15.561
Information Technology Essentials

Session 2
Fundamentals of Computing II

Copyright © 2003 Thomas Malone, Chris Dellarocas

Acknowledgments:.
Adapted from slides by Chris Dellarocas, U. Md..

Outline: Fundamentals of Computing

- Computer architecture
  - Hardware Components
    - » CPU, Memory, I/O, Buses
  - Understanding PC specs

- Operating Systems
  - What is an OS?
  - OS Functions
    - » Multitasking, Virtual Memory, File Systems, Window systems
  - Microcomputer operating systems

The Little Man Computer

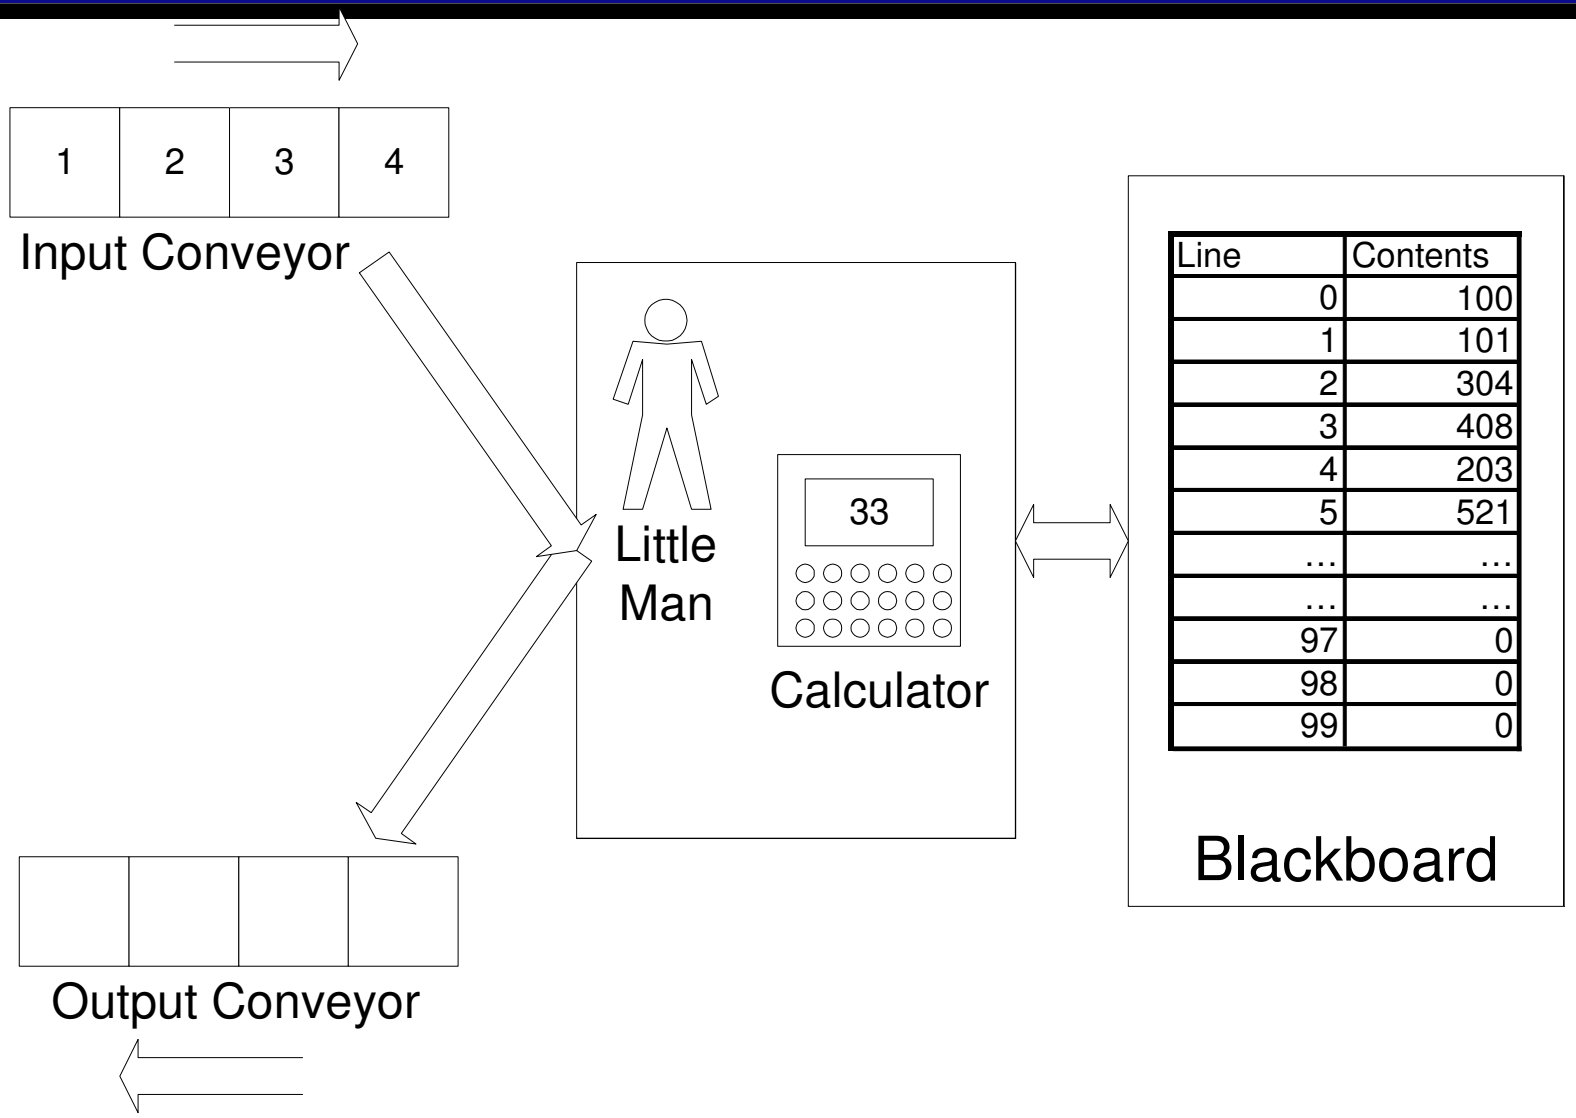

BASIC FACTS TO ASK ABOUT ANY COMPUTER

LMC ANSWERS

1. **MEMORY**
   - (A) BASIC UNIT 3 DECIMAL DIGIT NUMBER
   - (B) MAXIMUM SIZE 100 LOCATIONS

2. **REGISTERS**
   - (A) HOW MANY 1
   - (B) NUMBERS 3 DECIMAL DIGIT NUMBER

3. **INSTRUCTIONS**
