System Buses

EE2222 Computer Interfacing and Microprocessors

Partially based on
Computer Organization and Architecture by William Stallings
Computer Electronics by Thomas Blum

Connecting

- All the units must be connected
- Different type of connection for different type of unit
  - CPU
  - Memory
  - Input/Output

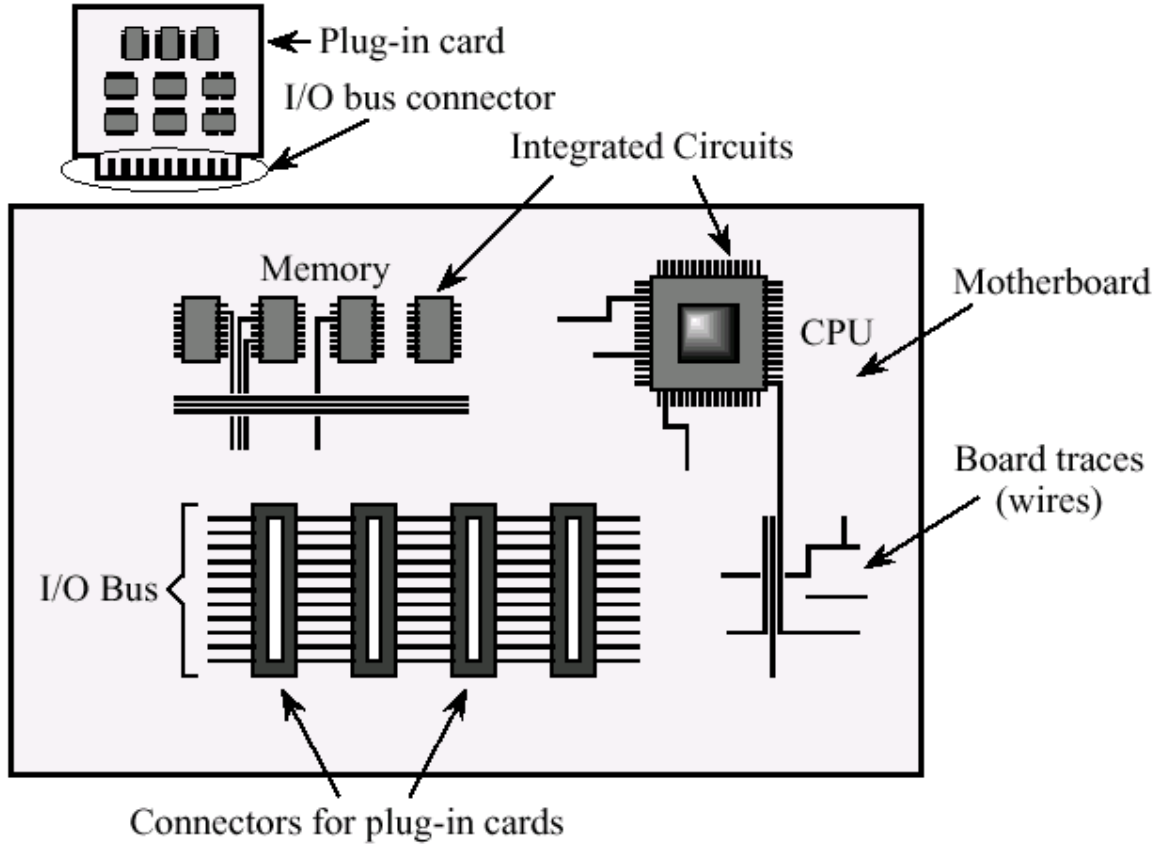

CPU Connection

- Reads instruction and data
- Writes out data (after processing)
- Sends control signals to other units
- Receives (& acts on) interrupts

