CISC 7310X

C02: Interrupts and I/O

Hui Chen

Department of Computer & Information Science

CUNY Brooklyn College

Von Neumann Computers

- Process and memory connected by a bus (Von Neumann, 1945)

An x86 Realization

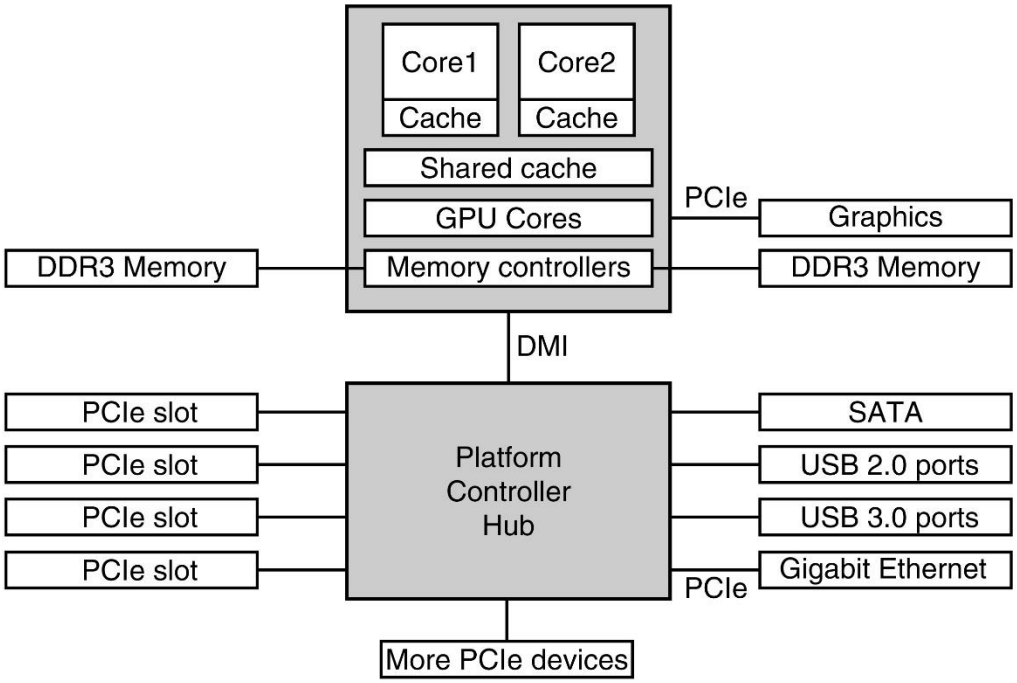

- A structure of a large x86 system [Figure 1-12 in Tanenbaum & Bos, 2014]

Architecture, OS, and Programming

- Architecture underpins design of OS and programming
  - Discussion on existing architecture & design
  - How about the future?

"I propose to call this tube the von Neumann bottleneck. The task of a program is to change the contents of the store in some major way; when one considers that this task must be accomplished entirely by pumping single words back and forth through the von Neumann bottleneck, the reason for its name becomes clear."

– John Backus, 1977

I/O Devices

- A few general categories
  - Storage devices
    - Examples: Disks, tapes, solid state drives
  - Transmission devices
    - Examples: network adapters, modems
  - Human-interface devices
    - Examples: display screens, keyboard, mouse, touch screen
- Specialized devices
  - Examples: control cars, robots, aircrafts, spacecrafts

Port and Bus

- Devices communicate with a computer via a connection point
  - (Physical) port
    - Examples: USB port, serial port, parallel port
- A common set of wires with a protocol that specifies commands that can be transmitted
  - Bus
    - Examples: PCI bus, SCSI bus

Device Controller

- Devices
  - Example: hard disk drives have motors, magnetic headers, and disks
- Controller
  - A collection of electronics that operate a port, a bus, or a device (some contain small embedded computer)
    - Accept and act on commands from the OS
    - Present a simpler interface to the OS
  - Examples: SATA controller

Device Driver

- Device driver
  - Each type of controller is different
  - Software communicates to the controller, and the OS
  - Adhere to some standard when communicating to the OS

