

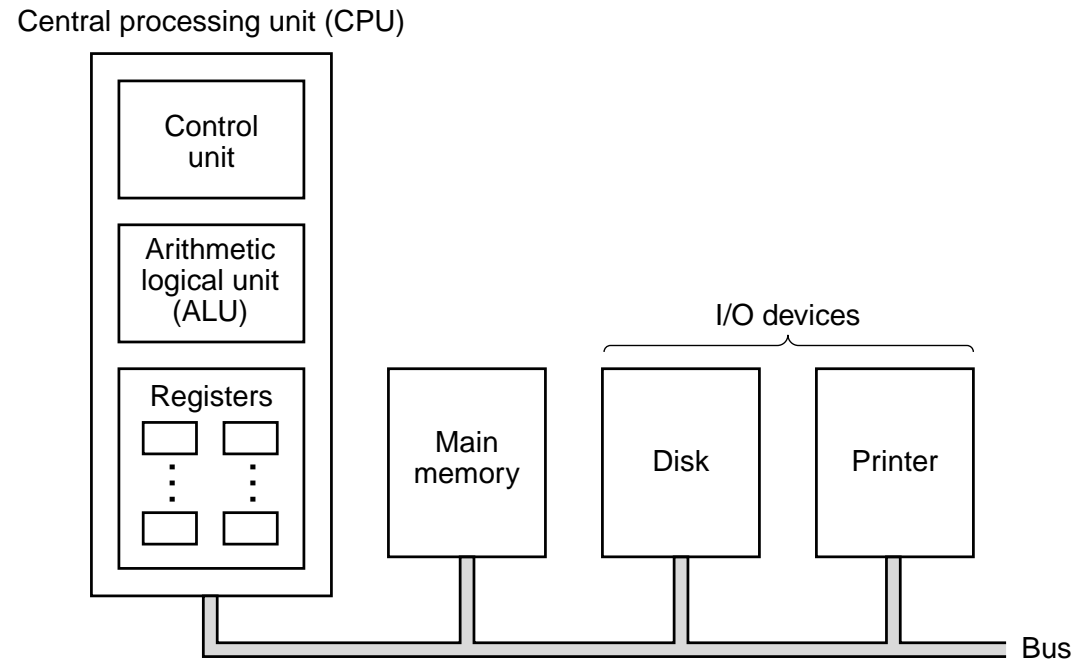
# Computer Systems Organization

Wolfgang Schreiner  
Research Institute for Symbolic Computation (RISC-Linz)  
Johannes Kepler University

Wolfgang.Schreiner@risc.uni-linz.ac.at  
<http://www.risc.uni-linz.ac.at/people/schreine>

# Computer Organization

Interconnected system of processor, memory, and I/O devices.



Later we will study these components on all levels.

## Example

A modern Personal Computer (PC) in 2001 may consist of

- a 933 MHz Pentium III processor,
- 256 MB RAM,
- a 30 GB hard disk,
- a keyboard and a mouse,
- a floppy disk drive,
- a 40x speed CD-ROM drive,
- a 19" monitor with  $1280 \times 1024$  pixels resolution,
- a 56 Kbit Modem,
- a 100 Mbit Ethernet card.

The last six items are I/O devices.

# Processor

## Processor

- Components are connected by a **bus**.
  - Collection of parallel wires.
  - Transmits address, data, and control signals.
- Central component is the **CPU**.
  - Central Processing Unit.
  - Executes program stored in main memory.
    - \* Fetches its instructions.
    - \* Examines them.
    - \* Executes them one after another.

The CPU is the “brain” of the computer.

## CPU Organization

CPU is composed of several distinct parts.