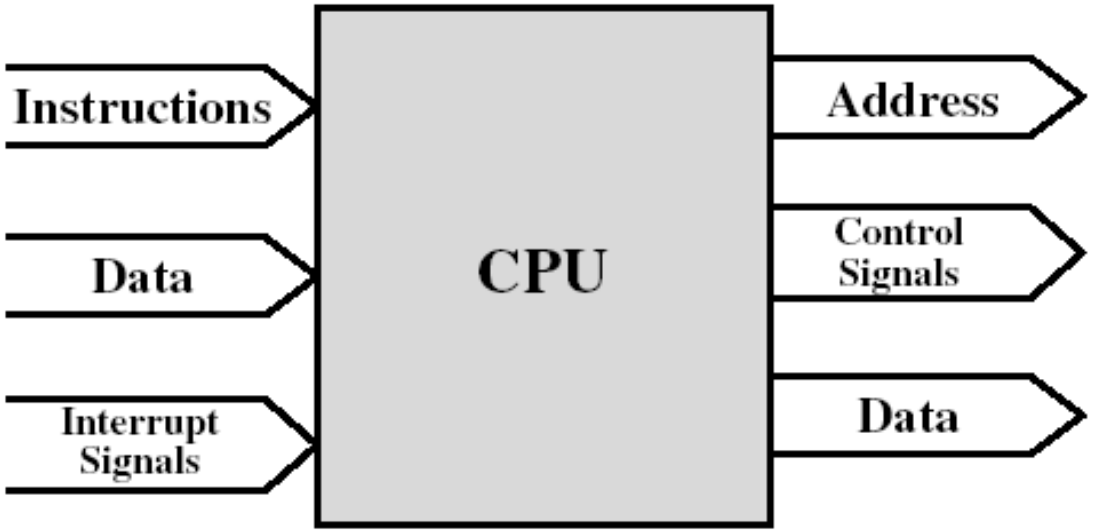

Memory Connection

- Receives and sends data
- Receives addresses (of locations)
- Receives control signals
  - Read
  - Write
  - Timing

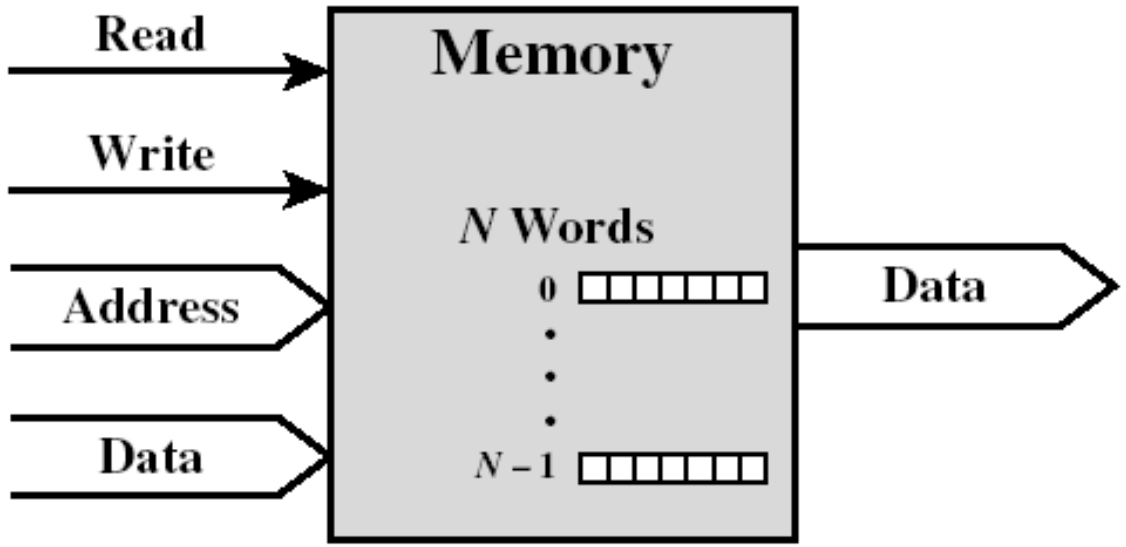

Input/Output Connection(1)

- Similar to memory from computer's viewpoint
- Output
  - Receive data from computer
  - Send data to peripheral
- Input
  - Receive data from peripheral
  - Send data to computer

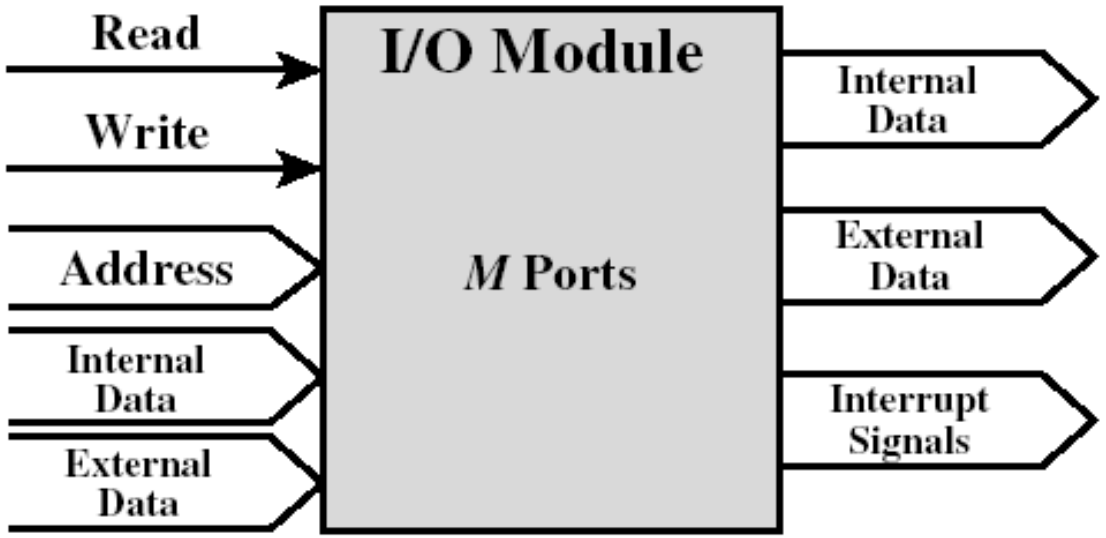

Input/Output Connection(2)