I/O Devices and OS

- Reduce complexity, increase uniformity and reliability

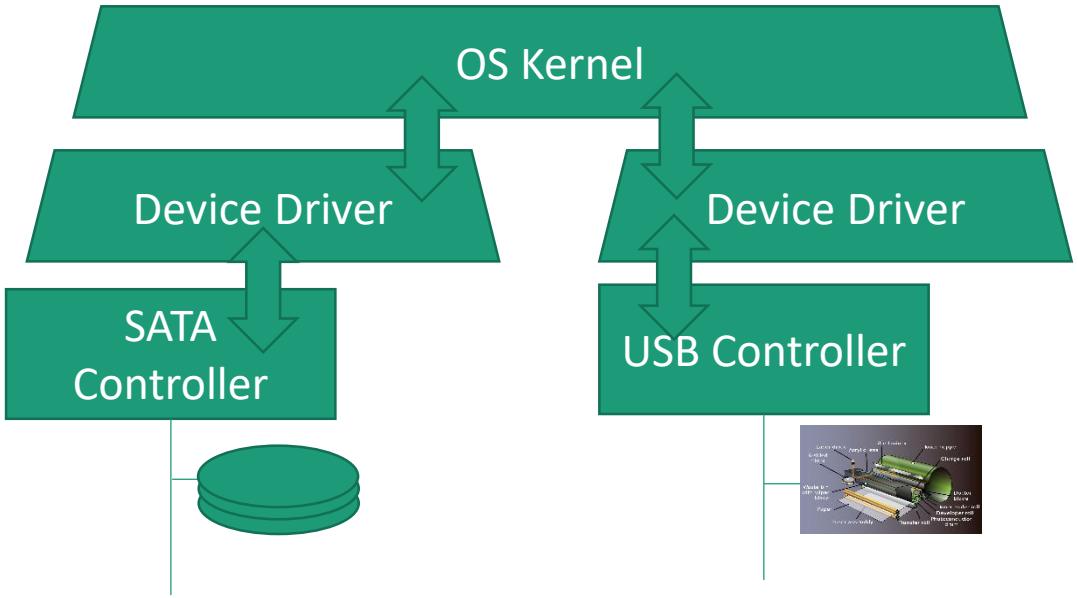

Design Consideration: Access Right

- A design consideration
  - What kind of access right should we give to device drivers?
  - Unrestricted
    - Kernel mode
    - Relatively easier to design, can affect the others
  - Restricted
    - User mode
    - More difficult to design, isolated from the others

Design Consideration: Load Device Drivers

- Relink the kernel with the new deriver
  - Require reboot
- Add to the kernel an entry indicating a new driver is needed
  - Load the driver during reboot
- Install and run the device driver on the fly
  - Hot-pluggable

Questions?

- Overview of I/O devices

Controller Registers

- Typically have 4 registers or more
  - Data-in register
    - Read by the host
  - Data-out register
    - Written by the host
  - Status register
    - A number of bits indicating the status of the device (e.g., busy, error)
  - Control register
    - A number of bits indicating the mode of the device
- May have a data buffer
  - Examples: video adapter (video memory)

