



Tools of computational physics

A. Godunov

1. Major tools of computational physics
2. Hardware
3. Software
4. Programming languages
5. Libraries (numerical, plotting, codes)

1


Part 1: Major parts

2

Computational Physics

Computing has become a central tool in physics!

From R. Landau "Computational Physics" (2015)



- Physics: methods for computer simulations
- Applied math: solving numerical problems
- Computer science: hardware and software (subject of this lecture)

3

Computer science

The basic ideas behind computational physics are hardware and software independent

However, for solving a problem in hand one may need to consider advantages and disadvantages various choices.

Critical or important

- Development time
- Computational time
- Computational memory (RAM)
- Storage memory

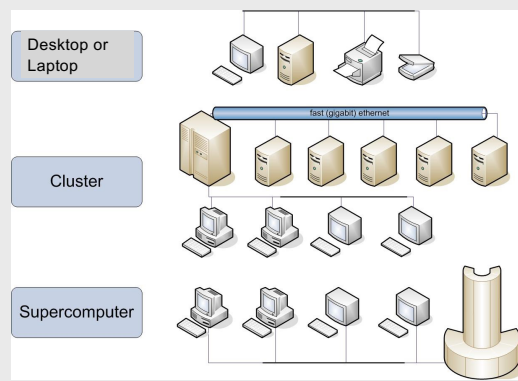
Use right tools for your problem!

4

Part 2: Tools: Hardware

5

Hardware 1.



6

Internal hardware

Most important for computing

- CPU(s) - central processing unit (how many + how fast)
also - cache memory: cache 1, cache 2
- RAM -random-access memory (GB)
communication with CPU by bus
- GPU – Graphic Processor Unit (how many + how fast)
Can speed up calculations considerably!
- Network Interface (MB/GB per sec)
- HDD – Hard Disk Drive (TB)

7

Top 500

The TOP500 lists the 500 most powerful commercially available computer systems <https://top500.org>

TOP500 LIST - JUNE 2021

R_{max} and R_{min} values are in TFlops. For more details about other fields, check the TOP500 description. R_{max} values are calculated using the advertised clock rate of the CPU. For the efficiency of the systems you should take into account the Super-CPU clock rate when applicable.

Rank	System	Cores	Max TFlops	Rank TFlops	Power kW
1	Supercomputer Fugaku - Supercomputer Fujitsu, A64FX-4C2 2.25Peta, Taka International Co., Fujitsu RIKEN Center for Computational Science Japan	7,430,848	462,070.0	527,372.0	29,891
2	Summit - IBM Power System AC920, IBM POWER9 20C 2.25Peta, NVIDIA Tesla V100S Dual-in-Meckel EDR Infinidat, IBM ORNL Center for Computational Science United States	2,414,592	148,400.0	200,794.9	10,094
3	Sierra - IBM Power System AC920, IBM POWER9 20C 3.10Peta, NVIDIA Tesla V100S Dual-in-Meckel EDR Infinidat, IBM NVIDIA / Mellanox ORNL Center for Computational Science United States	1,972,480	94,460.0	120,712.0	7,428
4	Summit TaihuLight - Sunway SP7 Sunway SW62010 20C 1.6Peta, Sunway, NXP National Supercomputing Center in Wuji China	10,649,480	10,074.4	120,400.9	10,371

TOP500 LIST - JUNE 2022

R_{max} and R_{min} values are in PFlops. For more details about other fields, check the TOP500 description. R_{max} values are calculated using the advertised clock rate of the CPU. For the efficiency of the systems you should take into account the Super-CPU clock rate when applicable.

Rank	System	Cores	Max PFlops	Rank PFlops	Power MW
1	Frontier - HPZ Epyc 7002, AMD Optimized for Datacenter EPYC 7002, AMD Infinity MCOSE, Singapore 11, HPE ORNL Center for Computational Science United States	8,730,112	1,102.00	1,005.65	21,103
2	Supercomputer Fugaku - Supercomputer Fujitsu, A64FX-4C2 2.25Peta, Taka International Co., Fujitsu RIKEN Center for Computational Science Japan	7,430,848	442.01	527.21	29,891
3	Summit - IBM Power System AC920, IBM POWER9 20C 2.25Peta, NVIDIA Tesla V100S Dual-in-Meckel EDR Infinidat, IBM ORNL Center for Computational Science United States	2,414,592	148.40	200.79	10,094
4	Sierra - IBM Power System AC920, IBM POWER9 20C 3.10Peta, NVIDIA Tesla V100S Dual-in-Meckel EDR Infinidat, IBM NVIDIA / Mellanox ORNL Center for Computational Science United States	1,972,480	94.46	120.71	7.428
5	Summit TaihuLight - Sunway SP7 Sunway SW62010 20C 1.6Peta, Sunway, NXP National Supercomputing Center in Wuji China	10,649,480	10.07	120.40	10,371