- Receive control signals from computer
- Send control signals to peripherals
  - e.g. spin disk
- Receive addresses from computer
  - e.g. port number to identify peripheral
- Send interrupt signals (control)

What is a Bus?

- A communication pathway connecting two or more devices
- Usually broadcast (all components see signal)
- Often grouped
  - A number of channels in one bus
  - e.g. 32 bit data bus is 32 separate single bit channels
- Power lines may not be shown

Bus Interconnection Scheme

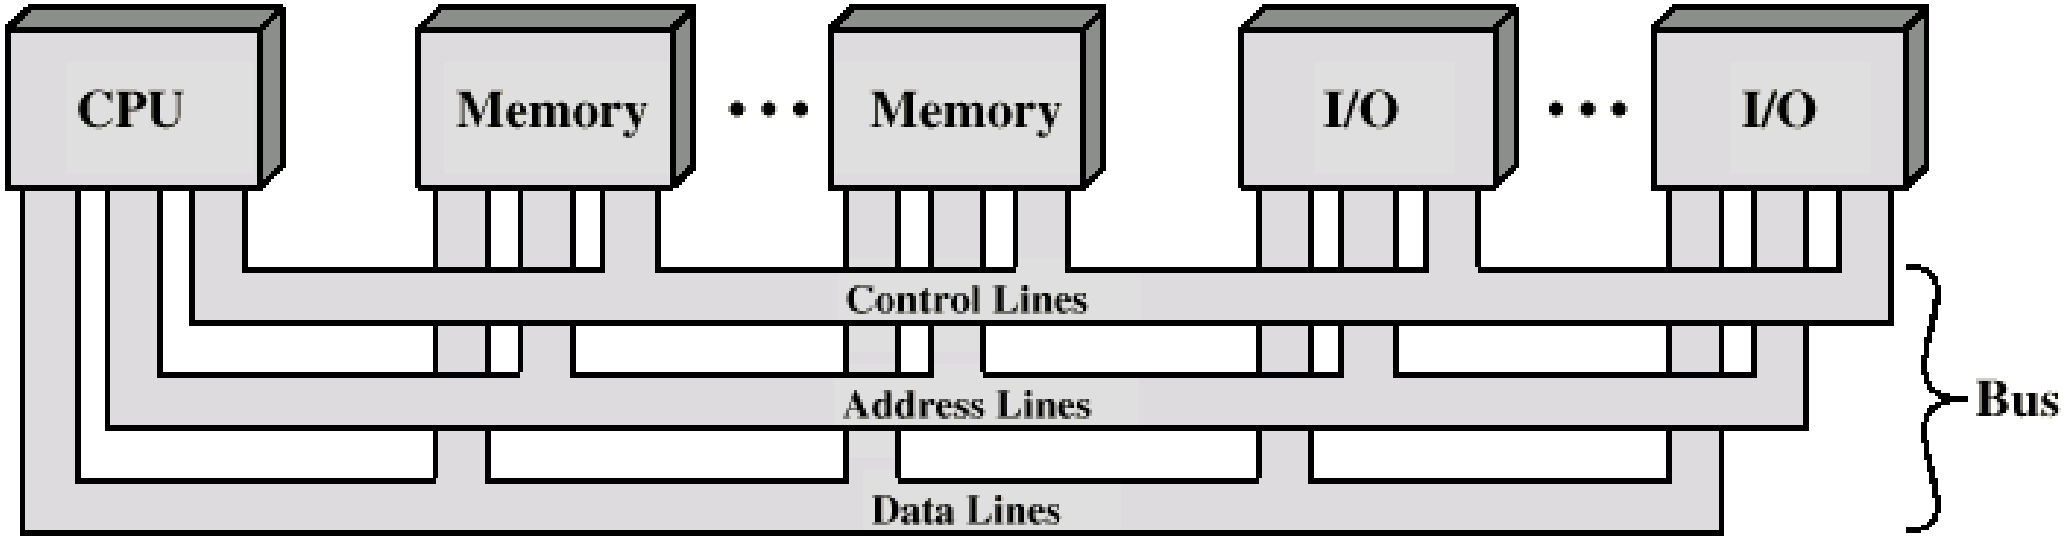

Data bus

- Carries data
  - Remember that there is no difference between “data” and “instruction” at this level
- Width is a key determinant of performance
  - 8, 16, 32, 64 bit