   - (A) NUMBER 7 INSTRUCTIONS

A "Real" Computer

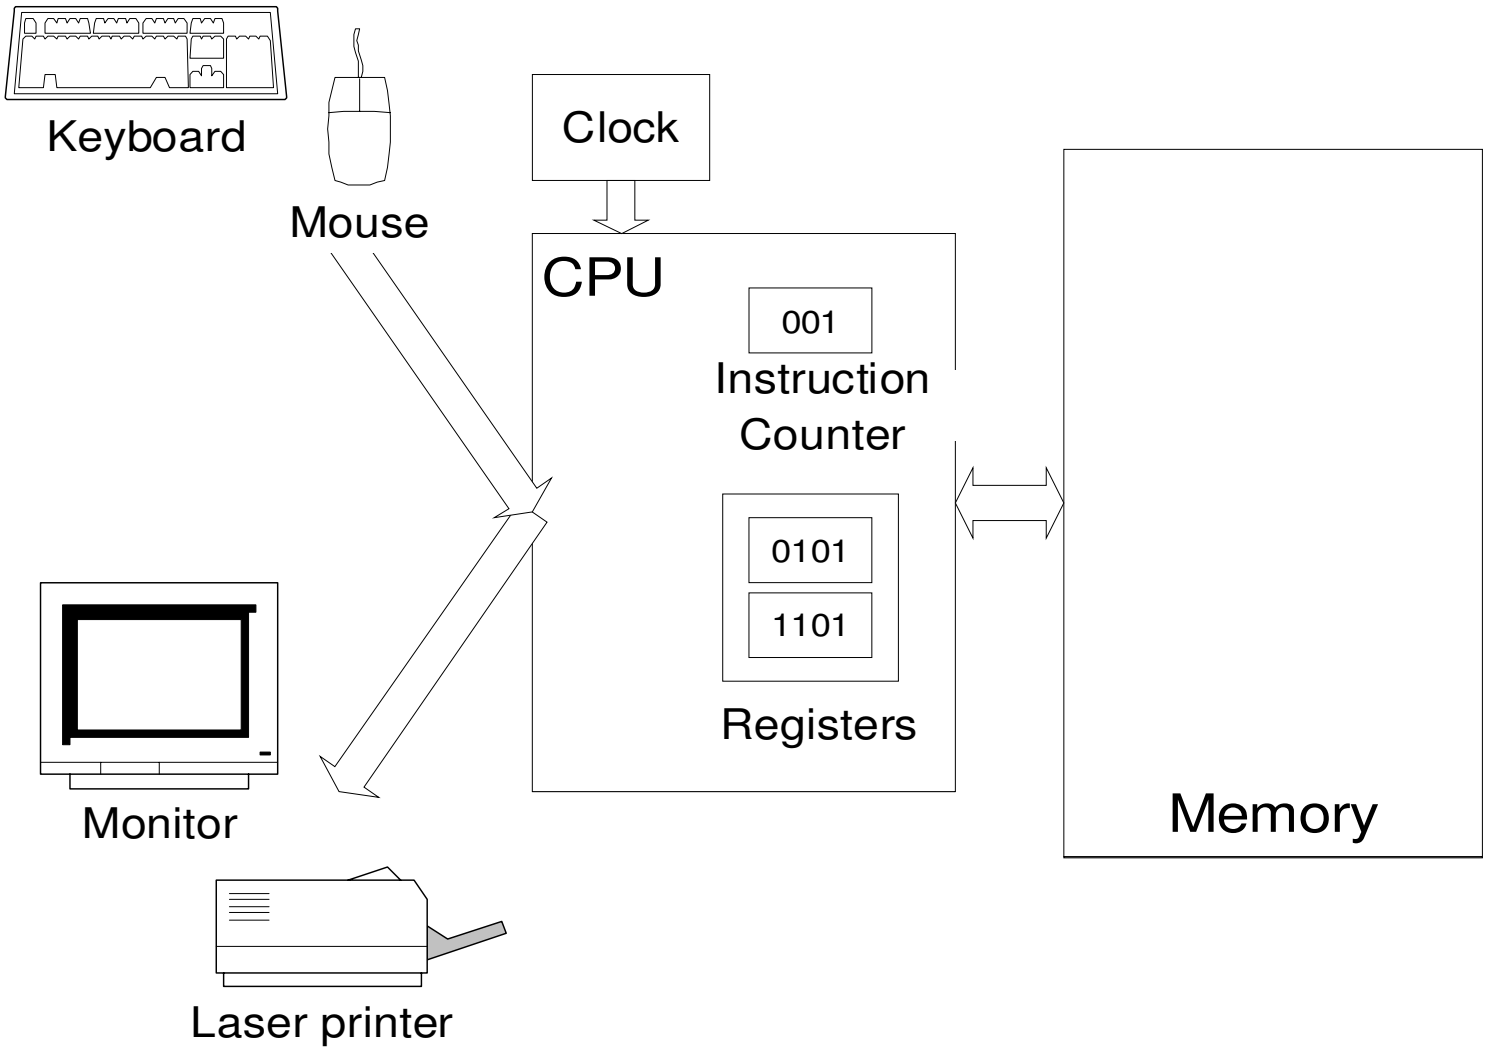

INTEL PENTIUM 4 ANSWERS

1. MEMORY

(A) INDUSTRY BASIC UNIT — 8 BINARY DIGITS (BITS) = 1 BYTE

(B) BASIC UNIT — 32 BITS = 4 BYTES

(C) TYPICAL SIZE — MEMORY RAM: 128 MB – 1GB

2. REGISTERS

(A) HOW MANY — ABOUT 50 REGISTERS

(B) NUMBERS — VARIOUS TYPES

3. INSTRUCTIONS

(A) NUMBER — ABOUT 500

SM

Binary Computers

- Real computers don't store and calculate with 3-digit decimal numbers
- A bit (binary digit) distinguishes between two states
  - TRUE and FALSE
  - 1 and 0
- Bits are easier to implement in machines
  - Light bulb on or off
  - High vs. low voltage (on wires)
  - Magnetized or not (computer hard disks, floppies, tapes)
  - Pit or no pit detected by a laser (compact discs)

