High Performance Computing

A simple machine model

Martin Raum

Computation and data on a machine

To obtain a first machine model, we can group concepts from basic programming as follows:

CPU Calculations, memory operators (e.g. memcpy), branching (if– and for–statements).

Memory Variables, arrays, allocated memory.

IO Terminal in- and output. File access.

The first one is procedural, emphasizing *what* is done.

The last two are data-centered, emphasizing *on what* is operated.

Computation and data on a machine

The connection between these domains, can be visualized as

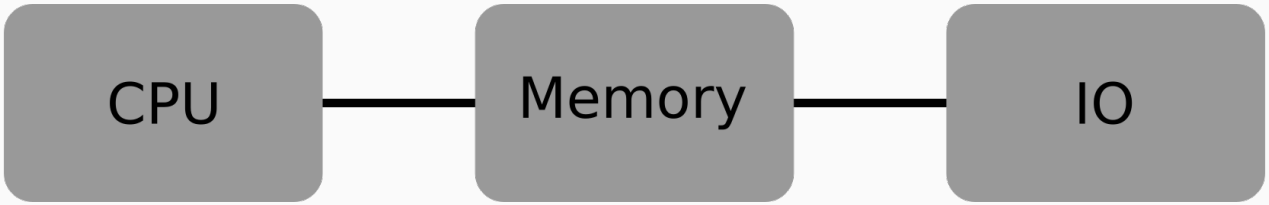

Hardware architecture is about refining this model and understanding its details.

There is no single truth about hardware architecture; it is descriptive. As hardware varies its architecture might vary.

Hardware architecture

Understanding hardware architecture will help us to

- understand behavior of programs in detail;
- pick most efficient algorithms and implementations;
- optimize for specific architectures.

CPU (central processing unit)

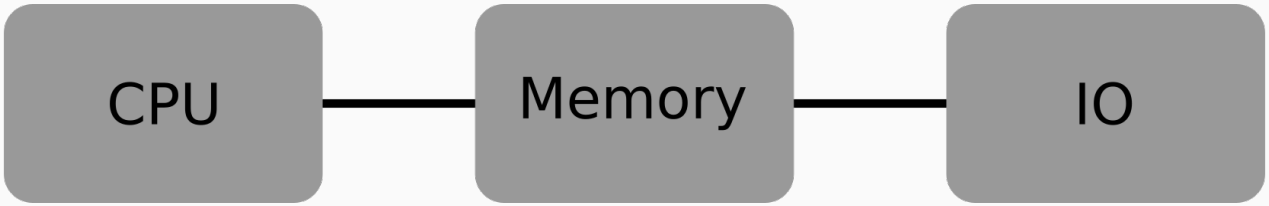

In most case, the CPU features at least

- the arithmetic and logic unit (ALU),
- the floating point unit (FPU), and
- the control unit (CU).

CPU–ALU

The arithmetic and logic unit performs:

arithmetic operations: $+$, -, *, / on integers;

logic operations: $+$(OR), *(AND) on boolean.

CPU–FPU

The floating point unit performs:

- basic arithmetic operations: $+$, -, *, $/$ of float and double;
- advanced arithmetic operations: sqrt, exp, sin, cos, etc.

CPU–CU

The control unit performs operations whose meaning we discuss later. For completeness, it

- fetches, decodes, preprocesses, and executes instructions;
- loads and stores data from memory and cache.

Memory (Cache and RAM)

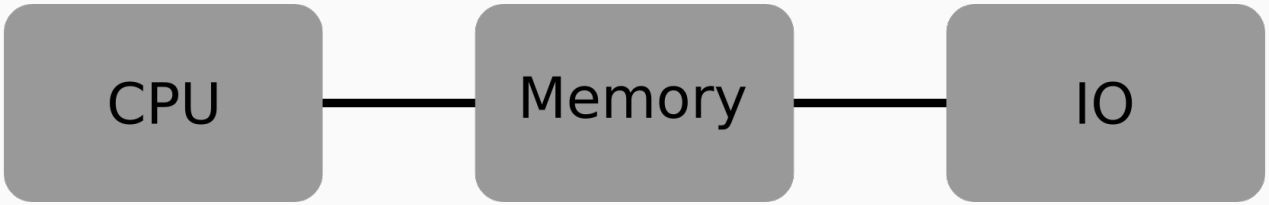

Memory stores and provides instructions and data.

Often “memory” refers to system memory (RAM), but there is also memory, called cache, that is associated closer to the CPU.

On modern systems areas in memory storing instructions and data may not overlap.

Storge units

Recall that $1\text{Byte} = 8\text{Bit}$. The following units are commonly used:

$$1\text{KiB} = 1024\text{Byte},\ 1\text{kB} = 1\text{KB} = 1000\text{byte},$$

$$1\text{MiB} = 1024\text{KiB},\ 1\text{MB} = 1000\text{KB},$$

$$1\text{GiB} = 1024\text{MiB},\ 1\text{GB} = 1000\text{MB},$$

$$1\text{TiB} = 1024\text{GiB},\ 1\text{TB} = 1000\text{GB}.$$

IO (Input/Output)

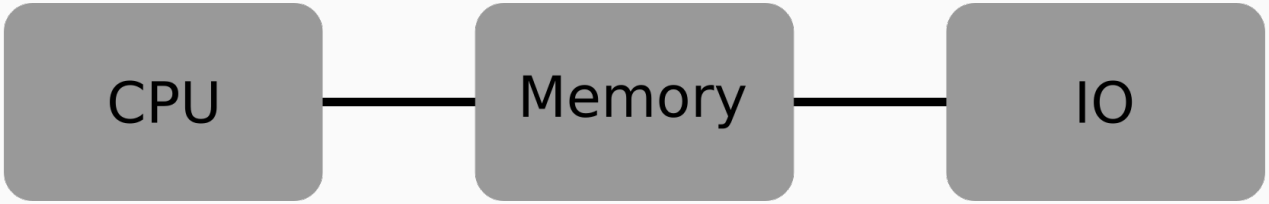

The input/output system covers:

screen (terminal), keyboards, printers;

disks (hard disk, solid state disk, magnet tape);

network.

Buses

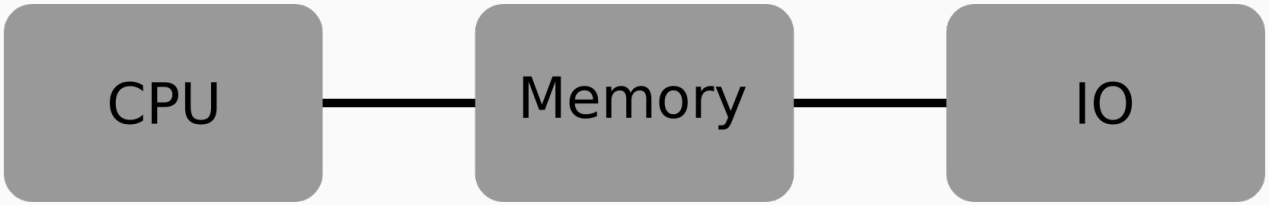

The connection between the CPU and memory is called the memory bus.

The connection between the memory and IO is an artifact of the simple model that we use.

Note that disks are attached indirectly to the CPU, insinuating that access is slower.