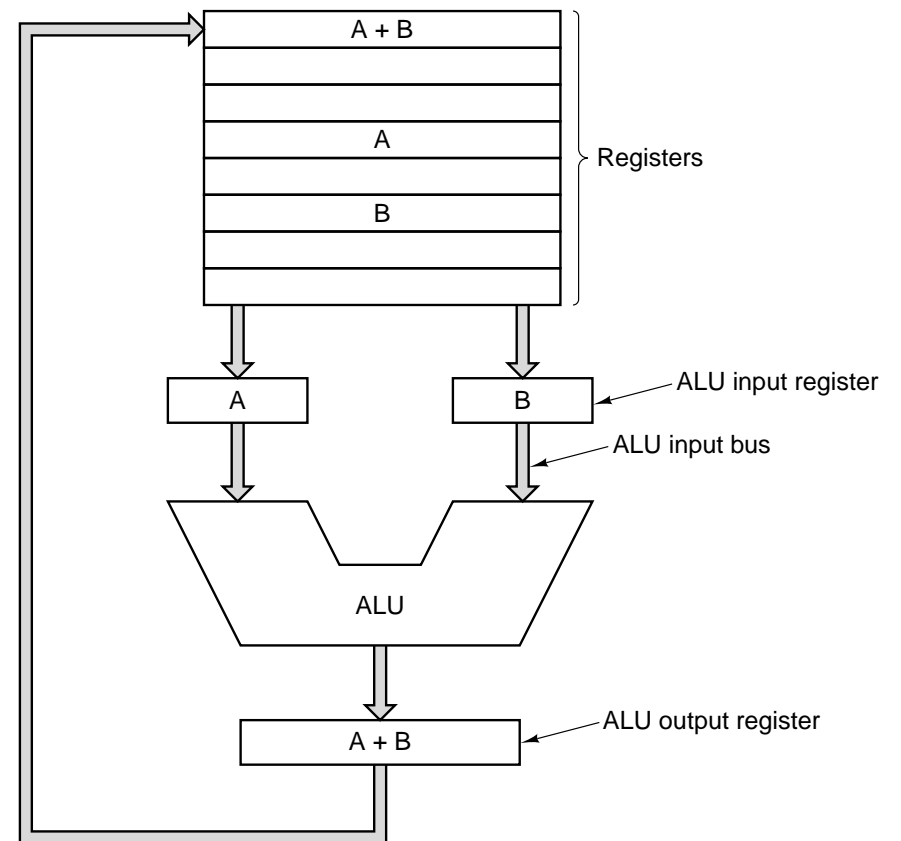
- **Control unit.**
  - Fetches instructions and determines their type.
- **Arithmetic logic unit:**
  - Performs operations needed to carry out instructions.
  - Example: integer addition (+), boolean conjunction ( $\wedge$ ).
- **Registers:**
  - Small high-speed memory to store temporary results and certain control information.
  - All registers have same size and can hold one number up to some maximum size.
  - There are general purpose registers and registers for special use.
  - **Program Counter (PC):** address of the next instruction to be fetched for execution.
  - **Instruction Register (IR):** instruction currently being executed.

## Data Path

ALU and registers connected by several buses.

- Registers feed into ALU input registers A and B.
  - Hold ALU input while ALU is computing.
- ALU performs operation and yields result in output register.
  - Content can be stored back into a register.
  - Register content can be stored in memory.

Data path cycle is core of CPU.



## Instruction Execution

CPU executes instruction in series of small steps.

1. Fetch the next instruction from memory into the instruction register.
2. Change the program counter to point to the following instruction.
3. Determine the type of the instruction just fetched.
4. If the instruction uses a word in memory, determine where it is.
5. Fetch the word, if needed, into a CPU register.
6. Execute the instruction.
7. Go to step 1 to begin executing the following instruction.

Fetch-decode-execute cycle.



## Interpreter for Simple Computer

```
static int PC, AC, instr, type, loc, data;
static boolean run_bit = true;

// interpret program stored in memory starting at start_address
public static void interpret(int memory[], int start_address) {
    PC = starting_address;
    while (run_bit) {
        instr = memory[PC];           // fetch next instruction
        PC = PC+1;                   // increment program counter
        type = getType(instr);       // determine instruction type
        loc = getLoc(instr, type);    // locate data (-1, if none)
        if (loc >= 0) data = memory[loc]; // fetch data
        execute(type, data);         // execute instruction
    }
}
```

## RISC versus CISC

- Since the late 1950s, instructions were **interpreted**.
  - Bugs could be easily fixed.
  - New instructions could be added at minimal cost.
  - More and more complex instruction sets emerged (**CISC**: 200–300 instructions).
- In the 1980s, a radically different concept emerged.
  - Patterson and Sequin (Berkeley): **RISC** design → SPARC architecture.
  - Hennessey (Stanford): MIPS architecture.
  - Simple instructions that could be **quickly issued** (started) and executed in **hardware**.
  - Small instruction sets (about 50 instructions).
- Today: CISC processors with a RISC core.
  - Simple instructions are executed in a single cycle (common instructions are fast).
  - More complex instructions are interpreted as usual (backward compatibility).
  - Example: Intel Pentium line.


## Modern Design Principles

- All common instructions are directly executed in hardware.
  - CISC instructions are broken into sequences of RISC instructions.
- Maximize the rate at which instructions are issued.
  - Less important how long execution of instruction takes.
- Instructions should be easy to decode.
  - Instructions have fixed length, regular layout, small number of fields.
- Only loads and stores should reference memory.
  - Operands of all other operations come from and return to registers.
- Provide plenty of registers.
  - Memory access is slow  $\Rightarrow$  32 registers or more.


Moreover: instruction-level parallelism (pipelining, superscalar).

# Pipelining

Divide instruction into sequence of small steps.



(a)