Interpretation of a decimal number

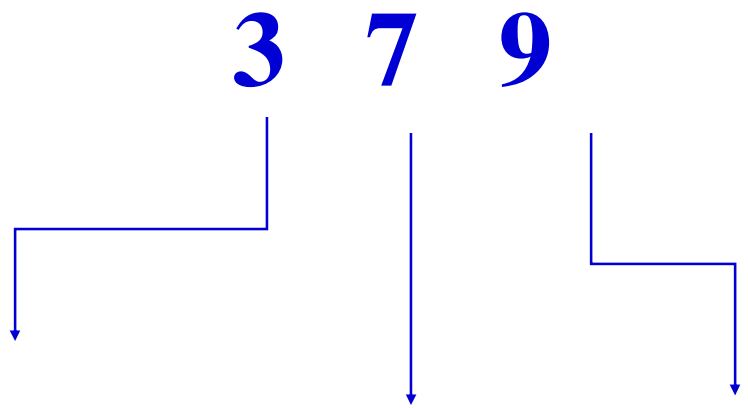

$3x100 + 7x10 + 9x1$

$3x10^2 + 7x10^1 + 9x10^0$

Interpretation of a binary number

1 1 0 1 1

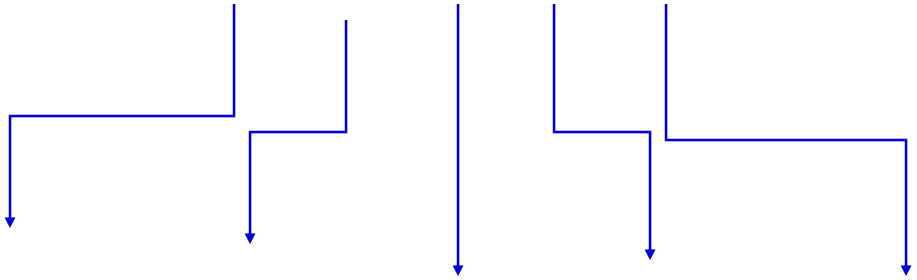

$$1x2^4 + 1x2^3 + 0x2^2 + 1x2^1 + 1x2^0$$

1x16 + 1x8 + 0x4 + 1x2 + 1x1

27

The CPU

- CPU = Central Processing Unit
- Internal clock ticks very fast (e.g., 1.6 GHz = 1.6 billion ticks per second)
  - activities are synchronized to start on a clock tick
  - some activities take more than one clock tick
- Instruction execution is automatic
  - (tick) find memory address of next instruction
  - (tick) retrieve instruction from memory
  - (tick) decode the instruction
  - (tick) fetch argument from memory if necessary
  - (tick) execute instruction
  - (tick) store result in memory if necessary