Virtual Box

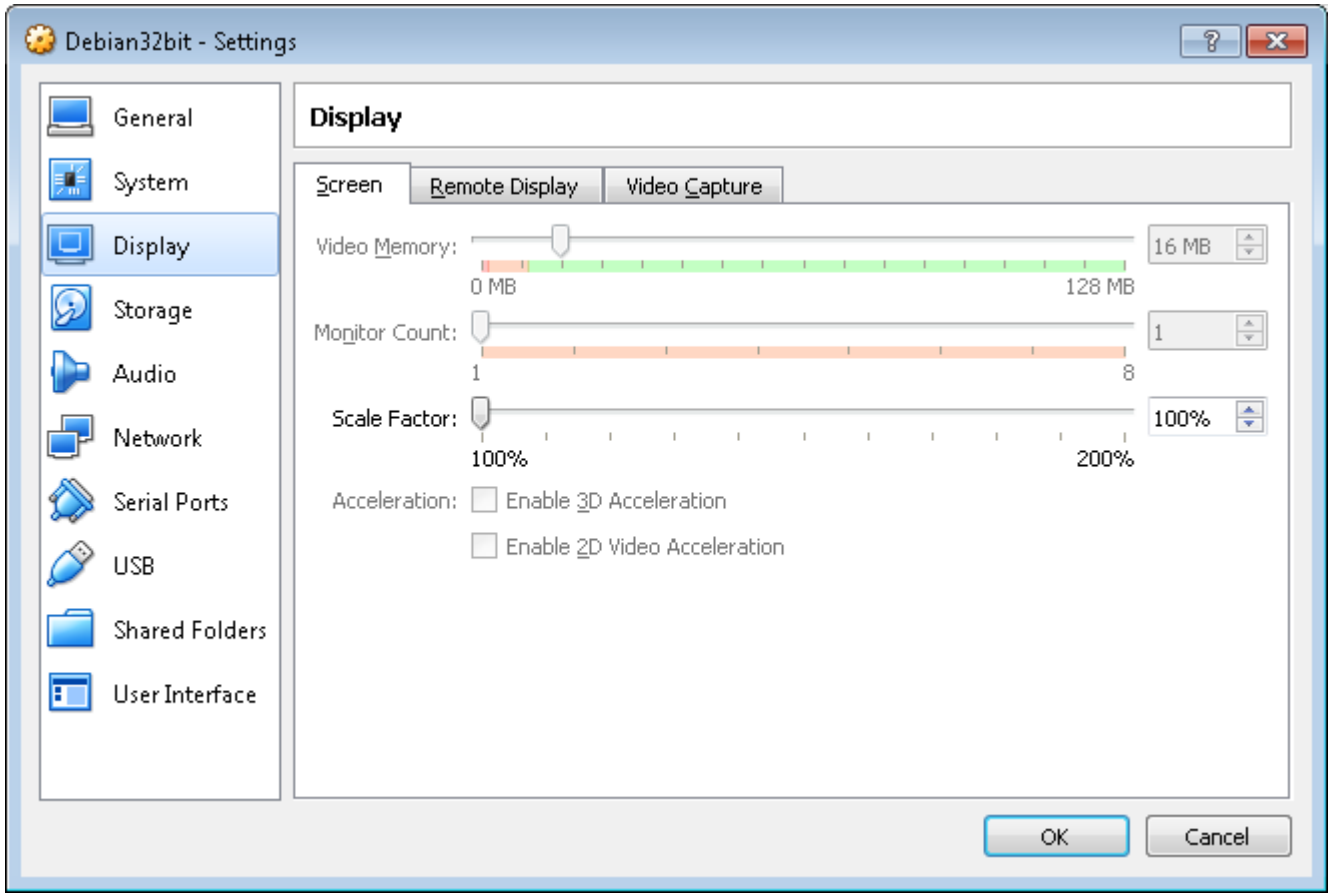
Debian32bit - Settings
General
System
Display
Storage
Audio
Network
Serial Ports
USB
Shared Folders
User Interface
Display
Screen
Remote Display
Video Capture
Video Memory:
16 MB
0 MB
128 MB
Monitor Count:
1
1
8
Scale Factor:
100%
100%
200%
Acceleration:
Enable 3D Acceleration
Enable 2D Video Acceleration
OK
Cancel

Access Device Controller

- CPU read and write to the device controller registers and data buffer
  - (Local) I/O ports
  - Memory mapped I/O

I/O Port Space

- Each register is assigned an I/O port number
  - Typically, a 8-bit or 16-bit integer
  - All I/O port numbers form the I/O port space
- A CPU has I/O instructions
  - Example instruction (in an assembly language):
    - IN REG, PORT
    - OUT PORT, REG

Example I/O Port Allocation

- Some default values on PCs

I/O Address Range	Device
000-00F	DMA Controller
020-021	Interrupt Controller
040-043	Timer
200-20F	Game Controller
2F8-2FF	Serial Port (Secondary)
320-32F	Hard-disk Controller
378-37F	Parallel Port
3D0-3DF	Graphics Controller
3F0-3F7	Diskette-drive Controller
3F8-3FF	Serial Port (Primary)

I/O Instruction

- Example from
  - http://www.tldp.org/HOWTO/text/IO-Port-Programming
- Source
  - https://github.com/CISC7310SP18/SamplePrograms/tree/master/W2_IO/ioport

Memory Mapped I/O

- Map all the control registers into the memory address space
  - A register is assigned to a unique memory address to which no memory is assigned
  - Accessing these registers as if they were main memory
- Hybrid scheme
  - Data buffers are mapped to memory address
  - Control registers have dedicated I/O ports

