High-Performance Mathematics

Objectives, Program and Organization

Progetto Speciale per la Didattica 2023/24

Fabio Durastante (Lo)

March 27, 2024

Dipartimento
di Matematica
Università di Pisa

Table of Contents

Scientific computing

"**Computational science** (also **scientific computing** or **scientific computation** (SC)) is a rapidly growing multidisciplinary field that uses advanced computing capabilities to *understand and solve complex problems*. It is an area of science which spans many disciplines, but at its core it involves the development of *models and simulations to understand natural systems*."

Wikipedia

Leonardo, CINECA

Green Data Center, UNIPI

What are the applications?

- Computational finance,
- Computational biology,
- Simulation of complex systems,
- Network analysis
- Multi-physics simulations,
- Weather and climate models,
- Artificial Intelligence,
- ...

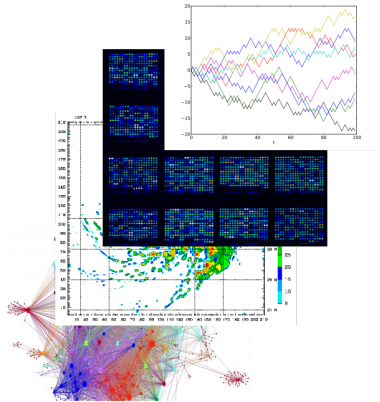

What are the applications?

- Computational finance,
- Computational biology,
- Simulation of complex systems,
- Network analysis
- Multi-physics simulations,
- Weather and climate models,
- Artificial Intelligence,
- ...

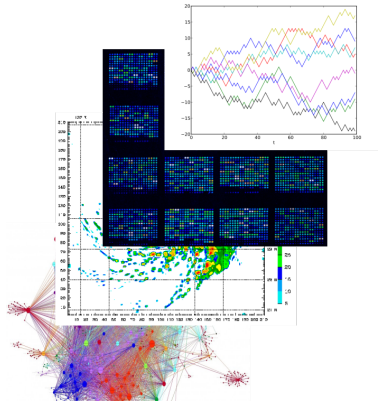

Why the need for parallelism?

Moore's law

1 High-Performance Mathematics: why?

"The complexity for minimum component costs has increased at a rate of roughly a factor of two per year. Certainly over the short term this rate can be expected to continue, if not to increase. Over the longer term, the rate of increase is a bit more uncertain, although there is no reason to believe it will not remain nearly constant for at least 10 years."

G. Moore, 1975

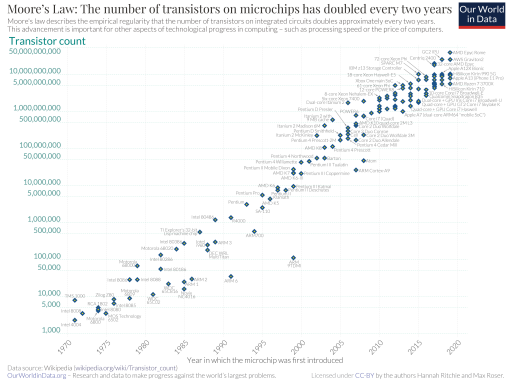

Moore's law

1 High-Performance Mathematics: why?

"The complexity for minimum component costs has increased at a rate of roughly a factor of two per year. Certainly over the short term this rate can be expected to continue, if not to increase. Over the longer term, the rate of increase is a bit more uncertain, although there is no reason to believe it will not remain nearly constant for at least 10 years."

G. Moore, 1975

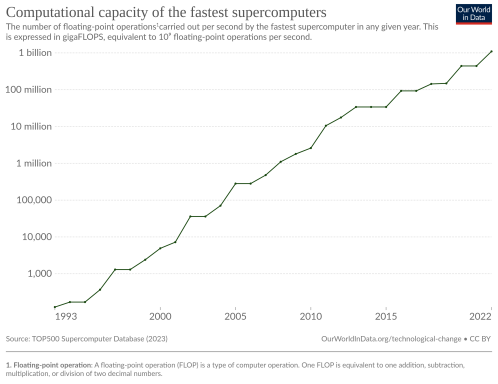

Computers *should* reach the physical limits of Moore's Law at some point in the 2020s…exponential functions saturates physical capabilities!

Parallel computing: why?

- We are hitting the wall of single processor transistor count/computing capabilities,

Parallel computing: why?

- We are hitting the wall of single processor transistor count/computing capabilities,
- Some applications needs more memory than the one that could be available on a single machine,

Parallel computing: why?

- We are hitting the wall of single processor transistor count/computing capabilities,
- Some applications needs more memory than the one that could be available on a single machine,
- Optimization of sequential algorithms can bring us only to a certain extent

Parallel computing: why?

- We are hitting the wall of single processor transistor count/computing capabilities,
- Some applications needs more memory than the one that could be available on a single machine,
- Optimization of sequential algorithms can bring us only to a certain extent

“διαίρει καὶ βασίλενε“

(diáirei kái basíleue)

Parallel computing: why?

- We are hitting the wall of single processor transistor count/computing capabilities,
- Some applications needs more memory than the one that could be available on a single machine,
- Optimization of sequential algorithms can bring us only to a certain extent

“διαίρει καὶ βασίλενε“

(diáirei kái basíleue)

Dividi et Impera

Parallel computing: why?