Address bus

- Identify the source or destination of data
  - e.g. CPU needs to read an instruction (data) from a given location in memory
- Bus width determines maximum memory capacity of system
  - e.g. 8080 has 16 bit address bus giving 64k address space

Control bus

- Control and timing information
  - Memory read/write signal
  - Interrupt request
  - Clock signals

What do buses look like?

- Parallel lines on circuit boards
- Ribbon cables
- Sets of wires
- Strip connectors on mother boards
  - e.g. PCI

Bus types

- Dedicated
  - Separate data & address lines
- Multiplexed
  - Shared lines
  - Address valid or data valid control line
  - Advantage - fewer lines
  - Disadvantages
    - More complex control
    - Ultimate performance

Bus characteristics

- The highway analogy:
  - moving data along the buses is like moving cars on the highway.
- Bus width (number of lanes)
  - How many bits are moving around in parallel
- Bus speed (speed limit)
  - How fast those bits are moving
- **Throughput** is the numbers of bits being handled per unit time, it combines bus width and bus speed into one measure.

Memory size

- The width of the system's address bus puts an upper limit on the amount of memory locations.
- For example, if the address bus width is 32, then there are 232 (4,294,967,296) addresses.
- Note that instead of addressing individual words, computers usually address individual bytes, so that would mean 4 GB.
- A few years ago 4 GB of memory seemed quite large, but now ….
  - To exceed this limit, one needs a wider bus or one needs to break the addresses into parts.

Bus speeds

- Measured in MHz (millions of cycles per second).
- It doesn’t make much sense to have a very fast processor speed and a slow bus speed; they should be compatible.
  - The bus speed is slower than the processor speed and often limits the speed of the computer.

Single bus problems

- Lots of devices on one bus leads to:
  - Propagation delays
    - Long data paths mean that co-ordination of bus use can adversely affect performance
    - If aggregate data transfer approaches bus capacity
- Most systems use multiple buses to overcome these problems