Accessing Device Controllers

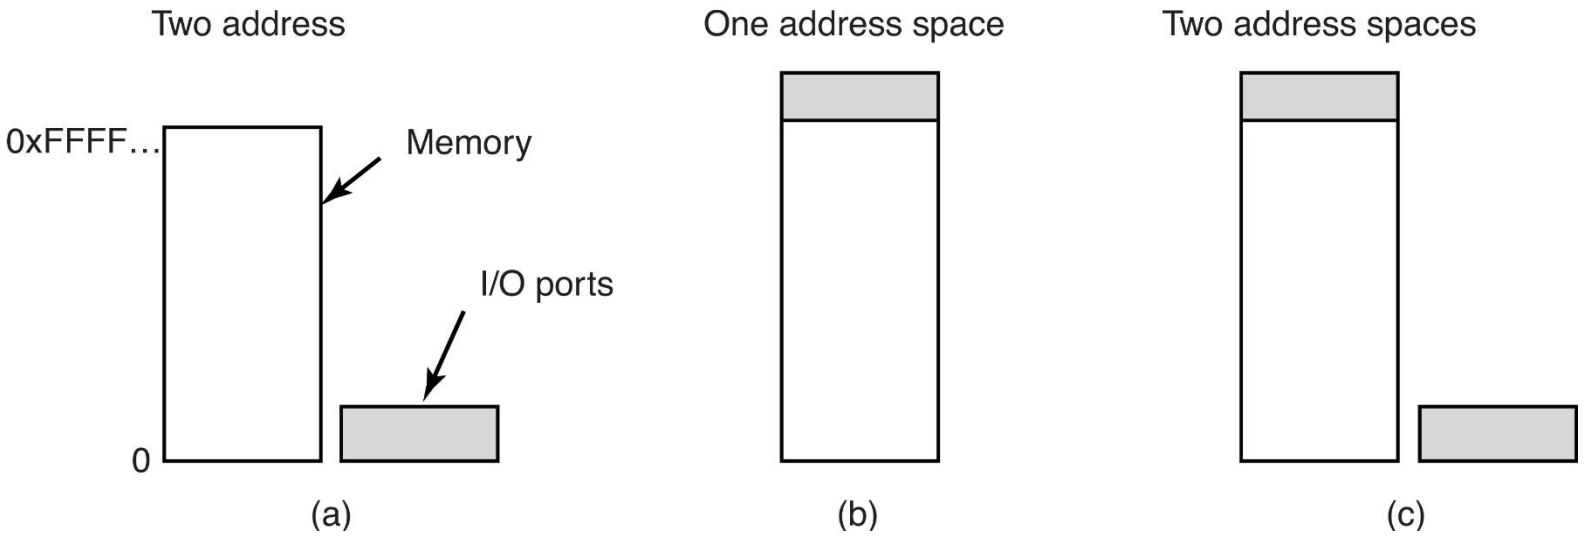

I/O Ports Memory-Mapped Hybrid

- Access controller registers [Figure 5-2 in Tanenbaum & Bos, 2014]

Strength and Weakness

- Strength of memory mapped I/O
  - Easier to program
  - Easier to protect
  - Faster to access
- Weakness (two addresses logically identical, but physically different)
  - More complex to design cache
  - More complex to design bus

Questions?

- Access devices controller registers
  - I/O ports
  - Memory-mapped I/O
  - Hybrid

I/O Schemes

- Busy waiting (polling)
  - while (busy) wait; do I/O;
- Interrupted I/O
  - do something else; when (interrupted) do I/O;
- Direct memory access (DMA)
  - initialize DMA; do something else; I/O done when interrupted;

Implementing Busy Waiting

- Illustrate it with writing a byte
  - Host

1. do
2. read the busy-bit in the device status register
3. while (busy)
4. set the write-bit in the control register
5. write a byte into the data-out register
6. set the command-ready bit in the control register

  - Device Controller

1. do
2. read the command-ready bit
3. while (not set)
4. set the busy bit
5. read the byte in the data-out register
6. write the byte to the device
7. if (success) clear the command-ready bit and the busy bit
8. else set the error bit

Interrupted I/O

- Conduct I/O in an asynchronous fashion

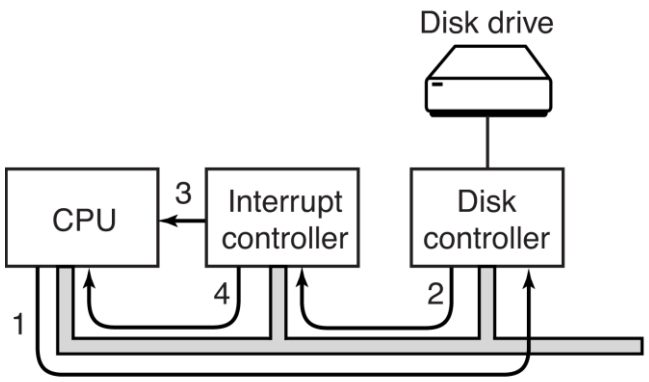

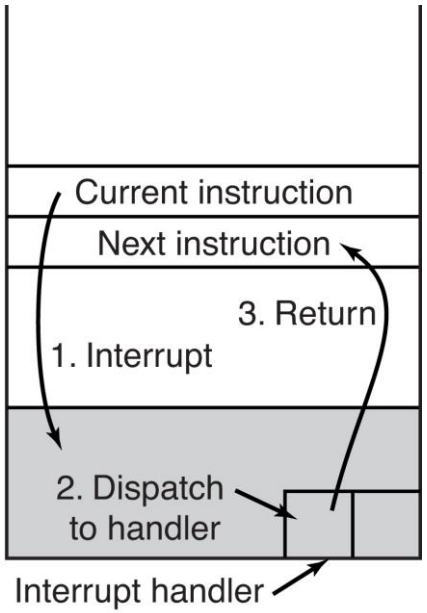

- Interrupted I/O[Figure 1-11 in Tanenbaum & Bos, 2014]