8

Part 3: Software

9

Software

Principal parts

1. Operating system
2. Ideally IDE integrated development environment + language(s)
 - a) Editor
 - b) Compiler(s) (for high-level languages)
 - c) Libraries (modules, toolboxes,...)
3. *Optional:*
 - a) Lint and other static code analyzers
 - b) Debugger
 - c) Profilers
 - d) Report generators

10

Two major families

1. Unix and Unix-like
 - a) Unix - commercial versions: AIX, HP-UX, Solaris, ...
 - b) Apple - macOS
 - c) Linux – many distributions
2. Windows

Share of operating systems for desktops/laptops



Share of operating systems for clusters/supercomputers:
Linux now runs on all the fastest 500 supercomputers in the world

11

UNIX family

Developed in around 1970 in Bell Labs research center

1. Powerful beyond imagination.
2. ALL the Top 500 supercomputers in the world run on Linux.
3. Linux is based on Unix. macOS is based on Unix.
4. Robust and small kernel.
5. Very safe: sandboxing and rich file permission system.
6. Plenty of tools
editors, programming languages, ...

Philosophy of Unix/Linux: "Building blocks" + "glue"

- Building blocks: programs do only one thing, but do it well
- Glue: easy combine various blocks.

12

Windows

Generally available since 1992

1. Most popular OS in the world for personal computers and laptops.
2. The latest version is Windows 11.

13

Part 4: Software: Programming languages

14

Programming languages

Important questions

1. Which language to learn?
2. Which language to use?
3. Do I need to learn new language(s)?

Most common in physics

- C and C++ (current standard C++20)
- Python
- Fortran
- Matlab
- *Julia*
- *R*

15

Computer languages

Three basic modes to run a code

1. Interpreted: Python, Matlab, Mathematica, R.
2. Compiled: Fortran, C/C++.
3. JIT (Just-in-Time) compilation: Julia

Interpreted languages can we used with:

1. A command line
2. A script file.

16

C

Created in the early 1970s by Dennis Ritchie at Bell Labs,

- General purpose, multi-paradigm, **compiled** language
- By design, C's features reflect the capabilities of the targeted CPUs
- It was designed to be compiled to provide low-level access to memory and language constructs that map efficiently to machine instructions
- The language was primary design for writing operating systems
- Many languages have based directly or indirectly on C, including C++, C#, Java, JavaScript, Julia, Perl, PHP, Python, Ruby, Swift, ...

17

C++

Developed by Bjarne Stroustrup at Bells Labs in the early 1980s

- General purpose, multi-paradigm, **compiled** language
- C/C++ is the infrastructure of much of the modern computing world.
- Powerful language: you can code anything in C++
- Easy integration with multiprocessor programming OpenMP, MPI, CUDA, OpenCL, ...
- If you know Unix/Linux and C/C++, you can master everything else
- Excellent compilers (including open-source) and tools.
- Top performance in terms of speed.

18

C++

Some disadvantages

- It was not designed for scientific calculations!
- Hard language to learn, even harder to master
- Large specification: C++20 (This causes, at times, portability issues)
- **Matrix indexing starts at zero**

19

FORTRAN (formula translator)

Grandfather of all modern languages – developed in 1957 (IBM)

At the beginning almost all computing was for physicists!

- General purpose, multi-paradigm, **compiled** language
- Last version Fortran 2018
- Lot of high-quality libraries (both numerical and applications)
- Still widely used in science in engineering
weather forecast, nuclear weapon research and development, ...
- Easy to learn, portable, nice array support, easy to parallelize
- Generally available on clusters and supercomputers

20

FORTRAN (formula translator)

Some disadvantages

- Small community of users
- Most Fortran compilers are proprietary

21

Python

Designed by Guido von Rossum around 1991

- General purpose, multi-paradigm, **interpreted** language
- Open source
- Intuitive – easy to learn
- Scientific computation modules: NumPy, SciPy, and SymPy
- Plotting modules: matplotlib and ggplot.
- Preinstalled on many systems (e.g. macOS)

22

Python

Some disadvantages

- Considerable time penalty
- Python's memory usage is high
- Python's functional programming can be difficult to read
- **Runtime Errors:** One of the major drawbacks of this language is that its design has numerous issues

23

MATLAB

Started in the late 1970s, released commercially in 1984.

