

Computer Organization


What's in the box of a typical PC?

- motherboard
- CPU
- power supply
- RAM memory
- hard drive
- GPU graphics processor
- network card
- optionally, other drives: CD/DVD drive, solid-state drive

motherboard

contains sockets and connectors:

- a socket for the CPU
- connectors for memory and other hardware
- common forms are called "form factors", ex: the ATX advanced technology extended



CPU

- modern CPUs have many cores
- modern CPUs employ multithreading that can perform two execution threads per core
- clock speeds range from 3 to about 5 GHz




power supply

- supplies power to different components at different voltages



RAM memory

- plugs into motherboard
- the CPU has limited internal memory, RAM serves as auxiliary memory
- memory is swapped between CPU and RAM as needed
- RAM loses its contents upon power loss



drives

- internal hard drives: magnetic disk or solid state
- CD/DVD drives
- long-term permanent storage



device interface

hard disk - traditional drive with internal platters and read heads

interfaces for connecting:

- IDE, EIDE connects to motherboard with ribbon cables
- Serial ATA (SATA) serial advanced technology attachment - faster and less bulky cable
- SCSI - high-performance disk drives (more expensive)
- SAS (serial attached SCSI) - higher speed
- IEEE 1394 (FireWire) - high speed data transfer (audio/video)
- USB - universal serial bus for communication with devices

USB types



USB TYPE A



USB TYPE B



USB TYPE C



USB MINI A



USB MINI B



USB MICRO A



USB MICRO B



USB MICRO B SUPER SPEED


device interface continued

- PCIe - Peripheral Component Interconnect Express - is a technology for I/O
- PCIe allows anywhere from 1 to 32 lanes connected to I/O devices
- There used to be separate chips to connect the microprocessor to memory and high-speed I/O devices
- More recently, these connections are internal to the chip, as in Intel's Sandy Bridge


GPU graphics processor

- rapidly processes memory in parallel
- primarily used for graphics
- also used for machine learning and more
- some older or lower-end PCs have an integrated GPU on the motherboard
- Nvidia and AMD are the top GPU creators





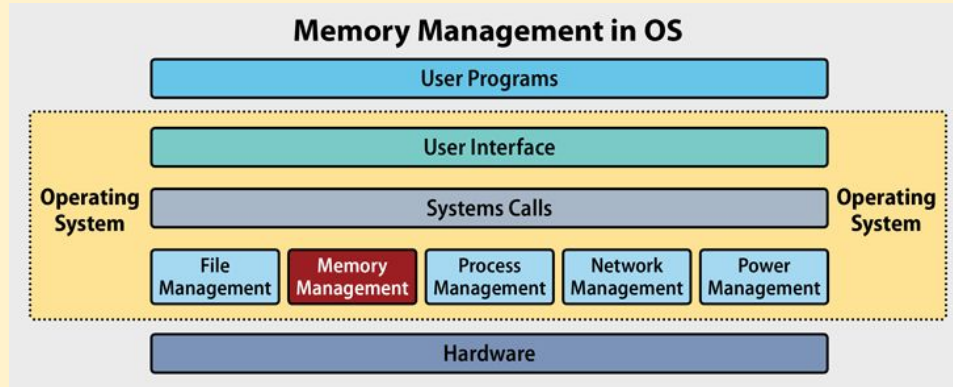
Operating system



- software can be application software
- or system software like the os
- the os works closely with the CPU to manage processes and memory
- provides the interface between hardware and humans
- applications must compile to a specific operating system


Memory management

- memory management swaps memory from disk to RAM and from RAM to internal CPU cache memory coordinating which processes get which memory blocks
- combines hardware, the operating system, the CPU and applications



I/O

- input devices: keyboard, mouse, scanner
- output devices: monitor, speakers, printer
- both input and output: disk drive, internet connection