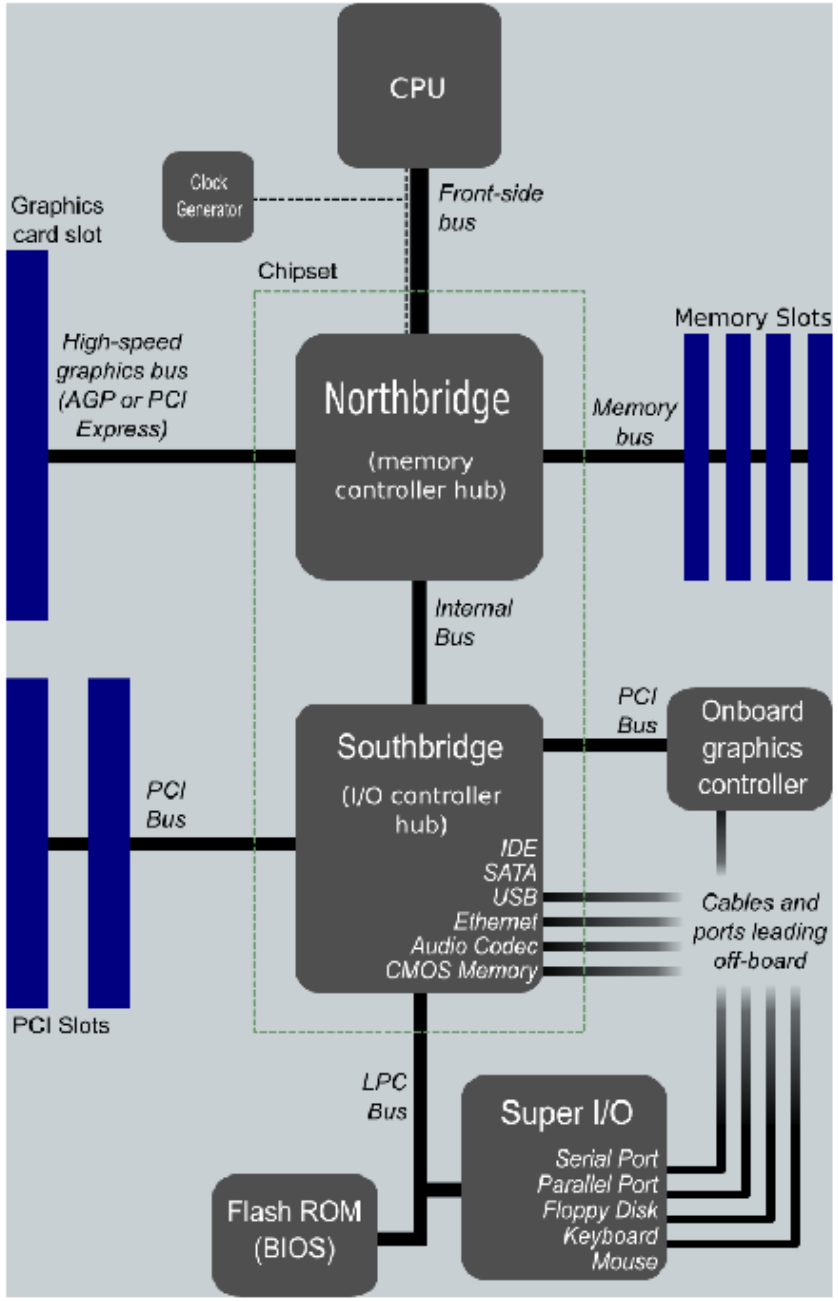

System bus

- The system bus connects the CPU, memory and other motherboard parts.
- This bus should be well coordinated with the processor and memory access speeds.
- Other buses must interface with the system bus if they want to interact with the processor.

Expansion bus

- The expansion bus connects the system bus to the expansion slots (where cards are inserted to expand the computer's capabilities)
  - This bus usually works at slower speeds
- Early PCs used an expansion bus called the ISA bus.
- Most PCs today have a much faster PCI bus but may have an ISA bus for backward compatibility.

Local bus

- If a device or devices require a great deal of speed (e.g. video), then one solution is for the device to have its own high-speed, direct (or nearly direct) connection to the processor.
- Such a connection is called a **local bus**.
- Can only support a few devices.

Traditional (ISA) (with cache)