Interrupted I/O Cycle

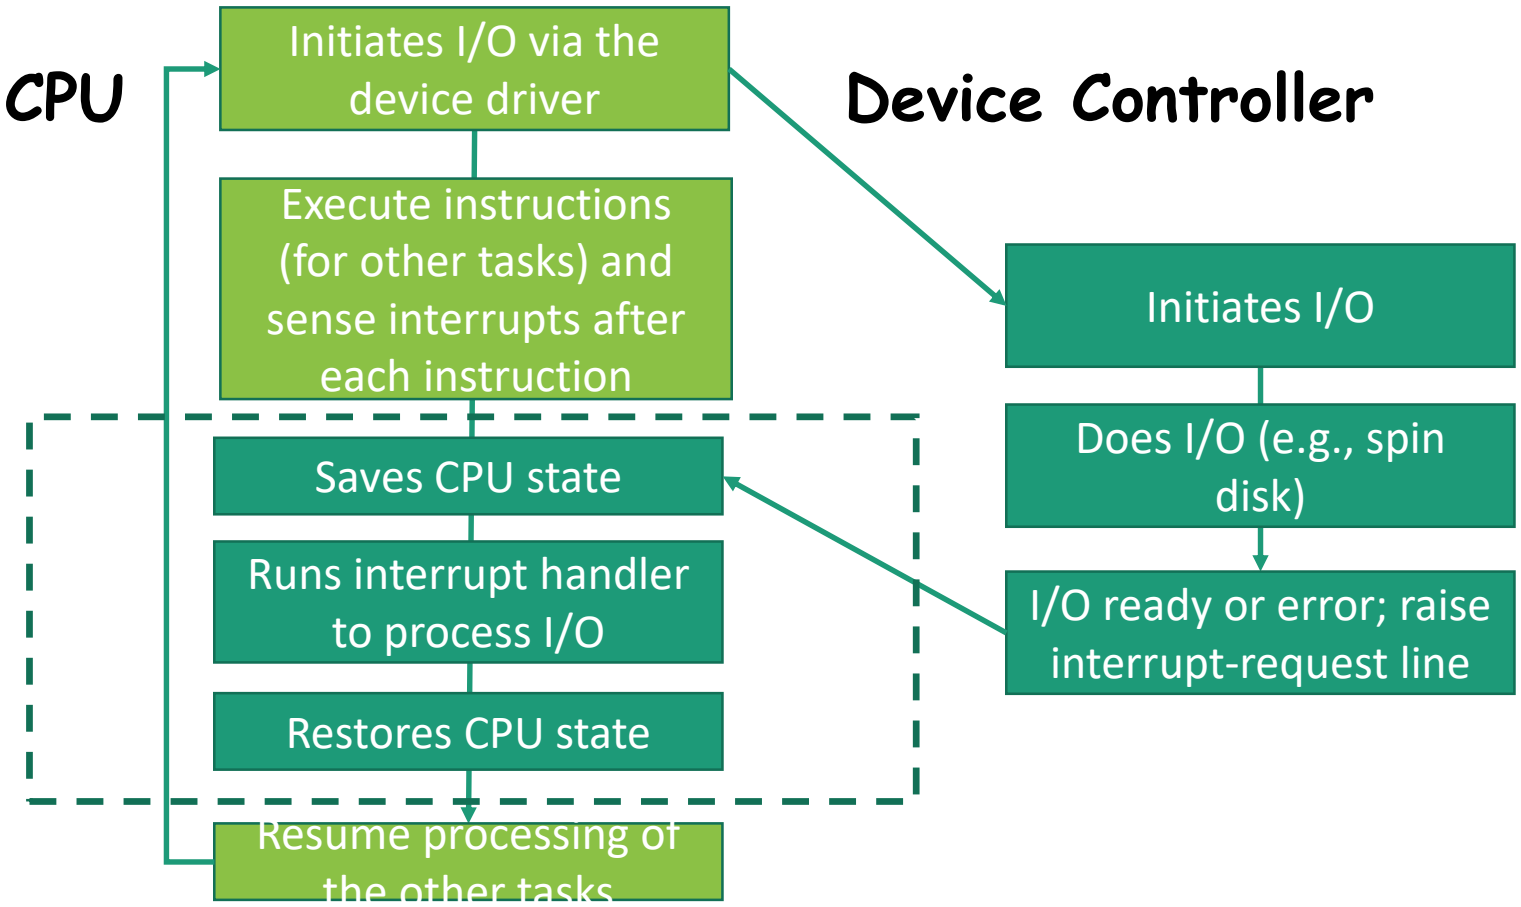

Direct Memory Access

- Aided by a special purpose processor called direct-memory-access (DMA) controller
  - CPU writes a DMA command block into memory
    - Pointer to the source of transfer
    - Pointer to the destination of transfer
    - A count of the number of bytes to be transferred