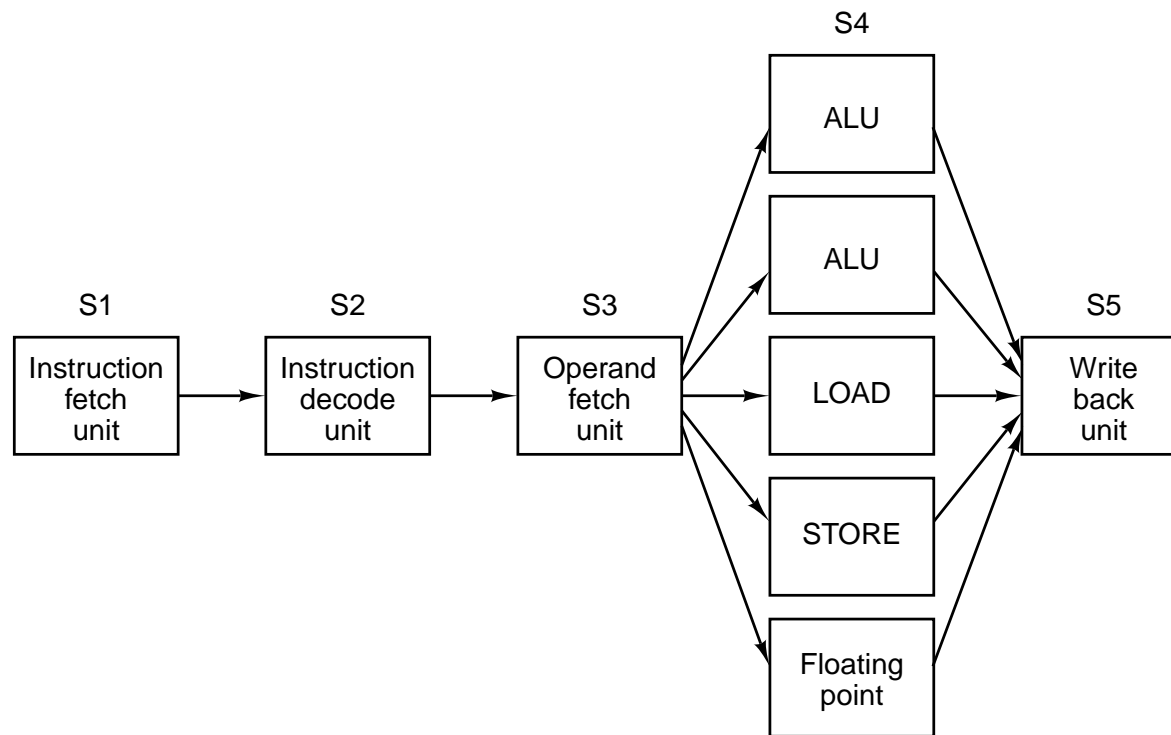


(b)

Processor operates at cycle time of the longest step.

## Superscalar Architectures

Have multiple functional units that operate independently in parallel.



**S3 stage must issue instructions fast enough to utilize S4 units.**

## Processor Clock

Execution is driven by a high-frequency processor clock.

- Example:

- 933 MHz processor (MHz = million Herz).
- Clock beats 933 million times per second.
- Clock speed doubles every 18 months (**Moore's law**).

- Clock drives execution of instructions:


- Instruction may take 1, 2, 3, ... clock cycles.
- 933 Mhz: at most 933 million instructions can be performed.
- Different processors have instructions of different power.
  - \* 550 MHz PowerPC processor may be faster than 933 Mhz Pentium III processor.

**Clock frequency is not a direct measure for performance.**

# Primary Memory

## Primary Memory

- Memory consists of a number of **cells**.
  - Each cell can hold  $k$  bits, i.e.,  $2^k$  values.
- Each cell can be referenced by an **address**.
  - Number from 0 to  $2^n - 1$ .
- Each cell can be independently read or written.
  - **RAM**: random access memory.




Memory is volatile: content is lost when power is switched off.



# Primary Memory Today

- Memory cells contain 8 bits.
  - Byte = smallest addressable memory unit.
- Bytes are grouped to **words**.
  - 4 bytes (32 bit computer) or 8 bytes (64 bit computer).
  - Most instructions operate on entire words (e.g. for adding them together).
- Bytes in a word are ordered in one of two ways.
  - **Big endian**: right-most byte has highest number.
    - \* Word 260 (= 256 + 4) is byte sequence |4|1|0|0|.
  - **Little endian**: left-most byte has highest number.
    - \* Word 260 (= 256 + 4) is byte sequence |0|0|1|4|.
  - Problem: data transfer between machines!



## Memory Units

Hierarchy of units of memory size.

Unit	Symbol	Bytes
byte		$2^0 = 1$ byte
kilobyte	KB	$2^{10} = 1024$ bytes
megabyte	MB	$2^{20} = 1024$ KB
gigabyte	GB	$2^{30} = 1024$ MB
terabyte	TB	$2^{40} = 1024$ GB

Multiplication factor to next unit is  $1024 = 2^{10}$ .

## Error-Correcting Codes

Computer memories make occasionally errors.

- Errors use **error-detecting** or **error-correcting** codes.
  - $r$  extra bits are added to each memory word with with  $m$  data bits.
  - When a word is read from memory, bits are checked to see if error has occurred.
- **Hamming distance** (Hamming, 1950).
  - Minimum number of bits in which two codewords differ.
  - 11110001 and 00110000 have Hamming distance 3.
  - With Hamming distance  $d$ ,  $d$  single-bit errors are required to convert one word into the other.
- Construct  $n$ -bit codeword where  $n = m + r$ .
  - $2^n$  codewords can be formed.
  - $2^m$  codewords are legal.
  - If computer encounters illegal codeword, it knows that error has occurred.


## Code Properties

Hamming distance of a code is minimum distance of two codewords.