CPU and Memory Interaction

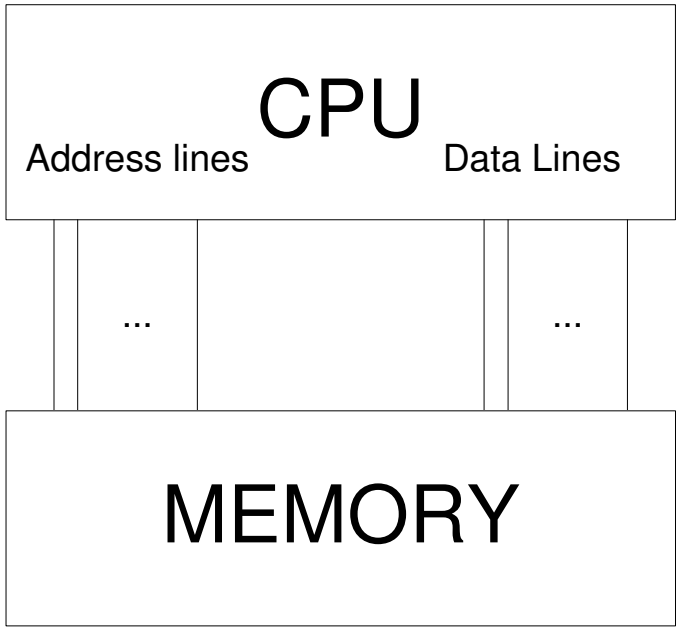

CPU Issues an Address

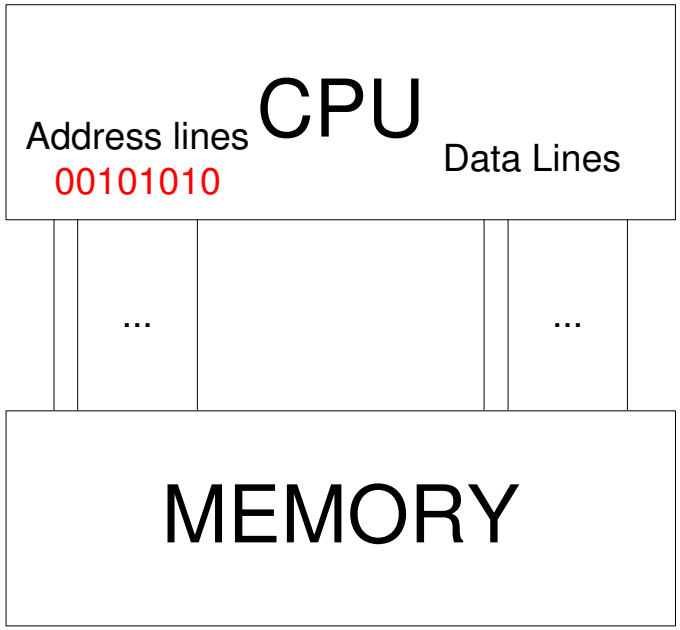

CPU: I need the contents of memory location 50 !

Memory makes the data available

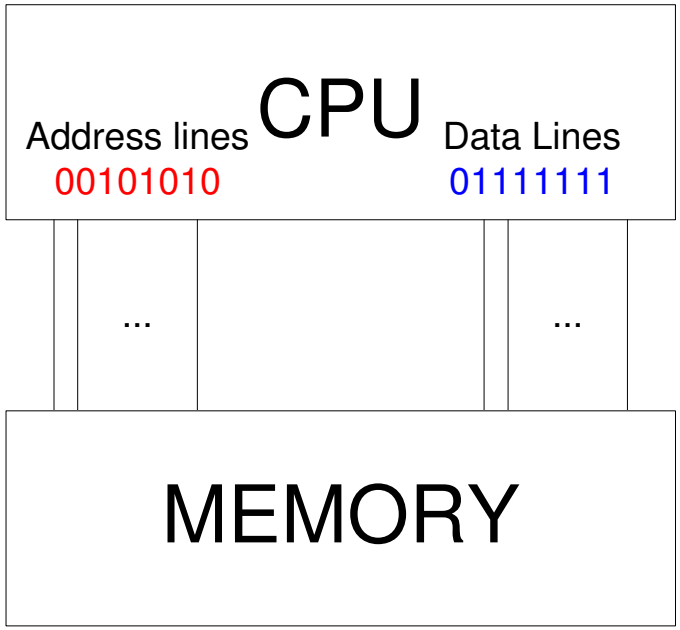

Memory: Location 50 contains the number 127

CPU Characteristics

- Family: Determines the set of instructions it understands
  - » Intel 80386, 80486, Pentium, Pentium II,…
  - » Motorola: 68030, 68040
- Clock Speed
  - » Pentium: 500 MHz – 2.2 GHz
- Data bus width: Size of data that can be manipulated at one time
  - » 80486: 32 bits, Pentium: 64 bits
- Address bus width: Limits the amount of memory that can be installed in the computer
  - » LMC: 3 decimal digits. Locations _____
  - » Pentium: 32 bits. Locations _____
  - » Itanium: 64 bits. Locations _____

Expressing Memory Capacity

- Measured in bytes (=groups of 8 bits)
- Each byte can store a binary number from 00000000 to 11111111 ($255 = 2^8-1$)
- More generally: n binary digits can store numbers from 0 to 2^n-1
- Frequently used multiples:
  - Kilobyte (KB) = 1,024 (2^{10}) bytes
  - Megabyte (MB) = 1,024 KB = 1,048,576 (2^{20}) bytes
  - Gigabyte (GB) = 1,024 MB ~ 1 billion (2^{30}) bytes

Semiconductor Memory

- RAM (Random Access Memory)
  - Can access any location equally fast
  - Loses contents without power
  - Two main types
    - » Static (SRAM): Faster, expensive
    - » Dynamic (DRAM): Slower, cheaper, consumes less power and space
- ROM (Read Only Memory)
  - Retains memory even without power
  - Useful to store programs executed upon system start-up (e.g. BIOS)

Hard Disks and Floppies

- Slower than main memory
- Bits stored as magnetic field of different polarity
- Magnetized surface of disk rotates under a magnetized head (read/write mechanism)
- Disk divided into tracks, each at different radius from center
- Tracks are divided into sectors

Hard Disk Geometry

- Head moves back and forth
- To read/write some data:

1. Head moves over desired track
2. System waits until desired sector passes under head
3. Data is read/written

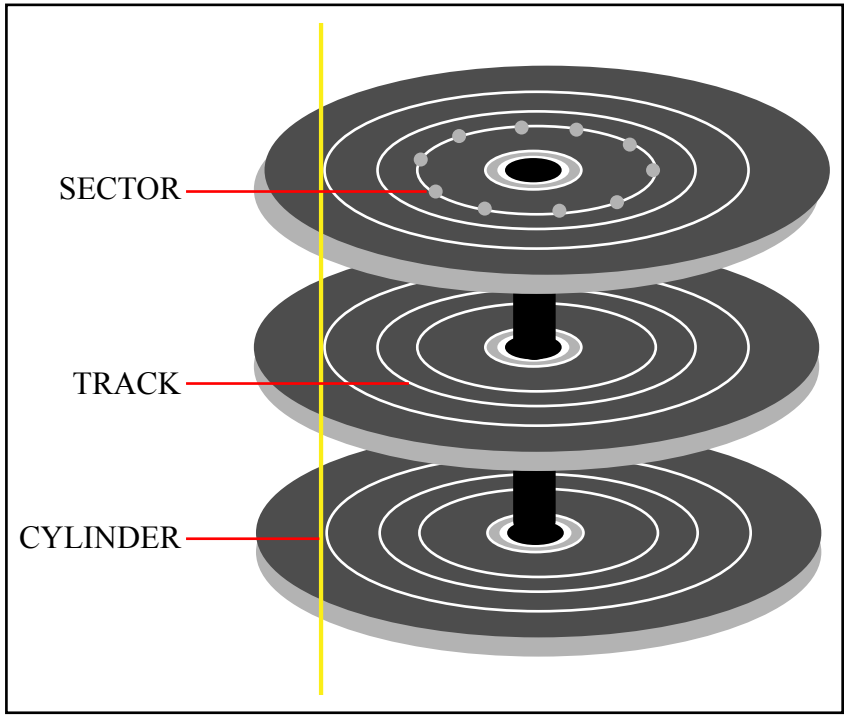

Figure by MIT OCW.

CD-ROMs

- Slower than hard disks
- Data is encoded by burning miniature “pits” on a photoreflective surface; read by laser
- CDs can hold up to 650MB of data.
- CD-ROM drive maximum transfer speed is expressed in multiples of 150KB/sec
  - 4X drive --> 600KB/sec
  - 20X drive --> 3000KB/sec

DVD (Digital Video Disk)

- **New, improved CD-ROM**
  - **smaller, denser "pits"**
  - **two layers of "pits" recorded on the same disk**