<https://www.mathworks.com/products/matlab.html>

- General purpose, multi-paradigm, **interpreted** language
- Widely used in engineering and industry
- Plenty of codes around for science, engineering and economics.
- Many useful toolboxes
- Great IDE (Integrated Development Environment)
- Interacts reasonably well with C/C++, Fortran, and R

Some disadvantages

- Can be expensive
- Tight integration with Java

24

Other languages

Julia

- Modern, high-performance programming language designed for scientific computation and data manipulation.
- Designed for parallelism and cloud computing. Syntax close to Matlab. However, at early stages of life (can be unstable)

R

- High level, open-source language for statistical computation
- Widely used for big data, easy to parallelize

Mathematica

- Mainly oriented toward symbolic computation
- Programming approach is different from other languages.

And more: C#, Javascript, PHP, Perl, Swift, Ruby, ...

25

Compare times of calculation ...

Results depends on a model/test but here are *some* average numbers

C++	1.00
Fortran	0.90
Python	50.0
Matlab	10.0
Mathematica	from 4.0 to 900
Julia	3.0
R	250

26

Summary

- **C++** good to learn (most powerful general programming language) if you master C++ you can quickly learn anything else.
- **Fortran** very powerful but learn only if needed (legacy codes or libraries)
- **Python** easy to learn, open source, but generally much slower than C++ and Fortran
- **MATLAB** easy to learn, convenient with great IDE, making graphs, multiple toolboxes available
- **Mathematica** good problem-solving environment but programming approach is different from other languages
- **Java** – rather no, unless the use of Virtual machine is important

27

C++ compilers and IDEs

Microsoft Visual C++ compiler

<https://visualstudio.microsoft.com/vs/features/cplusplus/>
Windows (IDE included)

Xcode (from Apple)

<https://developer.apple.com/xcode/>
macOS (IDE included)

Intel C++ compiler

www.intel.com/content/www/us/en/developer/tools/oneapi/dpc-compiler.html
Windows, macOS, Linux (works with Microsoft Visual Studio)

Dev-C++

<http://www.bloodshed.net>
Windows (IDE included) (open source)

28

Fortran compilers and IDEs

ABSoft

<https://www.absoft.com>
Windows, macOS, Linux (IDE included)

NAG (Numerical Algorithmic Group)

<https://www.nag.com/content/nag-fortran-compiler>
Windows, macOS, Linux (IDE included)

Intel Fortran compiler

www.intel.com/content/www/us/en/developer/tools/oneapi/fortran-compiler.html
Windows (IDE: Microsoft Visual Studio), macOS (IDE: Xcode),
Linux (IDE: Eclipse)

Most clusters and supercomputers have Fortran!

29

Python

Python.org

<https://www.python.org>
Windows, macOS

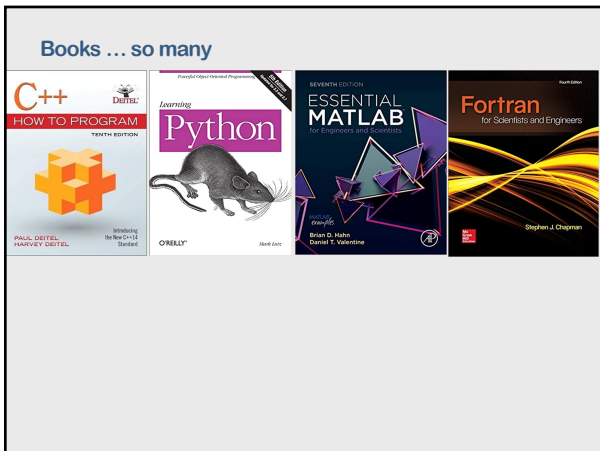
Anaconda

<https://www.anaconda.com/products/individual>
Windows, macOS

IDEs (for Windows, macOS and Linux)

- IDLE <https://docs.python.org/3/library/idle.html>
- Visual studio <https://visualstudio.microsoft.com/vs/features/python/>
- Spyder <https://www.spyder-ide.org>
- Atom <https://atom.io>
- And many more ...

