

## TCSS 422: OPERATING SYSTEMS

### I/O Devices

Wes J. Lloyd  
Institute of Technology  
University of Washington - Tacoma



## OBJECTIVES

- Polling vs Interrupts
- Programmed I/O (PIO)
- Direct memory Access (DMA)
- Port-mapped I/O (PMIO)
- Memory-mapped I/O (MMIO)

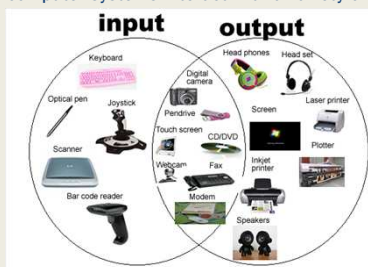
November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.2

## I/O DEVICES

- Modern computer systems interact with a variety of devices




November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.3

## COMPUTER SYSTEM ARCHITECTURE



**VERY FAST:** CPU is attached to main memory via a **Memory bus**  
**FAST:** High speed devices (e.g. video) are connected via a **General I/O bus**.  
**SLOWER:** Disks are connected via a **Peripheral I/O bus**.

November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.4

## I/O BUSES

- Buses
  - Buses closer to the CPU are faster
  - Can support fewer devices
  - Further buses are slower, but support more devices
- Physics and costs dictate "levels"
  - Memory bus
  - General I/O bus
  - Peripheral I/O bus
- Tradeoff space: speed vs. locality


November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.5

## CANONICAL DEVICE

- Consider an arbitrary canonical device



- Two primary components
  - Interface (registers for communication)
  - Internals: Local CPU, memory, specific chips, firmware (embedded software)

November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.6

## CANONICAL DEVICE: HARDWARE INTERFACE

- Status register
  - Maintains current device status
- Command register
  - Where commands for interaction are sent
- Data register
  - Used to send and receive data to the device

**General concept:**  
The OS interacts and controls device behavior by reading and writing the device registers.

November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.7

## OS DEVICE INTERACTION

- Common example of device interaction

```
while ( STATUS == BUSY) ← Poll- Is device available?
: //wait until device is not busy
write data to data register ← Command parameterization
write command to command register ← Send command
Doing so starts the device and executes the command
while ( STATUS == BUSY) ← Poll - Is device done?
: //wait until device is done with your request
```


November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.8

## POLLING

- OS checks if device is **READY** by repeatedly checking the **STATUS** register
  - Simple approach
  - CPU cycles are wasted without doing meaningful work
  - Ok if only a few cycles, for rapid devices that are often **READY**
  - BUT polling, as with "spin locks" we understand is inefficient




November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.9

## INTERRUPTS VS POLLING

- For longer waits, put process waiting on I/O to sleep
- Context switch (C/S) to another process
- When I/O completes, fire an interrupt to initiate C/S back
  - Advantage: better multi-tasking and CPU utilization
  - Avoids: unproductive CPU cycles (polling)



November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.10

## INTERRUPTS VS POLLING - 2

### What is the tradeoff space ?

- Interrupts are not always the best solution
  - How long does the device I/O require?
  - What is the cost of context switching?

If device I/O is fast → polling is better.  
If device I/O is slow → interrupts are better.

November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.11

## INTERRUPTS VS POLLING - 3

- One solution is a two-phase hybrid approach
  - Initially poll, then sleep and use interrupts
- Livelock problem
  - Common with network I/O
  - Many arriving packets generate **many many** interrupts
  - Overloads the CPU!
  - No time to execute code, just interrupt handlers !
- Livelock optimization
  - Coalesce multiple arriving packets (for different processes) into fewer interrupts
  - Must consider number of interrupts a device could generate

November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.12

## DEVICE I/O

- To interact with a device we must send/receive DATA
- There are two general approaches:
  - Programmed I/O (PIO)
  - Direct memory access (DMA)

"over-burdened"

Diagram of CPU utilization

November 28, 2016
TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma
L20.13