- **DVDs can hold up to 17GB of data.**

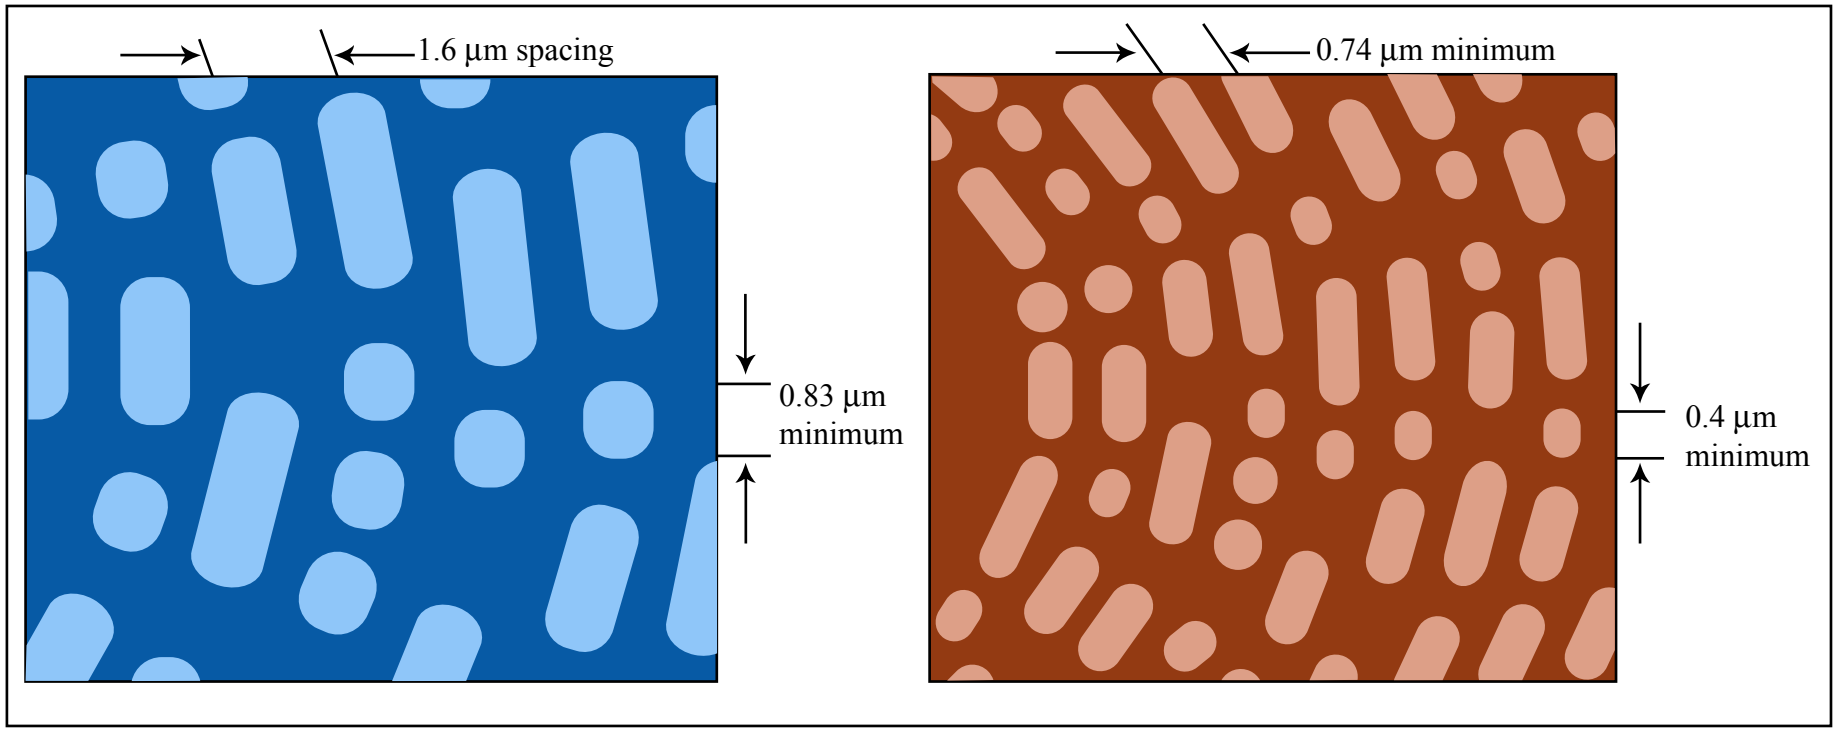

Figures by MIT OCW.

Keychain drives

- Hold 16 MB – 2 GB
- Attach to USB (Universal Serial Bus) port
- Usually use “flash memory”
  - A special kind of ROM that can be rapidly erased and re-recorded

I/O Devices

- Input
  - Keyboard
  - Mouse
  - Hard Disk
  - Floppy Disk
  - …
- Output
  - Printer
  - Screen
  - Speakers
  - …

Computer Displays

- Computer screen divided into small dots (pixels)
- Each pixel can be displayed in a different color
- Screen resolution: Number of pixels per screen
  - 640x480
  - 1024x768
- Color information for each pixel stored in memory, read and converted to video signal 60 times per second
  - To store information for a 1024x768 screen with 256 possible colors for each pixel we need _____ bytes

Buses: Connecting I/O to CPU

- One set of wires connect all devices and CPU
  - Transport of information is shared (public)
  - Hence called a bus (public transportation)
- Nearly all computers use a bus to connect CPU and I/O Devices
- Buses allow easy addition/replacement of I/O Devices
  - Modern PCs come equipped with expansion slots, directly connected to the bus
  - I/O Device controllers implemented as expansion cards
  - Examples: ISA, PCI, PCMCIA, IEEE 1394 (FireWire)