  - CPU writes the address of this block to the DMA controller
  - The DMA controller does I/O by directly access devices and system bus
  - CPU is interrupted when the DMA controller completes the transfer or encounters an error

Operating of a DMA Transfer

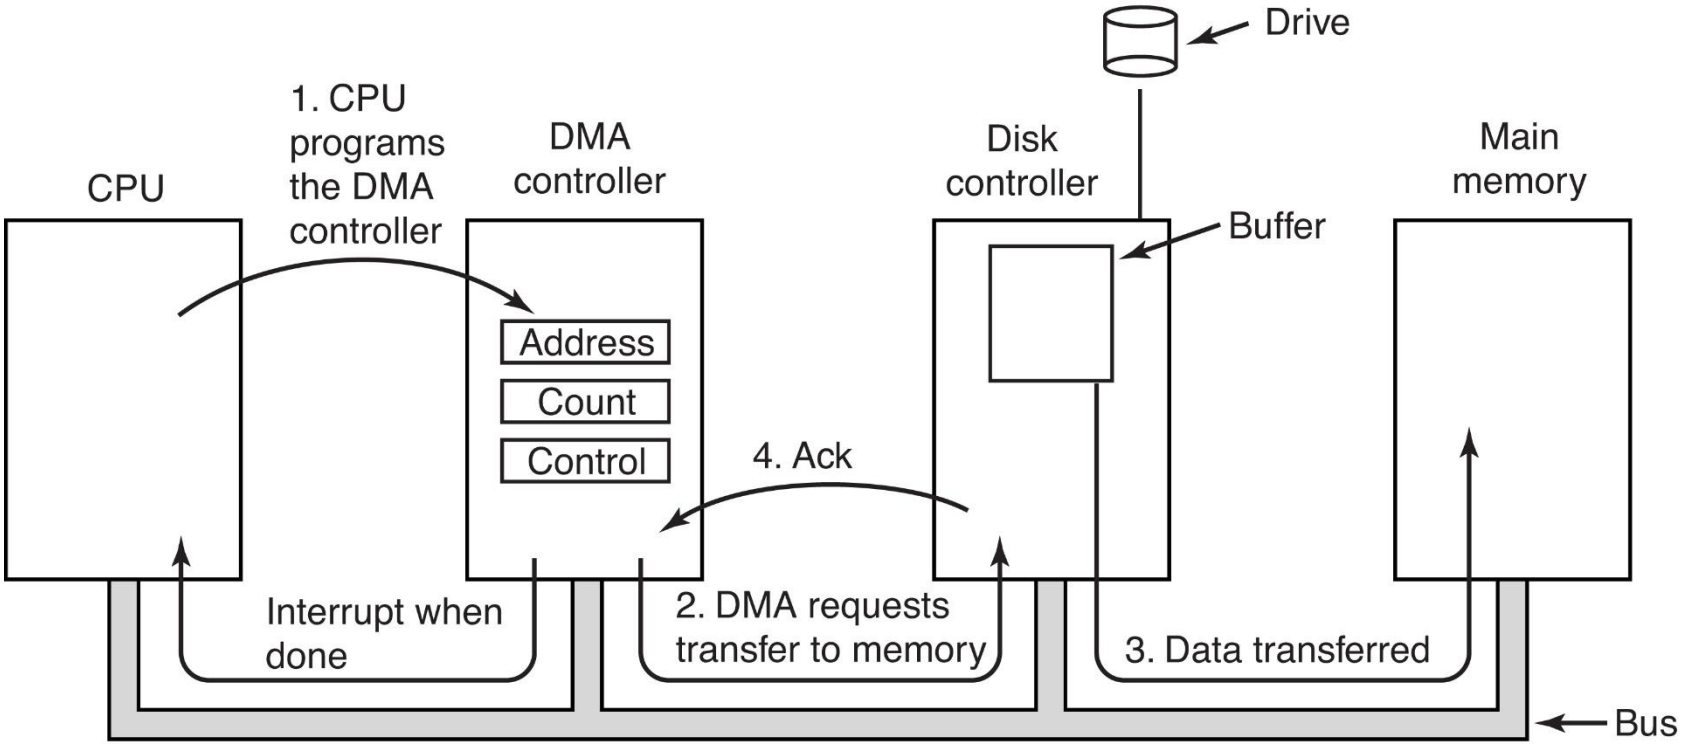

[Figure 5-4 in Tanenbaum & Bos, 2014]