Transfer Modes			
Mode	#	Maximum transfer rate (MB/s)	cycle time
PIO	0	3.3	600 ns
	1	5.2	383 ns
	2	8.3	240 ns
	3	11.1	180 ns
	4	16.7	120 ns
Single-word DMA	0	2.1	960 ns
	1	4.2	480 ns
	2	8.3	240 ns
Multi-word DMA	0	4.2	480 ns
	1	13.3	150 ns
	2	16.7	120 ns
	3 <sup>[34]</sup>	20	100 ns
	4 <sup>[34]</sup>	25	80 ns
Ultra DMA	0	16.7	240 ns + 2
	1	25.0	160 ns + 2
	2 (Ultra ATA/33)	33.3	120 ns + 2
	3	44.4	90 ns + 2
	4 (Ultra ATA/66)	66.7	60 ns + 2
	5 (Ultra ATA/100)	100	40 ns + 2
	6 (Ultra ATA/133)	133	30 ns + 2
	7 (Ultra ATA/167) <sup>[35]</sup>	167	24 ns + 2

From <https://en.wikipedia.org/wiki/Parallel ATA>

## PROGRAMMED I/O (PIO)

- Spend CPU time to perform I/O
- CPU is involved with the data movement (input/output)
- PIO is slow –CPU is occupied with meaningless work

"over-burdened"

Diagram of CPU utilization

November 28, 2016
TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma
L20.15

## PIO DEVICES

- Legacy serial ports
- Legacy parallel ports
- PS/2 keyboard and mouse
- Legacy MIDI, joysticks
- Old network interfaces

November 28, 2016
TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma
L20.16

## DIRECT MEMORY ACCESS (DMA)

- Copy data in memory by offloading to a "DMA controller"
- Many devices (including CPUs) have DMA controllers
- Give DMA memory address, size, and copy instruction
- DMA performs I/O independent of the CPU

Diagram of CPU utilization by DMA

November 28, 2016
TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma
L20.17

## DEVICE INTERACTION

- Two primary methods
  - Port mapped I/O (PMIO)
  - Memory mapped I/O (MMIO)

November 28, 2016
TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma
L20.18

## PORT MAPPED I/O (PMIO)

- Device specific CPU I/O Instructions
- Follows a CISC model: extra instructions
- x86-x86-64: **in** and **out** instructions
- **outb**, **outw**, **outl**
- 1, 2, 4 byte copy from EAX → device's I/O port

November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.19

## MEMORY MAPPED I/O (MMIO)

- Device's memory is mapped to CPU memory
- Tenet of RISC CPUs: instructions are eliminated, CPU is simpler
- Old days: 16-bit CPUs didn't have a lot of spare memory space
- Today's CPUs: 32-bit (4GB addr space) & 64-bit (128 TB addr space)
- Regular CPU instructions used to access device mapped memory
- Devices monitor CPU address bus and respond to their addresses
- I/O device address areas of memory are **reserved** for I/O
  - Must not be available for normal memory operations.

November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.20

## DEVICE INTERACTION

- The OS must interact with a variety of devices
- Example: for DISK I/O consider the variety of disks:
- SCSI, IDE, USB flash drive, DVD, etc.
- Device drivers use abstraction to provide general interfaces for vendor specific hardware
- In Linux: block devices


November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.21

## FILE SYSTEM ABSTRACTION

- Layers of I/O abstraction in Linux
- C functions (open, read, write) issue **block read and write** requests to the generic block layer



November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.22

## FILE SYSTEM ABSTRACTION ISSUES

- **Too much abstraction**
  - Many devices provide special capabilities
  - Example: SCSI Error handling
  - SCSI devices provide extra detail which are lost to the OS
- **Buggy device drivers**
  - 70% of OS code is in device drivers
  - Device drivers are required for every device plugged in
  - Drivers are often 3<sup>rd</sup> party, which is not quality controlled at the same level as the OS (Linux, Windows, MacOS, etc.)

November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.23

## QUESTIONS



November 28, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L20.24