A simple bus architecture

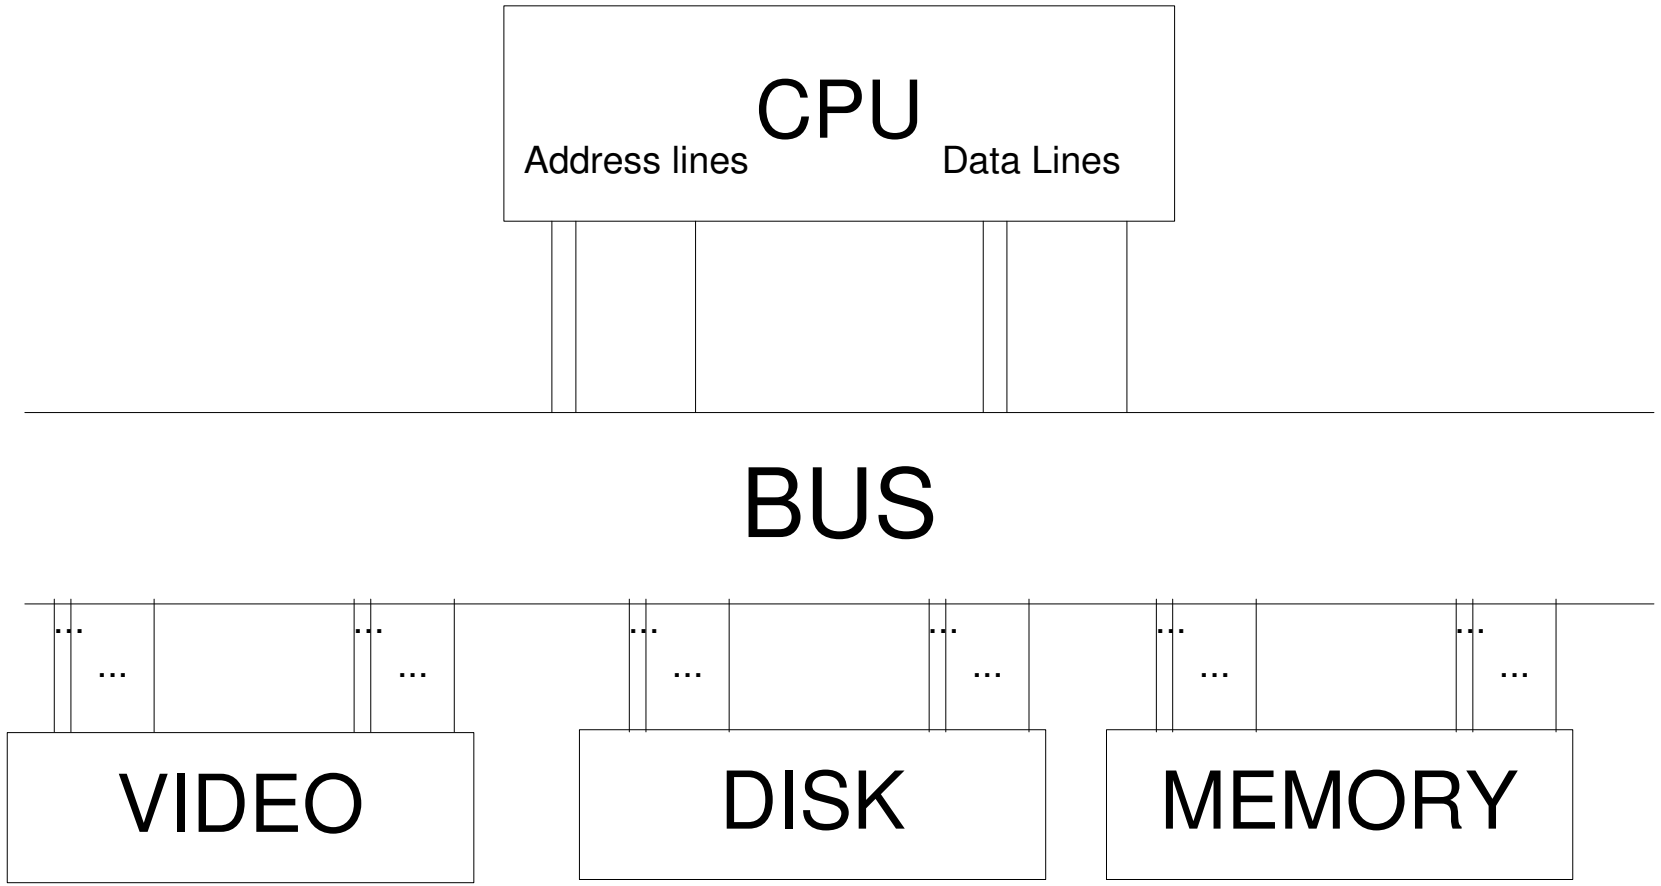

Cache Memory: Motivation

- Cheap main memory is slower than CPU
  - Example: Pentium PCs
    - » CPU 2ns (500MHZ)
    - » Main memory (100MHZ SDRAM) 10ns
  - Instructions that access main memory take many more clock ticks than those that don't

- Solution:
  - automatically keep copies of most frequently used memory locations in fast (but expensive) memory = cache memory

A modern PC architecture (simplified)

