CS1101: Lecture 17
Computer Systems Organization: Input/Output

Dr. Barry O'Sullivan
b.osullivan@cs.ucc.ie

Course Homepage
http://www.cs.ucc.ie/~osullb/cs1101

Lecture Outline

- Buses
  - A simple personal computer
  - Direct Memory Access
  - Interrupts
  - Bus Arbiter & Cycle Stealing
  - ISA, EISA & PCI Buses
- Terminals
- Keyboards
- Graphics
  - CRT Monitors
  - Flat Panel Displays
  - Types of Terminal
  - Screen Output on the PC
  - Character-Map Terminals
  - Bit-map Terminals
  - Video Memory & Bit-map Terminals
- **Reading**: Tanenbaum, Chapter 2, Section 4

Buses

- **Motherboard** – a large printed circuit board.
- The motherboard contains the CPU chip, some slots and various support chips.
- It also contains a **bus** etched along its length, and sockets into which the edge connectors of I/O boards can be inserted.

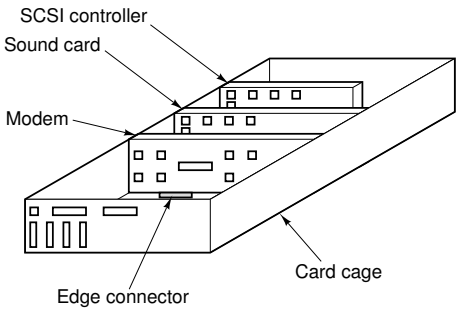

Figure 2-28. Physical structure of a personal computer.

A simple personal computer

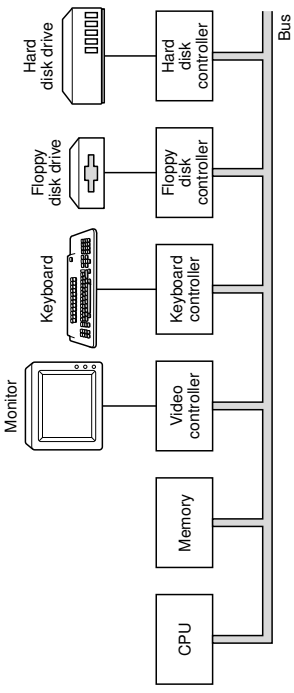

Figure 2-29. Logical structure of a simple personal computer.

Direct Memory Access

- Each I/0 device consists of two parts:
- One containing most of the electronics, called the **controller** and one containing the I/0 hardware device, such as a hard-disk.
- The job of a controller is to control its I/0 device and handle bus access for it.
- A controller that reads or writes data to or from memory without CPU intervention is said to be performing **Direct Memory Access**, better known by its acronym **DMA**.

Interrupts

- When the transfer is completed, the controller normally causes an **interrupt**
- This forces the CPU to suspend running its current program and start running a special procedure, called an **interrupt handler**
- This handler checks for errors, takes any special action needed, and inform the operating system that the I/0 is now finished.
- When the interrupt handler is finished, the CPU continues with the program that was suspended when the interrupt occurred.

Bus Arbiter & Cycle Stealing

- The bus is not only used by the I/O controllers, but also by the CPU for fetching instructions and data.
- What happens if the CPU and an I/0 controller want to use the bus at the same time?
- The **bus arbiter** decides who goes next.
- I/O devices are given preference over the CPU
- When no I/0 is in progress the CPU can have all the bus cycles for itself to reference memory.
- However, when some I/O device is also running, that device will request and be granted the bus when it needs it.
- This process is called **cycle stealing**.

ISA, EISA & PCI Buses

- One of the earliest bus architectures in the IBM PC was the **ISA (Industry Standard Architecture)** bus.
- However, this bus was too slow for market developments.
- This situation led to other companies developing machines with multiple buses, one of which was the old ISA bus, or its backward-compatible successor, the **EISA (Extended ISA)** bus.
- The most popular of these now is the **PCI (Peripheral Component Interconnect)** bus.

The PCI Bus

- Designed by Intel.
- In configurations based on a PCI bus, the CPU talks to a memory controller over a dedicated high-speed connection and not over the PCI bus.
- However, highbandwidth (i.e., high data rate) peripherals, such as SCSI disks, can connect to the PCI bus directly.
- The PCI bus would have a bridge to the ISA bus, so that ISA controllers and their devices can still be used.
- A machine of this design would typically contain several empty PCI slots and several empty ISA slots, to allow customers to plug in both old ISA I/0 cards (usually for slow devices) and new PCI 1/0 cards (usually for fast devices).

A Typical PCI bus Machine

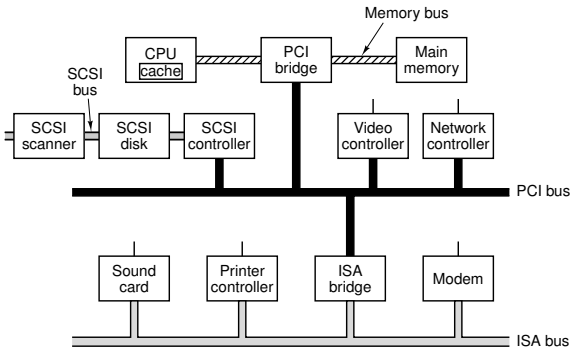

Figure 2-30. A typical modern PC with a PCI bus and an ISA bus. The modem and sound card are ISA devices; the SCSI controller is a PCI device.