- **Detect**  $d$  single-bit errors.
  - Errors must not change one valid codeword into another.
  - Hamming distance  $d + 1$  is required.
- **Correct**  $d$  single-bit errors.
  - Codewords must be so far apart that error word is closest to original codeword.
  - Hamming distance  $2d + 1$  is required.
- **Example: parity bit code.**
  - Add single bit such that total number of 1 bits is even.
  - Each single-bit error generates codeword with wrong parity.
  - Hamming distance 2, single-bit errors can be detected.

## Cache Memories

- CPUs have always been faster as memories.
  - When CPU issues a memory request, it has to wait for many cycles.
- Problem is not technology but economics.
  - Memories that are as fast as CPUs have to be located on CPU chip.
  - Chip area is limited and costly.
- **Cache**: small amount of fast memory.
  - Most heavily used memory words are kept in cache.
  - CPU first looks in cache; if word is not there, CPU reads word from memory.





## Cache Performance

- **Locality principle:** basis of cache operation.
  - In any short time interval, only a small fraction of total memory is used.
  - When a word is referenced for the first time, it is brought from memory to cache.
  - The next time the word is used, it can be accessed quickly.
  - 1 reference to slow memory,  $k$  references to fast cache.
- **Cache lines:** units of memory transfer.
  - When cache miss occurs, entire sequence of words is loaded from memory.
  - Typical cache line is 64 bytes (16 words a 32 bit).
  - Reference to address 260 loads cache line 256–319.
  - Chances are high that also other words in cache line will be used in near future.
- **Today CPUs have multiple caches.**
  - Primary cache on CPU chip.
  - Larger secondary cache in same package as CPU chip.

## Memory Packaging and Types

Group of chips is mounted on a circuit board and sold as a unit.

- **SIMM**: single inline memory module.
  - Single row of connectors on one side of the board (72 connectors delivering 32 bits).
  - For instance: eight chips with 4 MB each.



- **DIMM**: dual inline memory module.
  - Two rows of connectors on each side (168 connectors delivering 64 bits).


Error detection and correction are omitted for PC memory.

# Secondary Memory



## Secondary Memory


Main memory is always too small.



Non-volatile storage: keeps data without power.

## Magnetic Disk (Harddisk)


Magnetic medium: bits are represented by small magnetized particles.



Much larger but also much slower than primary memory.

## Harddisk Organization

- Multiple spinning platters are organized as a stack.
- Read/write head passes over platter.




All read/write heads are moved by the same arm.

## Platter Organization

Aluminium platter has magnetizable coating.

- Read/write head floats on a cushion of air over its surface.
  - Writing: current flowing through head magnetizes surface beneath.
  - Reading: magnetized surface induces current in head.
  - As platter rotates, stream of bits can be read or written.
  - Disk track: circular sequence of bits written in one rotation.



## Data Organization

- Each track is divided into **sectors** of fixed size.
  - **Preamble** that allows head to be synchronized before reading or writing.
  - 512 data bytes (typically).
  - Error-correcting code (ECC).
  - Intersector gap.
- **Formatted capacity:** size of data areas (only).
- Two densities define storage capacity.
  1. **Radial density:** number of tracks per radial cm (800–2000).
  2. **Linear bit densities:** number of bits per track cm (up to 100.000).
- **Cylinder:** set of tracks at a given radial position.
  - All read/write heads can access one cylinder at a time.

## Disk Performance

Performance is determined by various factors.

- **Seek time:** time to move the arm to the right radial position.
  - 5–15 ms average seek time (between random tracks).
- **Rotational delay:** time until desired sector rotates under head.
  - Average delay is half a rotation.
  - 7200 rpm (rotations per minute): 4 ms.
- **Transfer time:** time to read sector.
  - Depends on linear density and rotation speed.
  - 20 MB/s: 25  $\mu$ s to read a 512 byte sector.
- Seek time and rotational delay dominate transfer time.
  - **Burst rate:** data rate once the head is over the first data bit.
  - **Sustained rate** is much smaller than burst rate.

## Performance Comparison

Memory access times have different orders of magnitude.

Operation	Clock Cycles
CPU instruction	1–3
Register access	1
Cache access	1
Main memory access	10
Disk seek time	10.000.000

Magnetic disk is one million times slower than main memory.

## Disk Controllers

A CPU built into the disk controls its operation.

- Accepts commands from the main processor.
  - READ, WRITE, FORMAT.
- Controls arm motion.
- Detects and corrects errors.
- Converts bytes read from memory into serial bit stream.
- Cache sector reads and buffer sector writes.
- Remaps bad sectors by spare sectors.

Disk controller frees CPU from low-level details of operation.



## Disk Controller Interfaces

- **(E)IDE**: (Extended) Integrated Drive Electronics.
  - IDE: sectors are addressed by head, cylinder, sector numbers (512 MB limit).
  - EIDE: up to  $2^{24}$  consecutively addressed sectors (**LBA**: Logical Block Addressing).
  - ANSI standard: **ATAPI** (AT Attachment Packet Interface).
  - Technology for cheap PC harddisks, floppy drives, CD-ROM drives, tape drives, ...
- **SCSI**: Small Computer Systems Interface.
  - Bus to which controller and chain of up to seven devices can be attached.
  - All devices can run at the same time.
  - SCSI-1: 5 MHz clock rate, 8 bit data transfer.
  - Fast SCSI (10 MHz bus), Ultra SCSI (20 MHz), Ultra2 SCSI (40 MHz).
  - Wide SCSI: 16 bit transfer.
  - 5–80 MB/s transfer rate (Ultra2 Wide-SCSI).
  - Higher performance but more expensive than EIDE/ATAPI.