I/O approaches

- polling
 - CPU polls devices to see if they need attention
- interrupt-based
 - CPU interrupted from its work by I/O devices, handles the task, and gets back to work on other things
- DMA direct memory access
 - devices have permission to read/write to RAM without CPU involvement
 - frees the CPU from data transfer tasks
- Memory-mapped I/O
 - portions of the address space are assigned to I/O devices
 - the communication to the device is over some protocol like PCIe

CPU chipsets


The diagram shows a chipset with two parts:

- the Northbridge contains the memory controller and is connected to the CPU(s) via the Front Side Bus
- the Southbridge handles communication with other devices




CPU chipsets

- Northbridge handles this without the CPU
- many devices now have DMA, direct memory access, but CPU now competing with DMA requests for RAM access



CPU chipsets


- CPU cores/threads have to wait for access
- workaround: connect Northbridge to many external memory controllers



Integrated memory controller

Beginning with AMD64 and Intel Nehalem, the memory controller was moved inside the CPU die


Starting with Intel "Sandy Bridge" (2011) all Northbridge functionality moved into the CPU



UMA v. NUMA

Now all CPUs have 'local' RAM for each core but still have access to all memory

the memory access time is now not uniform: NUMA Non-Uniform Memory Architecture



Ethernet card

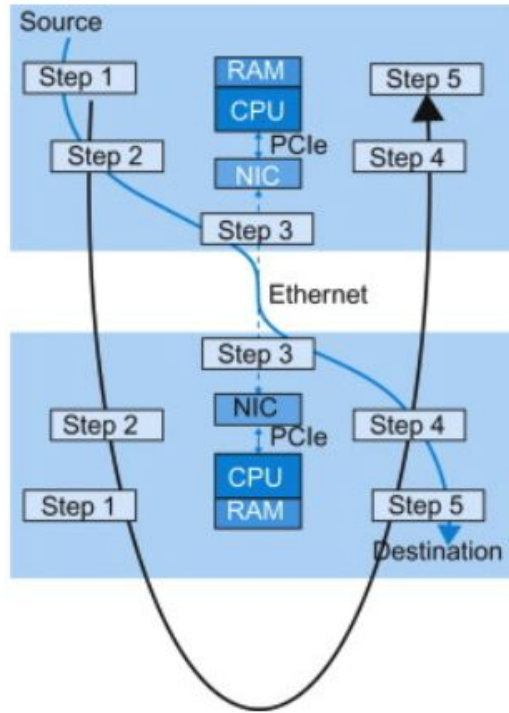


- a computer is connected to a network using a NIC network interface card
- every NIC has a hardware MAC (media access control) address that is given to the NIC when it is manufactured
- in contrast, IP addresses are associated with network software

Ethernet packets

Preamble		Addressing		Length	Data	CRC
Timing Frame		Destination	Source	Amount of data enclosed		Checksum
7 Bytes	1	6 bytes	6 bytes	2 bytes	46 to 1522 bytes	4 bytes
10101010101... alternating 1s & 0s						
SFD: 10101011						
		00:00:A1:FD:23:C9	00:00:A1:1E:04:12			

Figure 9.9.3: Relationship of the five steps of the driver when transmitting an Ethernet packet from one node and receiving that packet on another node (COD Figure e6.9.3).



sending and receiving packets

Transfer:

1. driver prepares a packet buffer in host memory; copies a packet from the user address space to the buffer in the os address space
2. driver writes an I/O descriptor to the NIC register; descriptor includes the address of the buffer and its length
3. The DMA in the NIC copies the outgoing Ethernet packet from the host buffer over PCIe
4. when transmission is complete, the DMA interrupts the processor with notification of successful transmission
5. the driver de-allocates the transmit buffer

Receive:

1. driver prepares a packet buffer in host memory
2. driver writes an I/O descriptor to the NIC register; descriptor has the buffer address and length
3. DMA in the NIC copies the incoming Ethernet packet over PCIe into the allocated host buffer
4. DMA interrupts the processor to notify of the reception of the packet and its size
5. driver copies the received packet into the user address space

