    t1, 0(t0)      #control
                andi  t1, t1,0x1
                beq   t1, zero, Waitloop
                lw    a0, 4(t0)      #data
```

- Output: Write to display from a0

```
                li    t0, 0xffff0000 #ffff0000
Waitloop:       lw    t1, 8(t0)      #control
                andi  t1, t1,0x1
                beq   t1, zero, Waitloop
                sw    a0, 12(t0)     #data
```

I/O – Interrupt

- Let the device notify the CPU when a data is ready

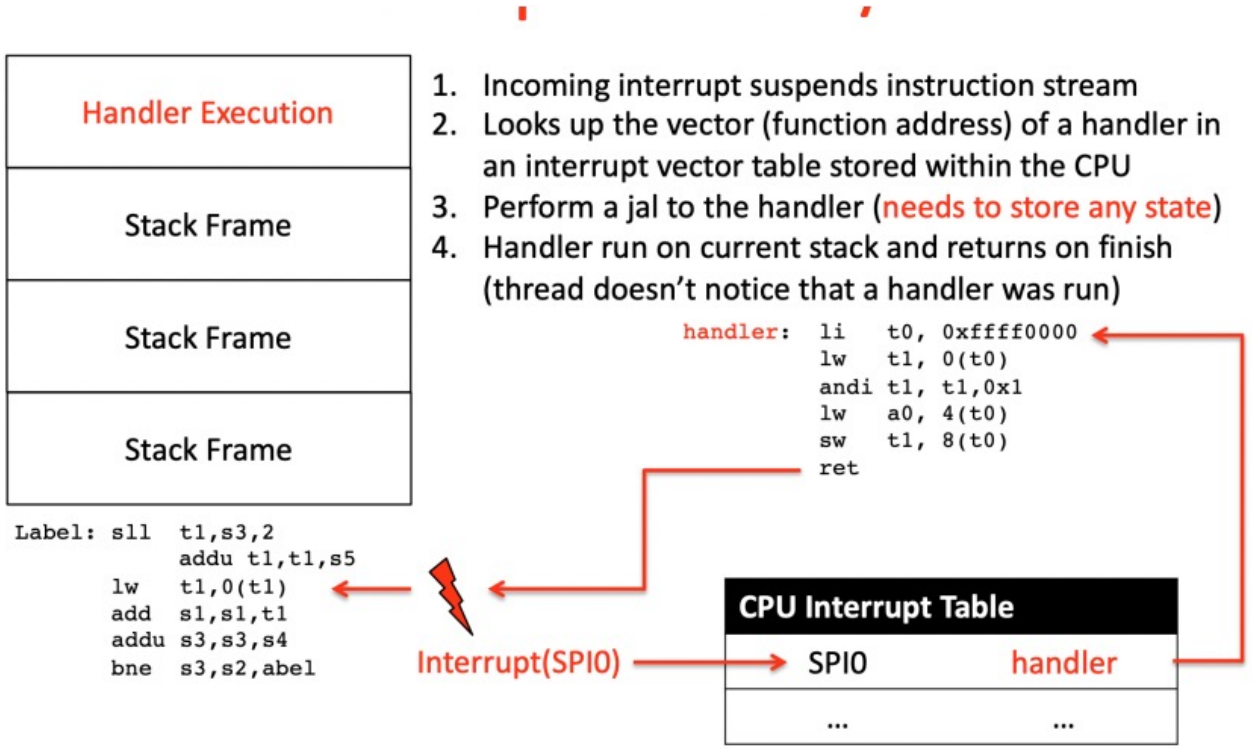

Memory mapped I/O

- Certain addresses are not regular memory
- Instead, they correspond to registers in I/O devices
  - Control Register, says if it's OK to read/write (I/O ready)
  - Data Register, contains data

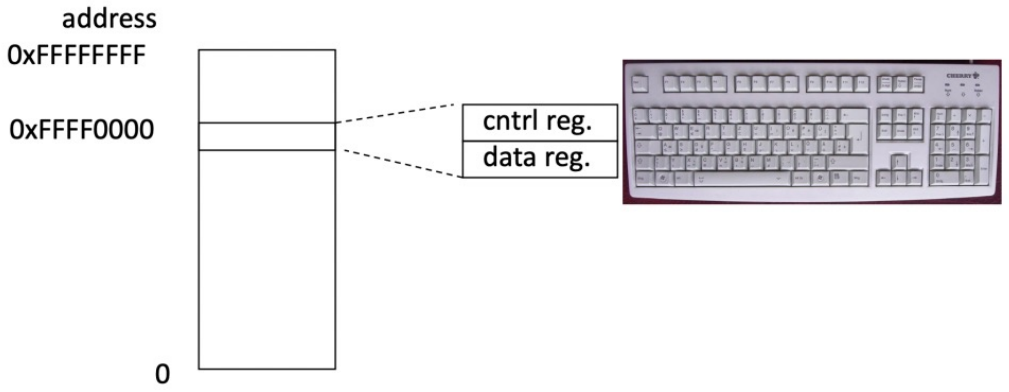

I/O is slow!