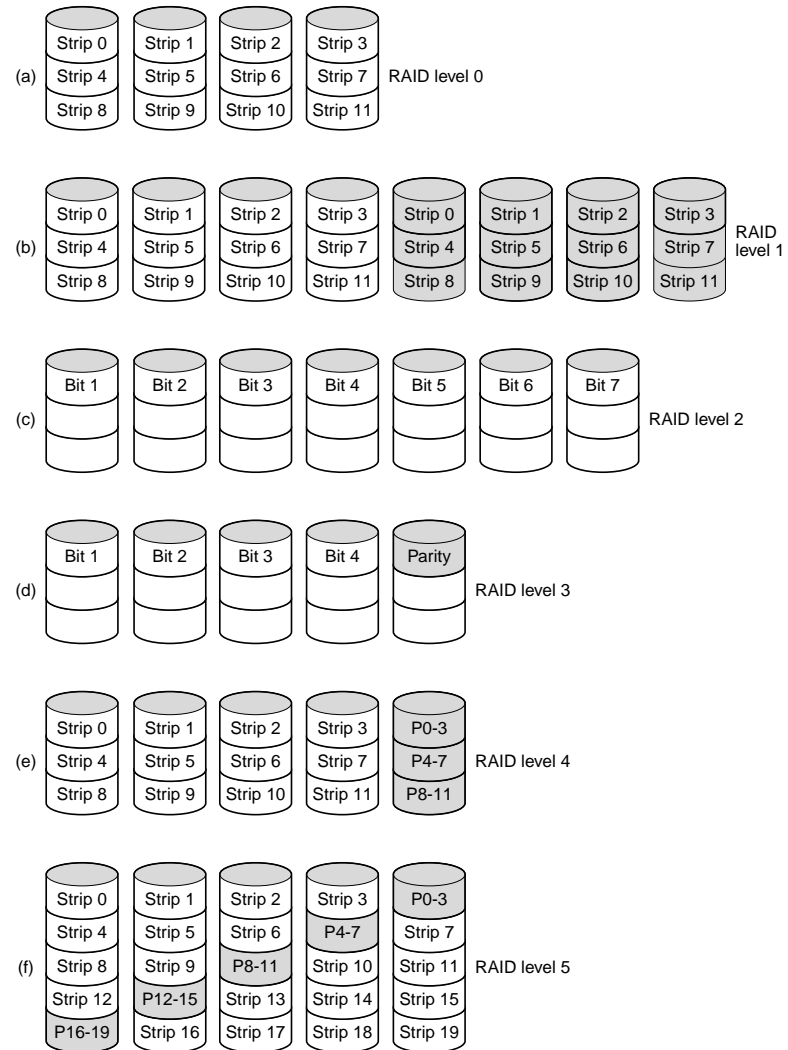
## RAID

### Redundant Array of Inexpensive Disks.

- Box of disks can be operated as a single disk.
  - RAID controller presents single disk to operating system.
  - Typically: SCSI disks operated by RAID SCSI controller.
- **RAID levels** define data distribution/redundance properties.
  - **RAID 0**: strips of  $k$  sectors are **distributed** round robin across disks (single disk).
  - **RAID 1**: strips are **mirrored** on two sets of disks (fault recovery).
  - RAID 2: bits are distributed across disks (data rate).
  - RAID 3: simplified version of RAID 2, parity bit written to extra drive.
  - RAID 4: like RAID 0, with parity bits written to extra drive (fault recovery).
  - **RAID 5**: like RAID 0, with parity bits distributed across drives (fault recovery).

**Also software implementations of RAID (Linux).**

# RAID



## CD-ROMs

Compact Disk (CD): optical disk by Philips and Sony (1980).

- Production:


- High-power infrared laser burns small holes in coated glass master disk.
- From master, mold is made with bumps where laser holes were.
- Molten polycarbonate is injected to form CD with same pattern as master (**pits** and **lands**).
- Thin layers of reflective aluminium and of protective lacquer are added.

- Play back:

- Low-power laser diode shines infrared light on pits/lands as they stream by.
- Lights reflecting off a pit is out of phase with light reflecting off the surrounding surface.
- Two parts interfere destructively, thus less light is returned to photodetector.
- Pit/land transition represents bit 0, land/pit transition represents bit 1.

**First successful mass market digital storage medium.**


## CD-ROM Physical Layout



- Pits/lands are written in a single continuous spiral inside-out.
  - 22.188 revolutions (600 per mm), 5.6 km length.
- Pits/lands must stream by at constant velocity.
  - Rotation rate of CD is reduced as reading head moves inside-out.
  - 530 rpm – 200 rpm.