Questions?

- I/O schemes
  - Busy waiting
  - Interrupted I/O
  - Direct-memory access

Interrupts

- CPU senses its interrupt-request line after each instruction
- When it is "lit", CPU saves the current state
  - Example: push registers PSW and PC to the stack
- CPU jumps to the interrupt-handler routine at a fixed address in the memory
- Interrupt-handler routine completes its task

and restore the CPU state

  - Pop the registers from the stack

How an Interrupt Happens?

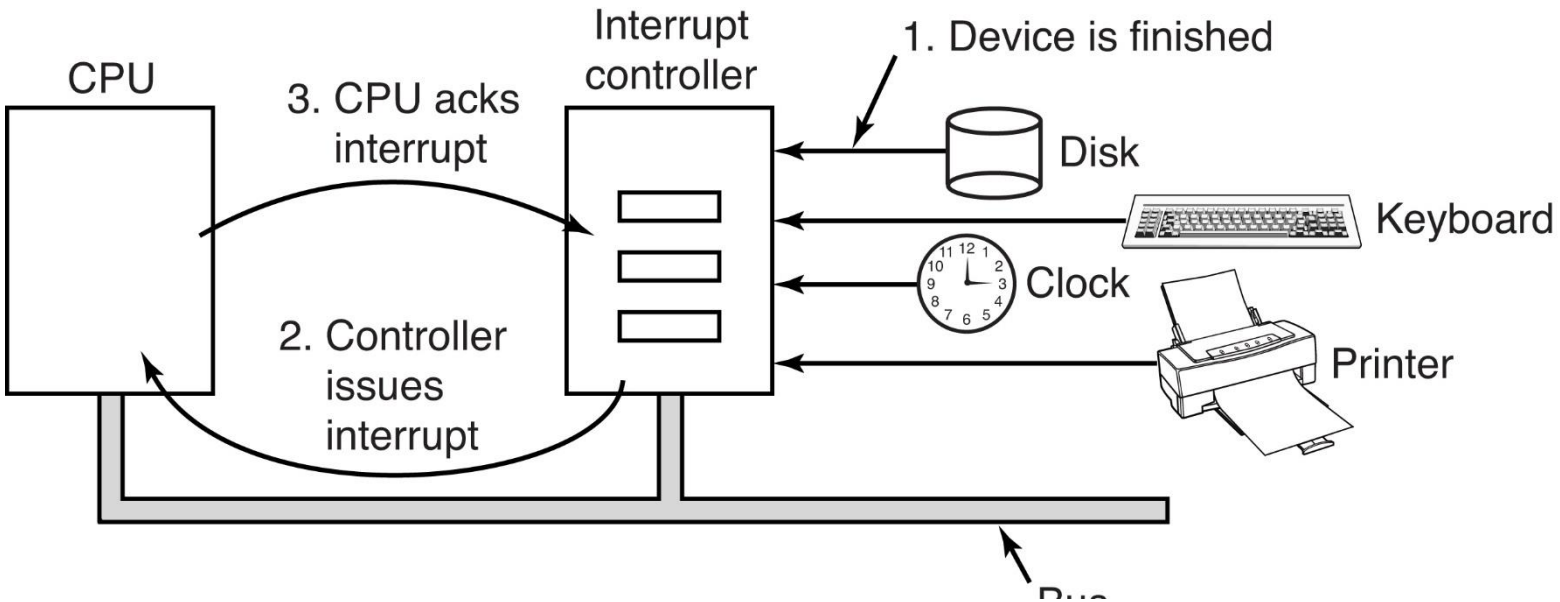

- [Figure 5-5 in Tanenbaum & Bos, 2014]

Interrupt Vectors

- Address to interrupt routines
  - Some PC event/interrupt-vector numbering

Vector Number	Description
0	Divide Error
1	Debug Exception
…	
6	Invalid Opcode
…	
32-255	Maskable Interrupts (deviced generated)

Design Consideration: Interrupts

- Maskable and nonmaskable interrupts
- Interrupt priorities and interrupt chaining
- Exceptions and software interrupts (traps)
- Precise and imprecise interrupts

Questions?