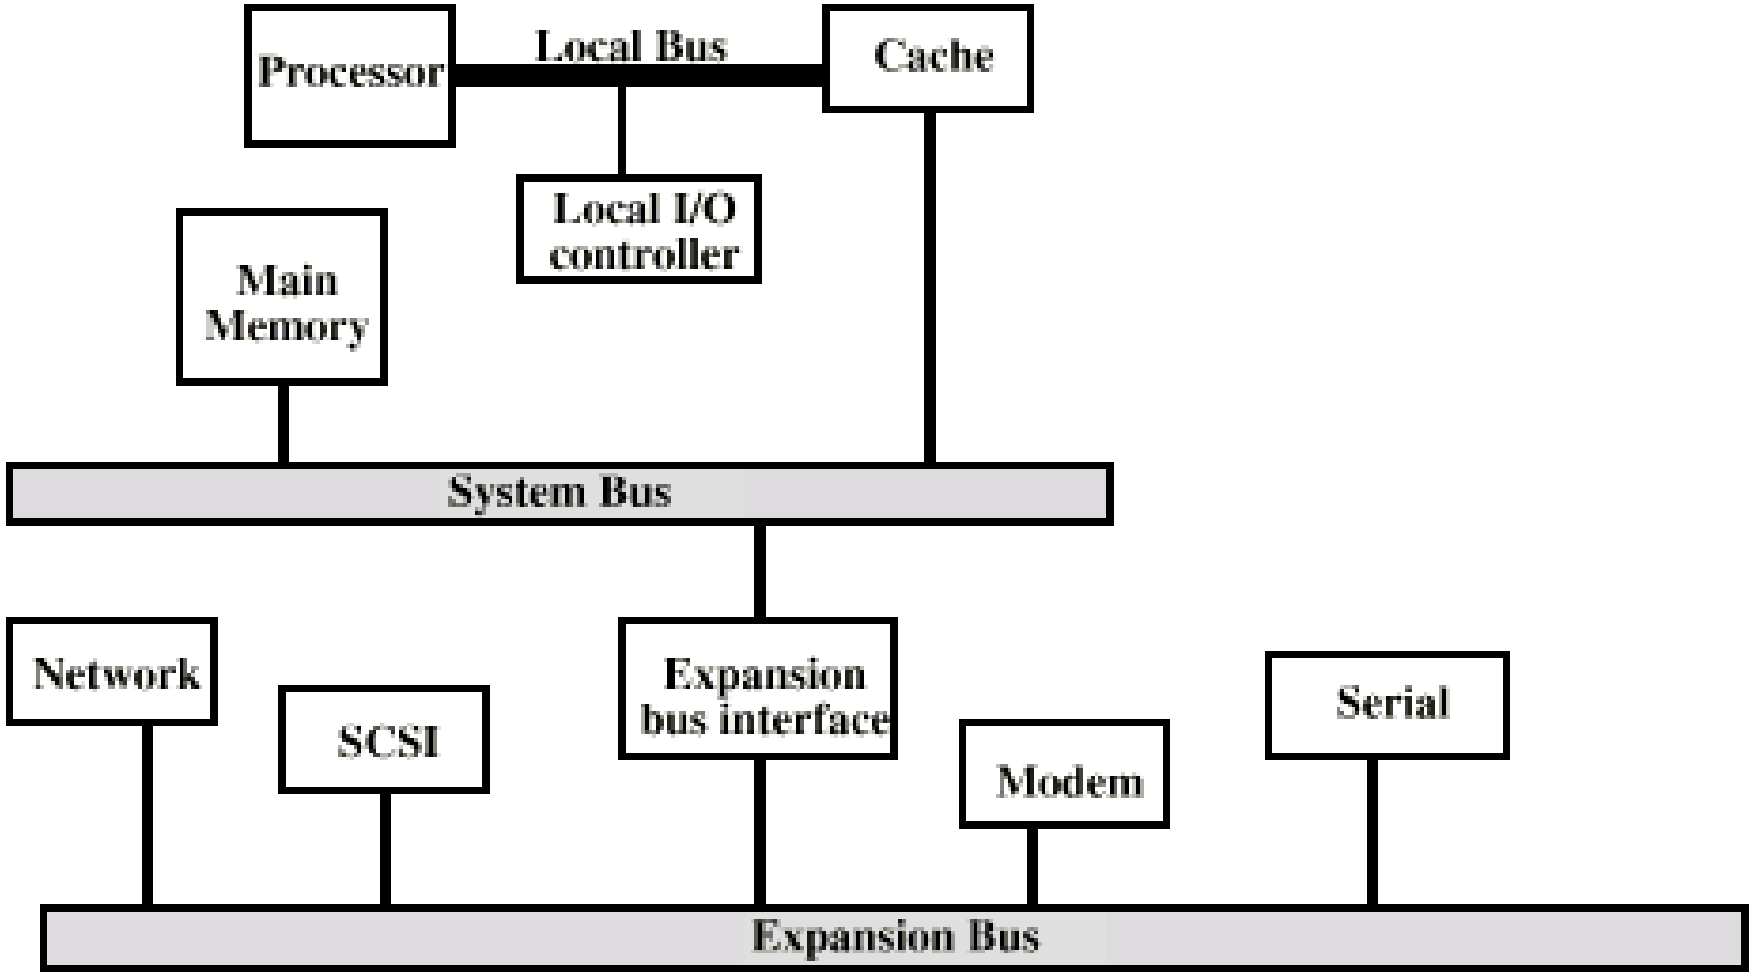

High Performance Bus

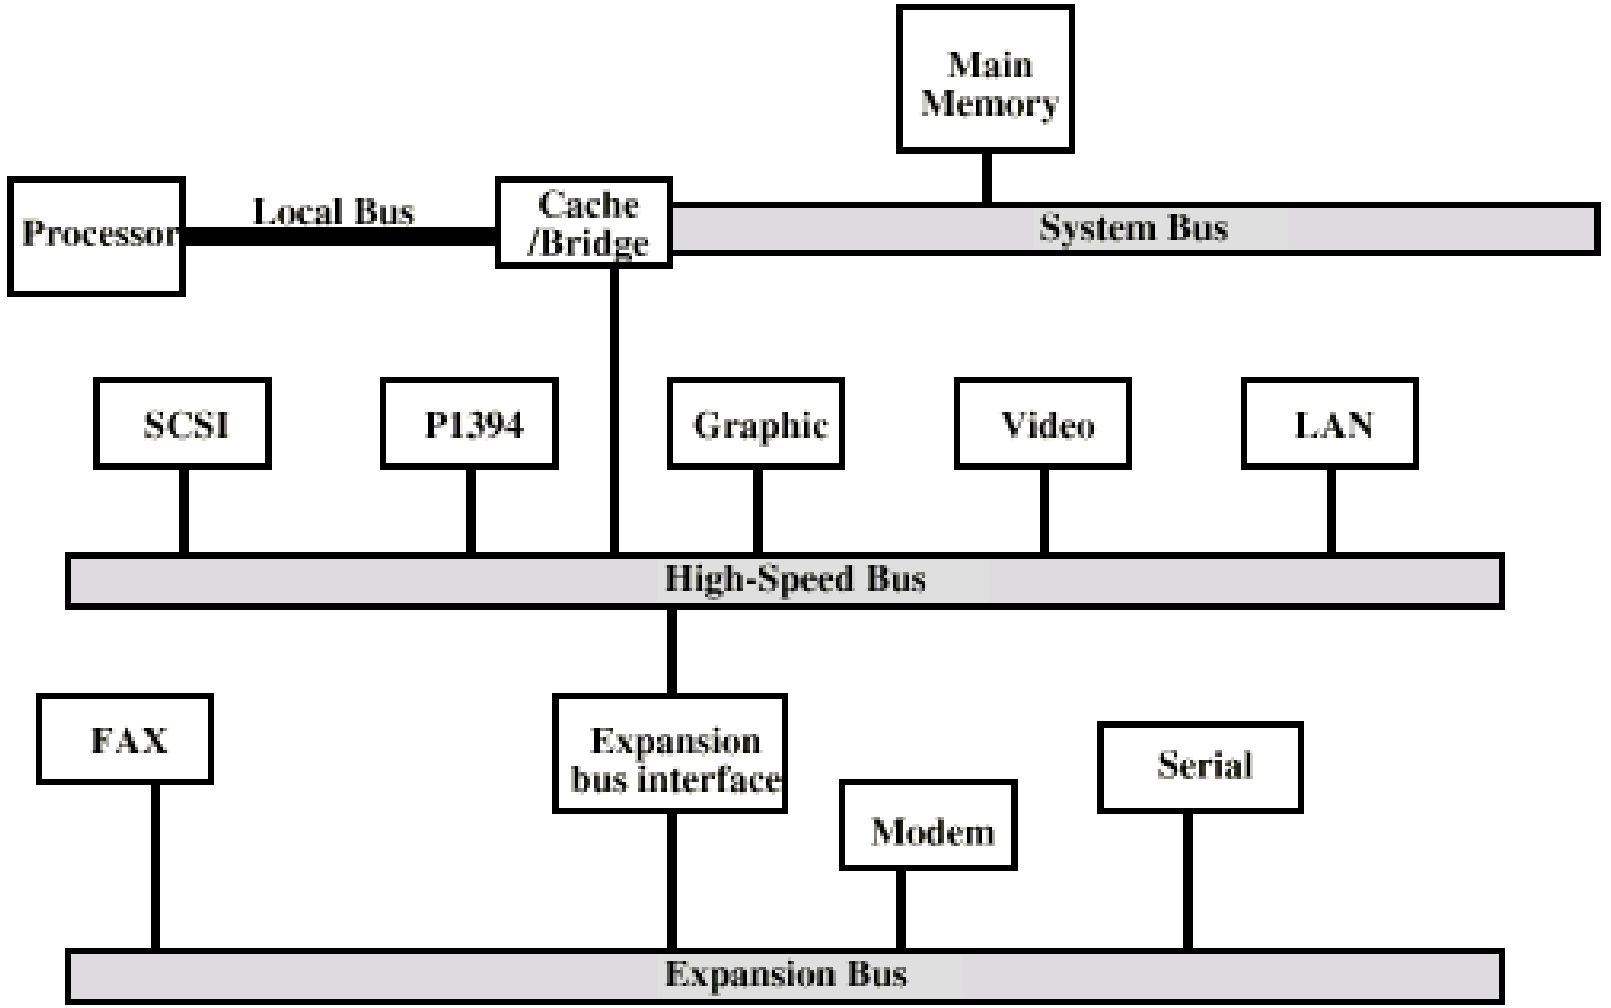

ISA

- **Industry Standard Architecture (ISA)** is the bus used in early IBM PC and their clones.
- The AT (advanced technology) version of the bus is called the “AT” bus and became an industry standard.
- Worked at 8.33 MHz

Plug and Play

- In 1993, Intel and Microsoft introduced a version of the ISA called **Plug and Play ISA**.
- Plug and Play ISA enables the operating system to do the configuring, instead of the user setting switches and jumpers.

PCI

- **Peripheral Component Interconnect (PCI)**, a standard introduced by Intel.
- PCI is a 64-bit bus, though it is usually implemented as a 32-bit bus.
- It can run at clock speeds of 33 or 66 MHz.
- At 32 bits and 33 MHz, it yields a throughput rate of 133 MBps (Mega bits per second).

PCI-X

- The **PCI** bus connects to the Southbridge.
- But as devices required more speed than the PCI standard allowed, a faster version, **Peripheral Component Interconnect Extended (PCI-X)** was developed by IBM, HP, and Compaq.
- Whereas PCI allows up to 532 MB per second, PCI-X allows up to 1.06 GB per second.
- PCI-X is backward-compatible,
  - You can install a PCI-X card in a PCI slot (it will operate at the slower PCI speed).