## CD-ROM Logical Layout

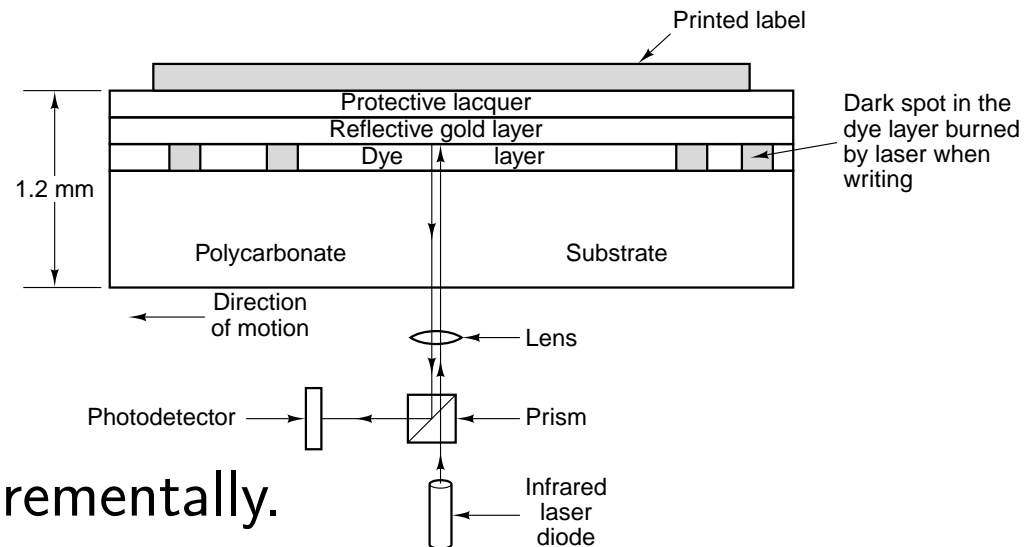
- 1984: CD-ROM (CD - Read Only Memory).
  - Data standard which is physically compatible with audio CDs.
  - Sector of 98 frames with 42 symbols of 14 bits each (1 byte + ECC).



- 1986: multimedia CD-ROMs.
  - Audio, video, data interleaved in same sector.
- Finally: ISO 9660 filesystem.

## CD-Recordables

- 1989: CD-Rs (recordable CDs).
  - Gold used instead of aluminium for reflective layer.
  - Layer of dye inserted between polycarbonate and reflective layer.
  - When high-power laser hits dye, it creates a dark spot which is interpreted as a pit.



- CD-Rs can be written incrementally.
  - Track: group of consecutive sectors written at once (without interruption).
  - Multi-session: multiple VTOCs (Volume Table of Contents).

## CD-ROM/R Characteristics

- Data transfer: 153 KB/s.
  - Single speed drive.
  - Today: up to 80× drives.
- Seek time: several 100s of ms.
  - No match for hard disk.
  - High streaming rates, low random access rates.
- Capacity: 74 minutes of music.
  - 650 MB of data.
  - Higher capacities available.


Primary purpose is data transfer and backups.



# DVD

## Digital Versatile Disk.

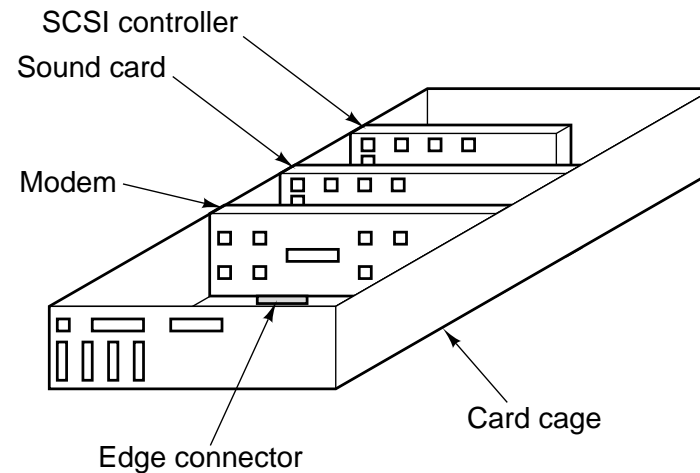
- Same physical design as CDs.
  - Smaller pits, tighter spiral, red laser.
  - 4.7 GB data capacity, 1.4 MB/s data transfer (single speed).
  - DVD drives can read also CD-ROMs.
- Further formats.
  - Capacity: Dual layer (8.5 GB), double sided (9.4 GB), double sided dual layer (17 GB).
  - Rewriting: DVD-R, DVD-RW, DVD+RW (let the battle begin!)



# Input/Output Devices

## Physical Computer Structure

Metal box with with the motherboard at bottom.