- More about interrupts

I/O Software

- Goals
- Programmed I/O
- Software layers
- Device drivers

Goals

- Device independence
- Uniform naming
- Error handling

Type of Devices

- Need to understand general characteristics to achieve device independent
- A couple of dimensions
  - Size of transfer: Character-stream or block
  - Access order: sequential or random access
  - Predictability and responsiveness: Synchronous and asynchronous
  - Shared or dedicated
  - Speed of operation, e.g., latency, seek time, transfer rate ec
  - Read-write, read only, or write only

Block Devices

- Naming
  - Examples on Linux
    - by label, by uuid, by id, and by path
    - Running examples
      - lsblk –f
      - ls /dev/disk/
- Read and write a block a time
- Essential behavior
  - read(), write()
  - For random-access block devices
    - seek()

Character Devices

- Read and write a character a time
- Essential behavior
  - get(), put()

I/O Software Layers

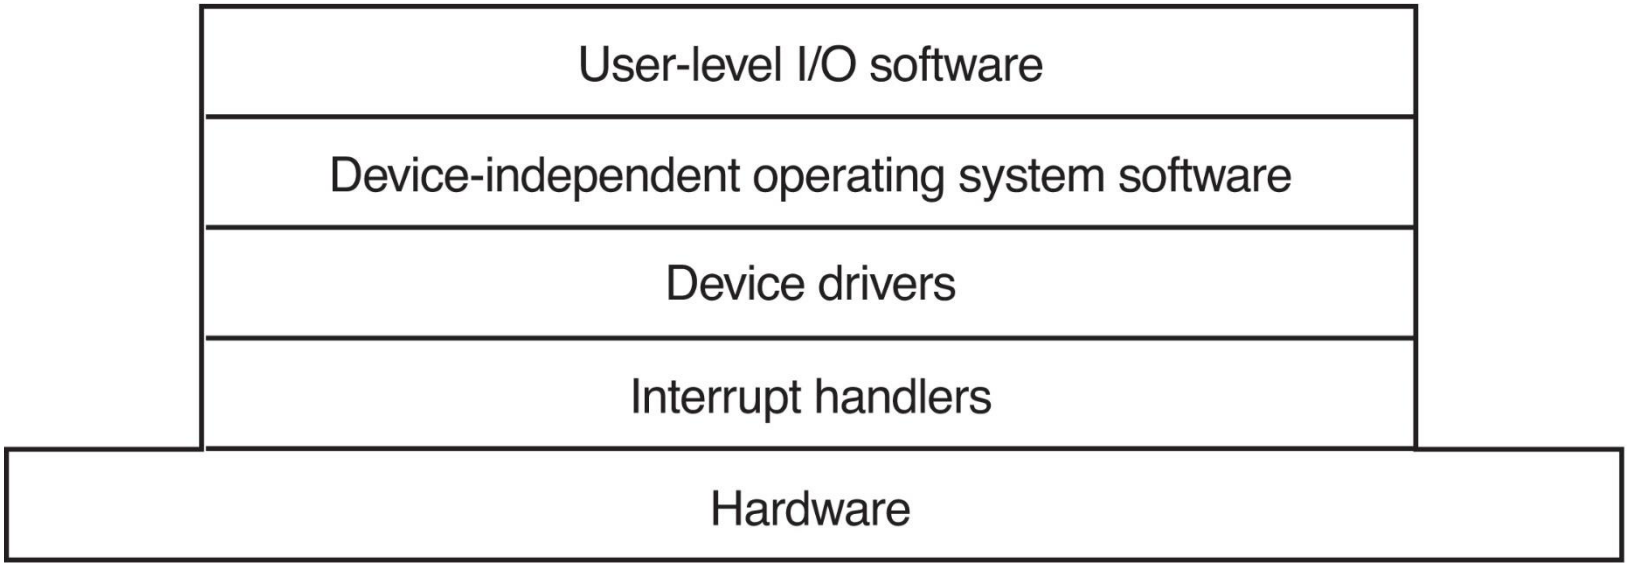

- [Figure 5-11 in Tanenbaum & Bos, 2014]

Questions

- Goals of I/O software and a few examples
- I/O software layers

Experiment Environment

- Debian Linux as a Virtual machine
- Need to install a few packages to make it useable
- Use the LKMPG as a guide
  - https://github.com/bashrc/LKMPG

Assignment

- Project 1