Terminals

- Computer terminals consist of two parts:
- A keyboard and
- A monitor.
- In the mainframe world, these parts are often integrated into a single device and attached to the main computer by a serial line or over a telephone line.
- In the personal computer world, the keyboard and monitor are independent devices.

Keyboards

- Keyboards rely on various methods to physically register the pressing of a key – mechanical, electromagnetic etc.
- On PCs, when a key is depressed, an interrupt is generated and the keyboard interrupt handler (a piece of operating system software) is started.
- The interrupt handler reads a hardware register inside the keyboard controller to get the number of the key (1 through 102) that was just depressed.
- When a key is released, a second interrupt is caused.
- Thus if a user depresses the SHIFT key, then depresses and releases the M key, then releases the SHIFT key, the operating system can see that the user wants an upper case ”M” rather than a lower case ”m.”

CRT Monitors

- A monitor is a box containing a **CRT (Cathode Ray Tube)** and its power supplies.
- Within the CRT an electron beam is accelerated onto a phosphorescent screen near the front of the tube.
- Color monitors have three electron guns, one each for red, green and blue.

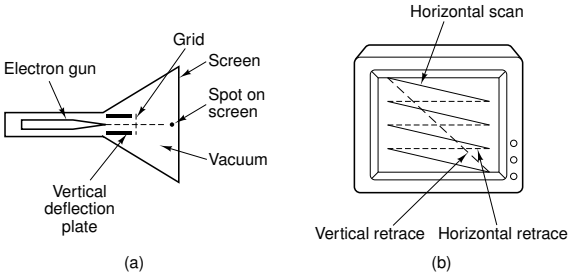

Figure 2-31. (a) Cross section of a CRT. (b) CRT scanning pattern.

Flat Panel Displays

- CRTs are far too bulky and heavy to be used in notebook computers.
- The most common screen for notebooks is based on the **LCD (Liquid Crystal Display)** technology.
- Two types are common:
  - Passive Matrix Displays
  - Active Matrix Displays

Flat Panel Displays

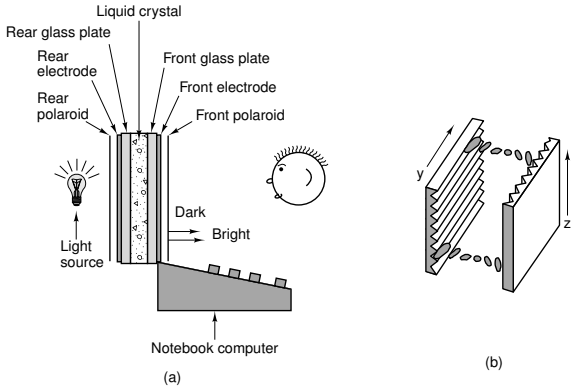

Figure 2-32. (a) The construction of an LCD screen. (b) The grooves on the rear and front plates are perpendicular to one another.

Types of Terminal

- Three kinds of terminals are in common use:
- character-map terminals
- bitmap terminals
- RS-232-C terminals
- They all can use any keyboard type
- The difference is in the way the computer communicates with them and how the output is handled.

Screen Output on the PC

- On a PC, there are two ways to organize the output to the screen: a character map and a bit map.
- On a serial communication board is a chunk of memory, called the **video memory**, as well as some electronics for accessing the bus and generating video signals.

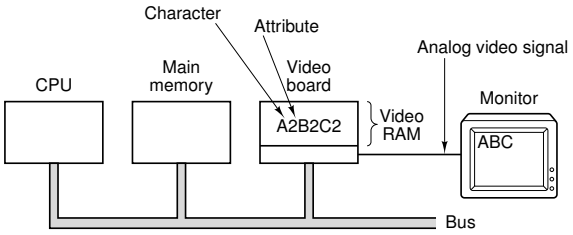

Figure 2-33. Terminal output on a personal computer.

Character-Map Terminals

- The screen is regarded as an array of characters
- To display characters, the CPU copies them to the video memory in alternate bytes.
- Associated with each character is an **attribute byte** that describes how that character is to be displayed.
- Attributes can include its color, intensity, whether it is blinking, etc.
- Thus a screen image of 25 x 80 characters requires 4000 bytes of video memory, 2000 for the characters and 2000 for the attributes.
- Most boards have more memory to hold multiple screen images.
- Example: MS-DOS based PC screen.

Bit-map Terminals

- The screen is regarded as an array of picture elements, called **pixels**.
- Each pixel is either on or off – it represents one bit of information.
- On PCs the screen may contain: 640 x 480 pixels, 800 x 600 pixels, 1024 x 768 pixels or 1280 x 960 pixels.
- The video RAM is just seen as a big bit array.
- Bit-map terminals are commonly used to support displays containing several windows.
- Examples: Windows, KDE on Linux.

Video Memory & Bit-map Terminals

- Although bit-map terminals are highly flexible, they have two major disadvantages:
  - RAM intensive
  - Performance issues.
- To get true color, 8 bits are needed for each of the three primary colors, or 3 bytes/pixel – thus a 1024 X 768 display, requires 2.3 MB of video RAM.
- *To display full-screen, full-color multimedia* on a 1024 x 768 display requires copying 2.3 MB of data to the video RAM for every frame.
- *To display full-motion video*, a rate of at least 25 frame/sec is needed, for a total data rate of 57.6 MB/sec.
- Thus, high-performance video cards on IBM PCs need to be PCI cards.