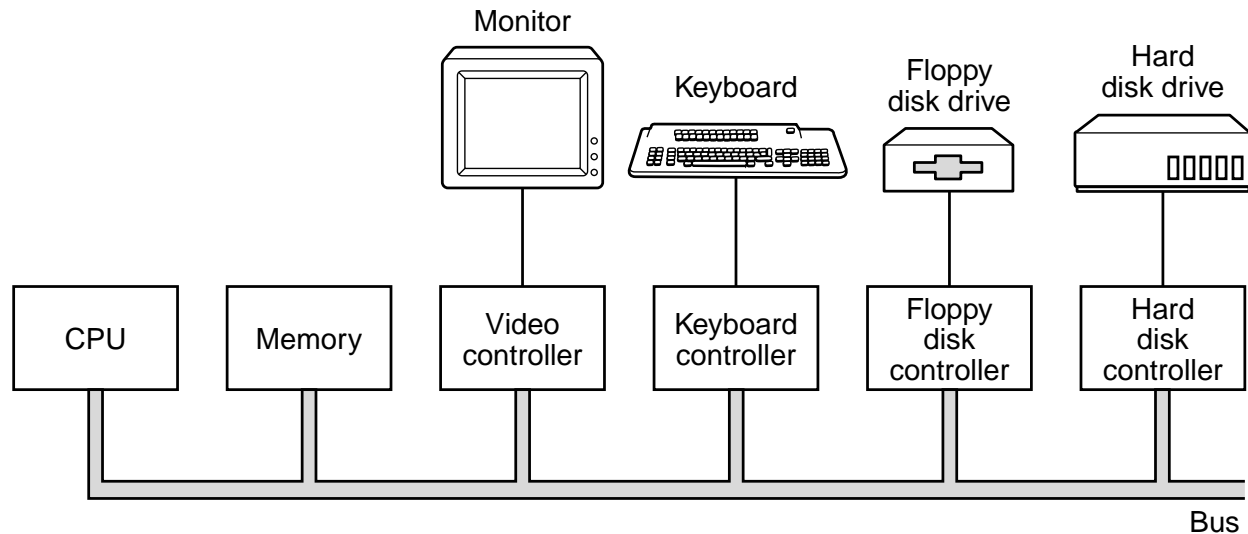
I/O boards can be connected to the motherboard.

## Motherboard



CPU, DIMM slots, support chips, bus, I/O slots.

## Logical Computer Structure



- Each I/O device has a **controller** which contains the electronics.
  - Typically contained on a board plugged into a free slot (video controller).
  - Sometimes contained on the board itself (keyboard controller).
  - Connects to device by a cable attached to a connector on the back of the box.

## Video Controller

- Monitor is passive device controlled by **video card**.
  - Active device connected to system bus (or special graphics bus).
  - Holds in local memory digital image.
  - Transfers it pixel by pixel to monitor for display.
  - Refresh rate e.g. 80 Hz, data transfer:  $80 \times 2560 \text{ KB} = 200 \text{ MB}$  per second!

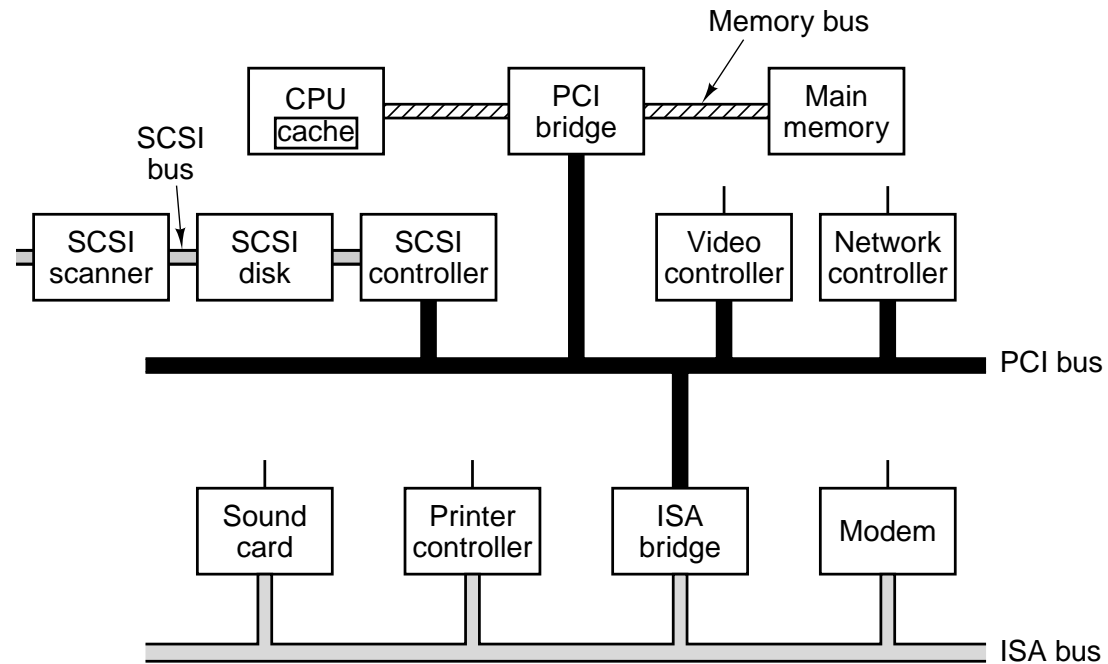


## Controller Operation

Controller controls I/O device and handles bus access for it.