- IO is very very slow
- If one CPU cycle (~0.5 ns)
- Read from main memory (~80-100 ns)
- Read from Pmem (~350 ns)*
- Read from SATA SSD (~50 us)
- Read from HDD (~1 ms)
- Internet RTT from Shanghai to Boston (~300 ms)
- In conclusion: We must free CPU from waiting for them…

* Data collected by Intel Optane persistent memory

DMA - Direct Memory Access

- As of now: CPU will do the read/write, but I/O is slow
- DMA: let the device do the I/O instead of the CPU
- CPU could do other things

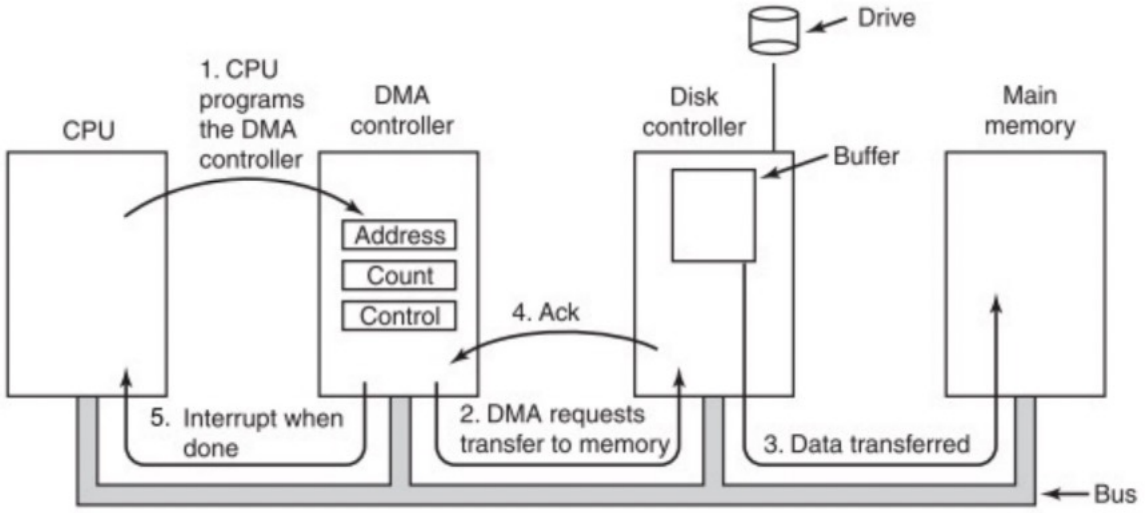

True or False?

- Interrupt has lower latency than polling in general.
- False. Polling has lower latency than interrupt: polling will take fewer steps and interrupt will require the RW request to be queued before the CPU actually do it.

True or False?

- Interrupt is more suitable for high volume data transfer
- True. Despite its higher latency, interrupt provides higher throughput under same CPU load, thus is more suitable for high volume data transfer.

True or False?

- User program can access OS routines with sys calls.
- True. Syscalls provides interface for user program to access some common interface from the OS: e.g., read from the disk.

Thanks!