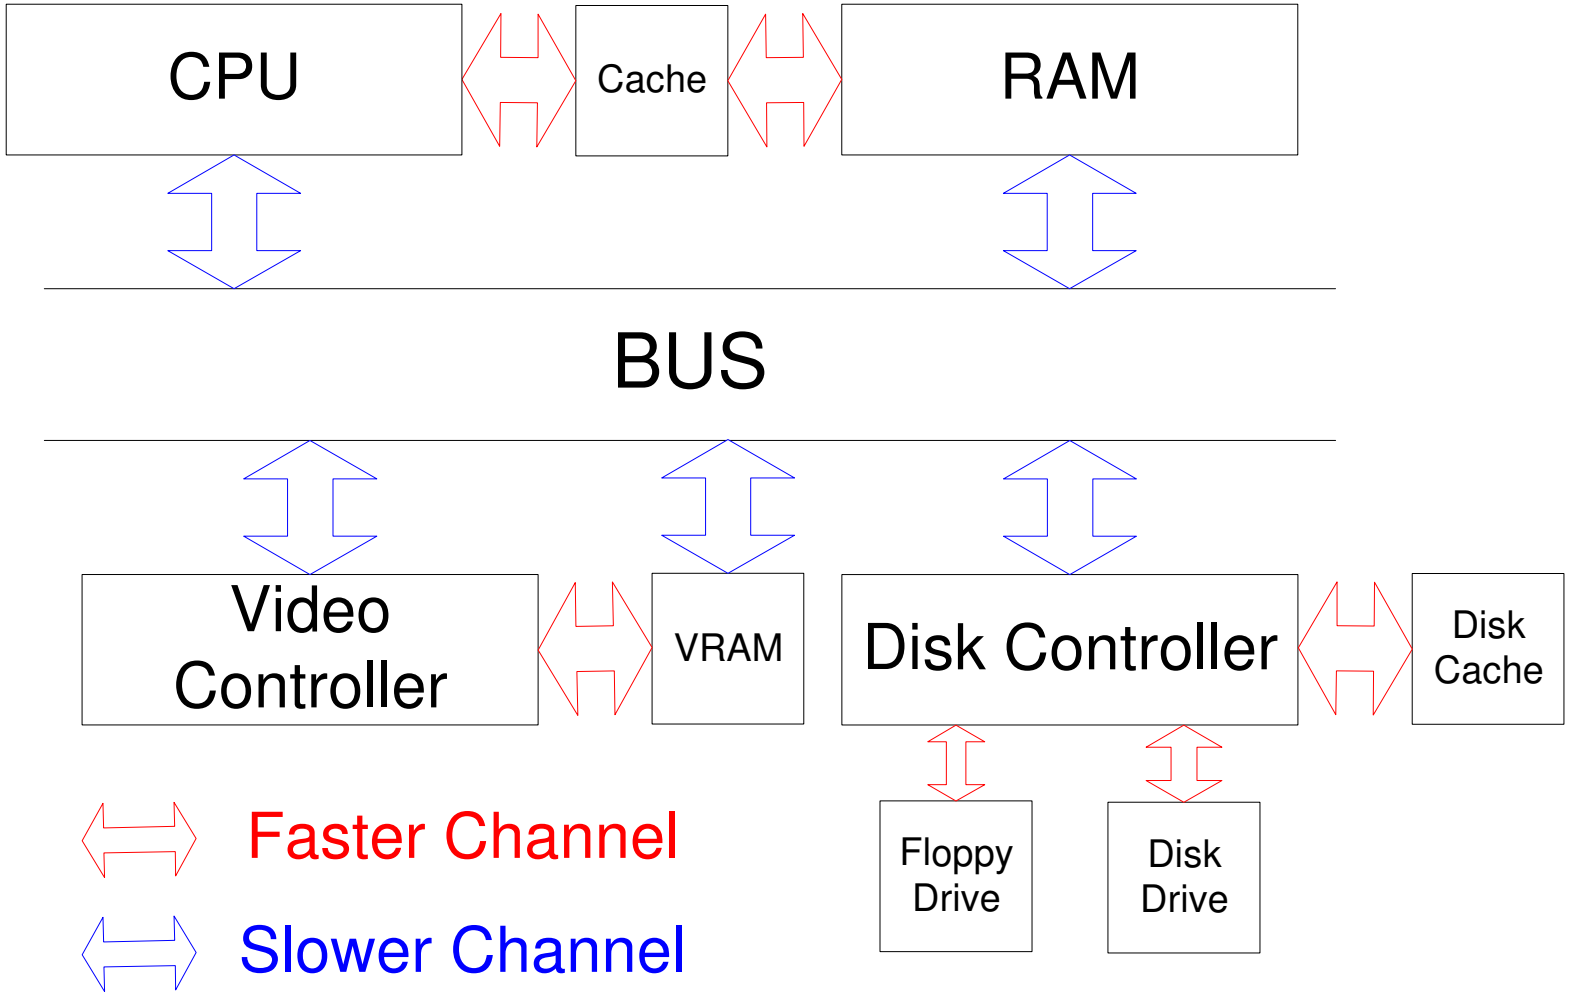

Summary: A modern PC

- Processor: Pentium (500 MHz – 3.6 GHz)
- Main Memory: 64 MB - 4 GB
- Floppy Drive: 1.44MB (3.5-inch disks)
- Hard Drive: 10 - 500GB
- Graphics: 640x480 – 2048x1536, 256 to 16 million colors
- Video Memory: 32 - 256MB

Types of software

- System software
  - Operating systems
  - Programming languages
  - Database systems
- Application software
  - General office tasks (word processing, etc.)
  - Accounting
  - Design
  - Factory automation
  - …

Operating systems - Examples

- DOS
- Windows (95, 98, NT, 2000, XP)
- Mac OS X
- Unix
  - Linux
- ...

Operating system

- Allocates and assigns:
  - memory
    - » e.g., file system, virtual memory
  - processor time
    - » e.g., multitasking, multiprocessing
  - input-output devices
    - » e.g., printer, keyboard, etc.
- May also provide other capabilities useful to many users or programs
  - Graphical User Interface (GUI) capabilities
  - Fonts, network protocols, ...
  - Web browser?

Operating system as magician: The four illusions

- Many separate computers, one for each process
  - "Multitasking"
- Large memory
  - "Virtual memory"
- Disks and other secondary storage are organized as collections of files
  - "File systems"
- Windows and menus
  - "Graphical User Interface (GUI)"

Illusion #1: Multitasking

- **Reality:**
  - One CPU
  - One instruction at a time
- **Illusion:**
  - Several application programs executing concurrently
- **Implementation:**
  - Operating system divides CPU time among application programs (time sharing)
    - » each program "thinks" it is the only one running
    - » OS copies Instruction Pointer and Registers back and forth as each program takes its turn

Multitasking issues

- How is control passed to next task?
  - Cooperative multitasking (Windows, Windows 95)
    - » Application explicitly passes control back to OS
    - » What if application never passes control back?
  - Preemptive multitasking (Unix, NT, XP)
    - » Operating system interrupts application when I/O requested or when preset time limit has passed
- Can one task access the memory of another one?
  - Preventing this is called “memory management”
  - Done by Unix, NT, XP, Mac OS X (Not by older versions of Mac OS and Windows)

Illusion #2: Virtual Memory

- Some data is not used for a long time
  - Why keep it all in memory?
- Copy a unit of data (called a “page”) to hard disk and use memory for other data
- Copy pages back from hard disk to main memory as they're needed
- Process (and its programmer) not aware that main memory is too small (the big memory illusion)
  - It asks for a main memory location (Page #, offset on page)
  - OS has to get that page into main memory if not already there

Illusion #3: File Systems

- **Reality:**
  - Disks are sets of tracks
  - Tracks are sets of sectors
  - Sectors can store fixed-sized byte blocks

- **Illusion:**
  - Disks are sets of directories
  - Directories contain other directories or files
  - Files are variable-size byte sequences
  - Directories and files have names

Illusion #4: Windows and Menus

- Reality: Screen is an array of pixels
- Illusion 1: Menus
  - Depending on where you click, different action happens
  - Technique: OS looks up location where mouse was clicked, executes appropriate action
- Illusion 2: Overlapping windows
  - A window may cover part or all of another
  - When a window is uncovered, its contents are redisplayed
  - Technique: OS saves bitmap of covered windows
    - » Application does not need to know how to redraw the contents of its window

Microcomputer Operating Systems

- DOS
  - text-based interface, no multitasking
- Windows
  - windows, cooperative multitasking
  - filenames restricted to 8 characters
  - bad memory management!
- Windows 98
  - large filenames
  - built-in networking capabilities
  - plug-and-play device configuration

Microcomputer Operating Systems (cont'd)

- Windows NT (including Windows 2000, Windows XP)
  - full multitasking
  - full memory management
- UNIX (including Linux)
  - great memory management, multitasking
  - complex, text-based interface
- Mac OS X
  - Based on Unix
  - Easy to use
  - can only run on Macintoshes

Selecting an Operating System

- Is our existing software compatible with the OS?
- Does the OS have a large base of compatible software?
- How reliable is the OS? Does it crash frequently?
- Is the OS available for a wide variety of hardware?
- How quickly does it run?
- How easy it is to learn and use?
- How easy is it to install and configure?
- How much does it cost?