- We are hitting the wall of single processor transistor count/computing capabilities,
- Some applications needs more memory than the one that could be available on a single machine,
- Optimization of sequential algorithms can bring us only to a certain extent

"διαίρει καὶ βασίλενε"
(diáirei kái basíleue)
Dividi et Impera

Therefore, we need

- Algorithms that can work in parallel,
- A communications protocol for parallel computation integrated with our programming languages,
- Parallel machines that can actually run this code.

The philosophy behind the effort

1 High-Performance Mathematics: why?

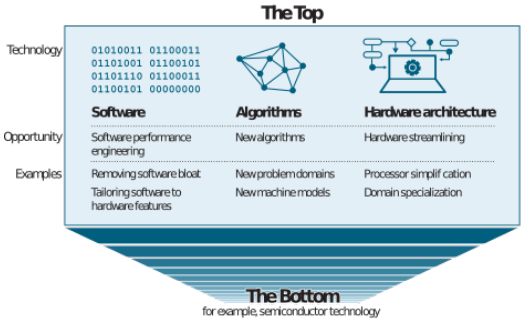

C. E. Leiserson, N. C. Thompson, J. S. Emer, B. C. Kuszmaul, B. W. Lampson, D. Sanchez, and T. B. Schardl, "There's plenty of room at the Top: What will drive computer performance after Moore's law?", *Science* (2020)

"As miniaturization wanes, the silicon-fabrication improvements at the Bottom will no longer provide the predictable, broad-based gains in computer performance that society has enjoyed for more than 50 years. Software performance engineering, development of algorithms, and hardware streamlining at the Top can continue to make computer applications faster in the post-Moore era."

Table of Contents

2 Yes, but why?

Yes, but why?

- It is good fun!
- It can be spent on the **programming/project part** of “Calcolo Scientifico”, “Laboratorio Computazionale” and “Metodi Numerici per le PDE” **exams**.
- It teaches you to use various computational resources of the Department of Mathematics, resources that can be used to carry out *thesis projects* even in disciplines *other than Numerical Analysis*.
- If you end up liking it, it can give you the opportunity to write **theses in computational mathematics** that are **very close to applications** and **current research**.

Table of Contents

3 High-Performance Mathematics: what?

Parallel computing: what?

What are the **topics** we will cover in these lectures?

1. Abstract models of parallel machines and their operation,

Parallel computing: what?

What are the **topics** we will cover in these lectures?

1. Abstract models of parallel machines and their operation,
2. Extension and improvement of the parallel computer prototype assembled in the previous edition,

Parallel computing: what?

3 High-Performance Mathematics: what?

What are the **topics** we will cover in these lectures?

1. Abstract models of parallel machines and their operation,
2. Extension and improvement of the parallel computer prototype assembled in the previous edition,
3. Laboratory introduction to programming techniques for the implementation of parallel algorithms,

Parallel computing: what?

What are the **topics** we will cover in these lectures?

1. Abstract models of parallel machines and their operation,
2. Extension and improvement of the parallel computer prototype assembled in the previous edition,
3. Laboratory introduction to programming techniques for the implementation of parallel algorithms,
   - MPI,
   - OpenMP,
   - NVIDIA/CUDA,

Parallel computing: what?

What are the **topics** we will cover in these lectures?

1. Abstract models of parallel machines and their operation,
2. Extension and improvement of the parallel computer prototype assembled in the previous edition,
3. Laboratory introduction to programming techniques for the implementation of parallel algorithms,
   - MPI,
   - OpenMP,
   - NVIDIA/CUDA,
4. Use of what has been discussed for the solution of selected applied mathematics problems:

Parallel computing: what?

3 High-Performance Mathematics: what?

What are the **topics** we will cover in these lectures?

1. Abstract models of parallel machines and their operation,
2. Extension and improvement of the parallel computer prototype assembled in the previous edition,
3. Laboratory introduction to programming techniques for the implementation of parallel algorithms,
   - MPI,
   - OpenMP,
   - NVIDIA/CUDA,
4. Use of what has been discussed for the solution of selected applied mathematics problems:
   - Solution of large and sparse linear systems,
   - Numerical solution of PDE/ODE,
   - Training and inference with neural networks.

Table of Contents

4 High-Performance Mathematics: how?

Parallel computing: how?

4 High-Performance Mathematics: how?

The idea is to

- organize **one meeting per week** lasting a **couple of hours**,
- have **practical programming sessions** in which **problem solving** is tackled,

employing:

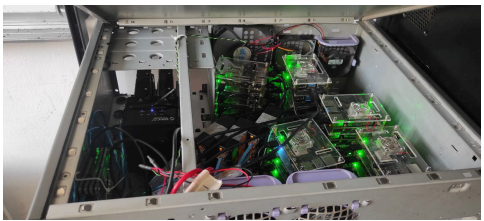

The `steffe` cluster.

The `NVIDIA T1000` GPUs in `Aula 3`

Table of Contents

5 High-Performance Mathematics: when?

Parallel computing: when?

5 High-Performance Mathematics: when?

The **QR code** leads to a `framadate` to schedule the next lecture.