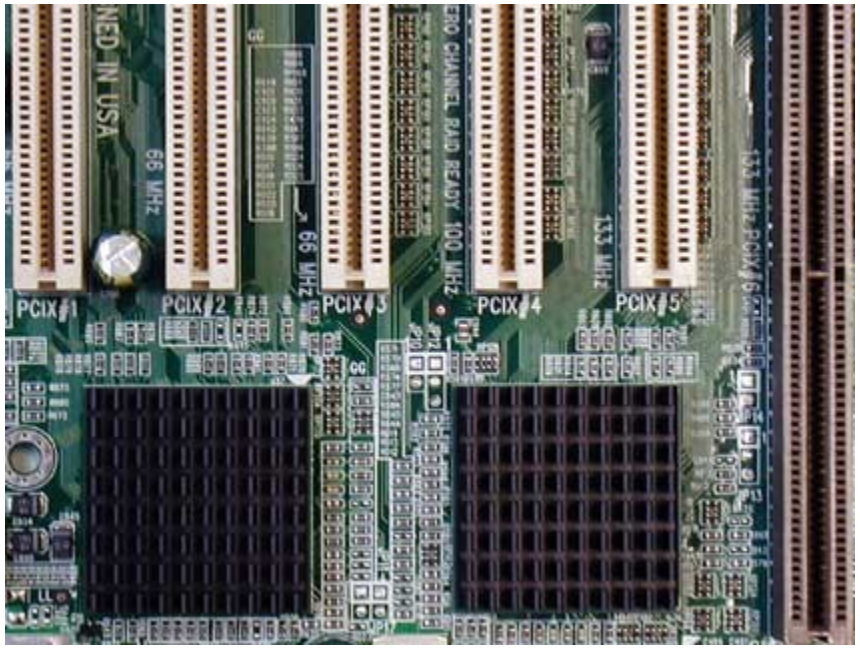

26

PCI Express

- **PCI Express**, **PCI-E** or **PCIe**
- Introduced by Intel in 2004.
- Built to be fast like AGP (faster actually) but general purpose like PCI.
- It has point-to-point rather than a (shared) bus structure. And has full-duplex serial connections called **lanes** –up to 32 lanes.
- In PCIe 1.1 (the most common for now) each lane carries 250 MB/s (per direction).
- PCIe 2.0 supports 500MB/s and PCIe 3.0 supports 1GB/s **per lane**.

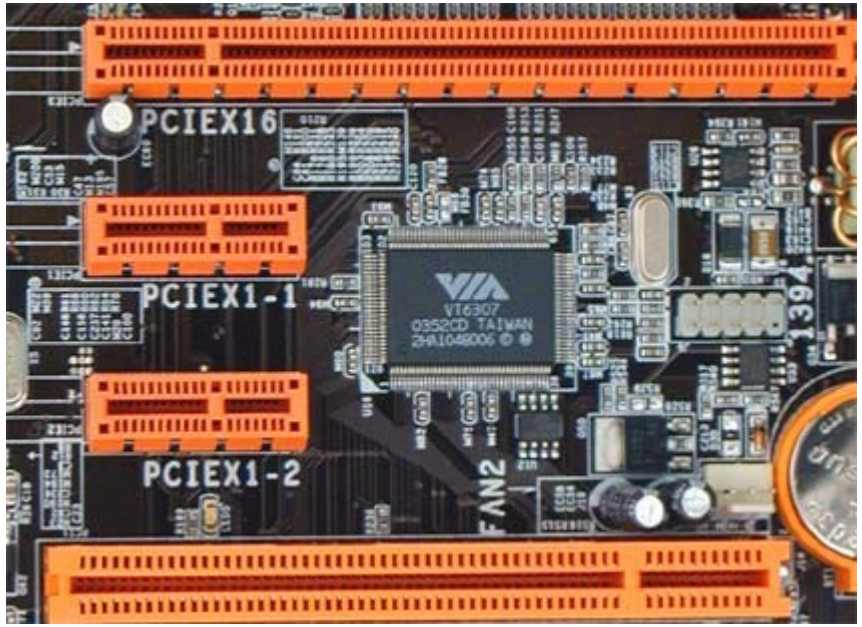

External bus

- A bus that connects a computer to peripheral devices.
- A common example is the **Universal Serial Bus (USB)**.
  - A single USB port can be used to connect up to 127 peripheral devices, such as mice, modems, and keyboards.
  - USB also supports Plug-and-Play installation and hot plugging.

Bus mastering

- Enables a controller connected to the bus to communicate directly with other devices on the bus without going through the CPU.
- Most modern bus architectures, including PCI, support bus mastering because it improves performance.

Further Reading

- Stallings, Chapter 3
- https://computer.howstuffworks.com/computer-buses-channel.htm