- Takes command from CPU and executes it.
  - Disk controller: issues commands to drive to read data from a particular track and sector.
- Controller returns result to computer memory.
  - **DMA (Direct Memory Access)**: controller can access memory without CPU intervention.
  - When finished, controller causes an **interrupt** in the CPU.
  - CPU suspends current program and executes an **interrupt handler**.
  - Interrupt handler informs operating system that I/O has finished.
  - Afterwards, CPU continues with suspended program.
- Bus access controlled by **bus arbiter**.
  - Chip which decides which device may access the bus.
  - Usually I/O devices are given preference over CPU.

## Modern PC Structure





PCI (Peripheral Component Interconnect) is the most popular bus technology today.



## Computer Monitors

- Typical characteristics:
  - 19" monitor (US inch = 2.54 cm, diagonal size).
  - Resolution:  $1280 \times 1024$  pixels (picture units).
  - 16 bit graphics: a pixel may have  $2^{16} = 65536$  colors.
  - Information content:  $1280 \times 1024 \times 16/8$  bytes = 2560 KB
- Pixel consists of 3 color elements.
  - RGB: red, green, blue.
  - Control of combination/brightnesses.
- Dominant technologies:
  - CRT (Cathode Ray Tube).
  - LCD (Liquid Crystal Display)




## Other I/O Devices

- Keyboard and mouse.

- Interactive input devices.
- Connected to controllers on main board.
- Signals (key pressed, mouse moved) transferred to controller.
- Controller notifies processor to handle request.

- Laser printer.

- Output device with local processor and memory.
- Controlled by printer programming language (e.g. PostScript).
- Computer constructs this program and sends it to printer.
- Printers paint image on electrically charged drum to which toner powder sticks.
- Print resolution measured in dpi (dots per inch); good quality  $\geq 300$  dpi.
- Color laser printers overlay images in subtractive primary colors and black (CYMK).



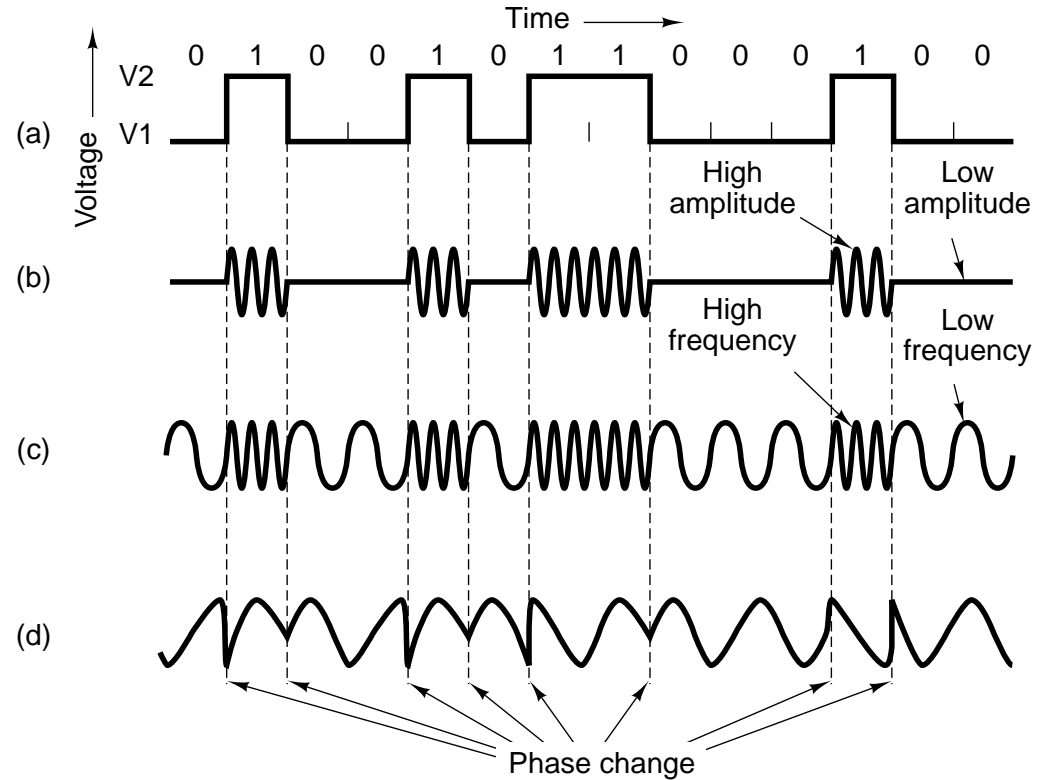
## Network Devices

Computers connected in various ways.

- **Modem** (modulator/demodulator)
  - Converts digital information to analog signals transferred via phone lines.
  - 56 Kbit/s (KBit = 1000 bit).
- **ADSL** (Asynchronous Digital Subscriber Line)
  - Conventional phone lines, special end devices.
  - Download speeds: 512 KBit/s and more (upload: 64 Kbit/s).
- **Ethernet**
  - Most prominent technology for Local Area Networks (LANs).
  - Copper wires used as buses.
  - 10–100 Mbit/s (FastEthernet), 1 Gbit/s (GigaEthernet).



# Modulation Principles



Modulation by amplitude, frequency, phase.