30



31

Example: a circle using C++

```
// calculation: the diameter, circumference, and area of a circle
#include <iostream>
#include <iomanip>
#include <cmath>
using namespace std;
int main()
{
    const double pi=3.1415926;
    double radius, diameter, circumference, area;
    cout << "enter radius as float " << endl;
    cin >> radius;
    diameter = 2.0*radius;
    circumference = 2.0*pi*radius;
    area = pi*pow(radius,2);
    cout.setf(ios::fixed | ios::showpoint);
    cout.width(10);
    cout.precision(5);
    cout << "radius = " << radius << endl;
    cout << "diameter = " << diameter << endl;
    cout << "circumf. = " << circumference << endl;
    cout << "area = " << area << endl;
    return 0;
}
```

32

Example: a circle using MATLAB

```
% calculation: the diameter, circumference, and area
% of a circle with a given radius

Pi = 3.1415926;

prompt = 'Enter radius of a circle \n';
radius = input(prompt);

diameter = 2.0*radius;
circumference = 2.0*Pi*radius;
area = pi*radius*radius;

fprintf(' radius      %8.4f \n',radius);
fprintf(' diameter   %8.4f \n',diameter);
fprintf(' circumference %8.4f \n',circumference);
fprintf(' area         %8.4f \n',area);

%end
```

33

Example: a circle using Python

```
# -*- coding: utf-8 -*-
""" From "COMPUTATIONAL PHYSICS" & "COMPUTER PROBLEMS in PHYSICS"
    by RH Landau, et all."""

# Area.py: Area of a circle, simple program
from math import pi

N = 1
r = 1.3
C = 2. * pi * r
A = pi * r**2

print ('Program number =', N, '\n r, C, A = ', r, C, A)
```

34

Example: Fortran – Fibonacci prime numbers

```
program fibonacci
! the program generates Fibonacci numbers and chooses only prime numbers
! f(0) = 0
! f(1) = 1
! f(n) = f(n-1) + f(n-2) for n>1
implicit none
integer :: f(0:100)
integer :: i, j
character :: prime*5
f(0) = 0
f(1) = 1
do i=2,40
    f(i) = f(i-1) + f(i-2)
! check for prime numbers
    prime = 'prime'
    do j=2,f(i)-1
        if (f(i) == (f(i)/j)*j) then
            prime = ' '
            exit
        end if
    end do
    write (*,102) i, f(i), prime
end do
102 format(i3, i12, a6)
stop
end
```

35

- ### Other items
1. Version Control
 2. Backups
 3. Dynamic notebooks (Jupyter, Markdown, ...)

36

Part 4b: High Performance Computing (HPC)

37

High-performance computing (HPC)

Deals with scientific problems that require substantial computational power.

Usually, but not always, HPC involves the use of several processors:

- Multi-core/many-core CPUs (in a single machine or networked).
- Many-core coprocessors.
- GPUs (graphics processing units).
- TPUs (tensor processing units).
- FPGAs (field-programmable gate arrays)

"Amateurs talk about the speed of their processors, but professionals study coding techniques" from Gen. Robert H. Barrow, USMC (27th Commandant of the US Marine Corps)

38

High-performance computing (HPC)

Resources

- Livermore National Lab <https://hpc.llnl.gov/training/>
- HPC carpentry <https://www.hpc-carpentry.org>
- More <https://pages.tacc.utexas.edu/~eijkhout/istc/istc.html>
- ...

And many books

39

Part 5: Libraries

40

You should not reinvent the wheel

- Computational Physics Libraries
- Numerical Libraries and Depositories
- Large packages (example: the COMSOL Multiphysics)

www.odu.edu/~agodunov/computing/lib_net.html

Note 1: do not use routines as black boxes without understanding

Note 2: Collect you own library!



41

Few quotes

"Spend your intellectual energies on the current problem - not on fancy tools. When the volume and sophistication of your problems demand these weapons you will know it. That is the time to learn a new tool - and learn it by re-doing an already-solved problem, not a new one." *F.S. Acton*

"I realized that a large part of my life from then on was going to be spent in finding mistakes in my own programs."

Maurice Wilkes, after the first attempts to write programs for the EDSAC computer

"I conclude that there are two ways of constructing a software design: One way is to make it so simple that there are obviously no deficiencies and the other way is to make it so complicated that there are no obvious deficiencies."

Charles Hoare, inventor of the QuickSort algorithm, in his 1990 ACM Turing Award Lecture

42