

High performance computing

Eugeniy E. Mikhailov

The College of William & Mary



Lecture 24

Notes

Energy considerations for memory storage

To store a bit, you need to change the energy state. Physics dictates that such energy change should be larger than the typical energy within the ambient. Since cosmic background radiation has a temperature of 2.725 K, we conclude that a bit should have energy difference between on and off states larger than

$$\Delta E_b = k_B T = 3.76 \times 10^{-23} \text{ J} = 0.235 \text{ meV}$$

Notes

Mass considerations for storage limit

A typical computer RAM memory size is about 10 GB (10^{10} bytes or $\sim 10^{11}$ bits).

If we flip all bits from 0 to 1, we need to deliver the energy change

$$\Delta E = \Delta E_b \times 10^{11} = 3.76 \times 10^{-12} \text{ J} \quad (1)$$

This according to $E = mc^2$ requires the mass change

$$\Delta m = \Delta E / c^2 = 4.18 \times 10^{-29} \text{ kg} \quad (2)$$

For comparison, the mass of the hydrogen atom is 1.67×10^{-27} kg. May be we should not worry about this limit. From the other hand, black holes might be alien's supercomputers.

Notes

Power considerations for limits to computing speed

Computing is energy shifting

$$P_{\text{consumed}} = f N_f \Delta E_b$$

f computer clock frequency

N_f number of bit flips per clock tick

A computer with typical modern day specs, which runs at room temperature (≈ 300 K) with the clock frequency 1 GHz and fully flips a 64 bit word per cycle, should consume 240 pW.

Based on this, we can expect about 12 orders of magnitude speed up for a computer consuming 200 W.

But, **there is a limit** at the end.

Notes

Practical limits due to heat dissipation requirement

Previous slides show the fundamental limits but, in practice, the main problem is the cooling of computers. They generate too much heat and, thus, consume a lot of energy. According to physics, radiation from the accelerating charge and, thus, the required computer power

$$P_{\text{radiated}} \sim f^4$$

So for the same architecture, the cooling system has to be adjusted accordingly.

- Increase of the clock frequency by the factor of 2 leads to the power consumption increase by the factor of 16.

For the last several years, the CPU clock frequency has not grown much beyond 4 GHz.

Notes

From single CPU to many

Because of energy and price considerations, it is impractical to boost the speed (i.e., the clock frequency) of a single CPU. An alternative approach is needed. Multiple interconnected computing nodes: **Clusters**

Node is a computer which might have a single or multiple cores (CPUs)

Such nodes/cores can be connected either in pipe lines (serially) or in parallel.

The pipe line approach is very suitable for a GPU, since they do series of relatively simple operations over different data sets over and over. The high end video cards have thousands of pipelines.

The parallel approach is used for problems which can be run independently on different nodes (think about Monte Carlo simulations).

Notes

Conditions for parallel problem

Not all problems can be parallelized. For example, a general ODE solver is naturally sequential. The state at the time t depends on the state at the time $t - \Delta t$, so you cannot calculate it independently, i.e., without obtaining the previous time state. Recall a planetary motion problem, you cannot calculate a planet position independently without taking in account other bodies.

Such problem are called **synchronous**.

However, there are problems which can be done in parallel, i.e., **asynchronous problems**.

Think about matrix multiplication

$$\vec{C} = \mathbf{A}\vec{B} \rightarrow C_i = \sum_{j=1}^N A_{ij}B_j$$

There is the "vector" computer term, since the very first super computers were designed to do such matrix products in the parallel fashion.

Notes

Speed up vs. latency

Even in the best case, where we have a problem which can be equally split to N cores, the speed up does not scale with N .

$$S(N) = \frac{T_1}{T_1/N + \tau(N)}$$

S the speed up

N the number of nodes

T_1 the execution time for a problem on a single node

τ latency or communication time, it usually grows with the number of nodes

There is no point to infinitely increase the number of nodes.

Notes

Human supercomputer

Solve quadratic equation

$$ax^2 + bx + c = 0$$

for multiple a , b , and c inputs.

Notes

Human supercomputer

Solve quadratic equation

$$ax^2 + bx + c = 0$$

for multiple a , b , and c inputs.

a	b	c	x_1	x_2
1	2	1		
1	0	-1		
0	-2	2		
2	-6	-20		
3	3	-18		

Notes

Human supercomputer

Solve quadratic equation

$$ax^2 + bx + c = 0$$

for multiple a , b , and c inputs.

a	b	c	x_1	x_2
1	2	1	-1	-1
1	0	-1	1	-1
0	-2	2	1	NaN
2	-6	-20	5	-2
3	3	-18	-3	2

Notes

MATLAB's parallel tools

parfor vs. for

```
N=5e5;
for i=1:N
    c=inv(rand(10));
end
```

```
% pool=parpool(2)
N=5e5;
parfor i=1:N
    c=inv(rand(10));
end
% delete(pool)
```

Important: for every loop iteration, the evaluation of the `for` body should not depend on previous calculations in the loop body.

Notes

See also `batch`

Human supercomputer