CSC5001 - High Performance Systems

Elisabeth Brunet – Raymond Namyst

Parallel architectures are everywhere

Laptop, phone, common PCs, …

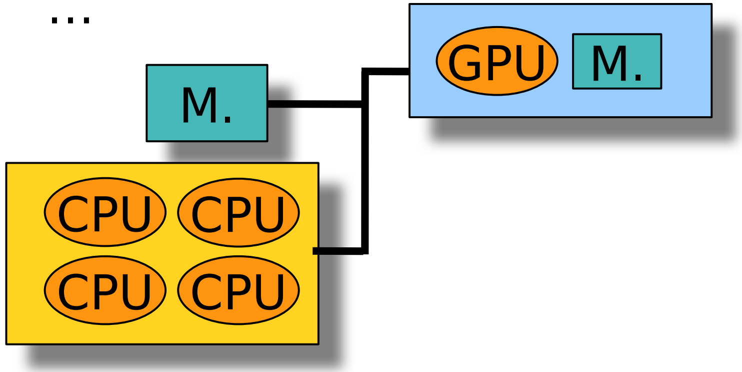

…, embedded components , …

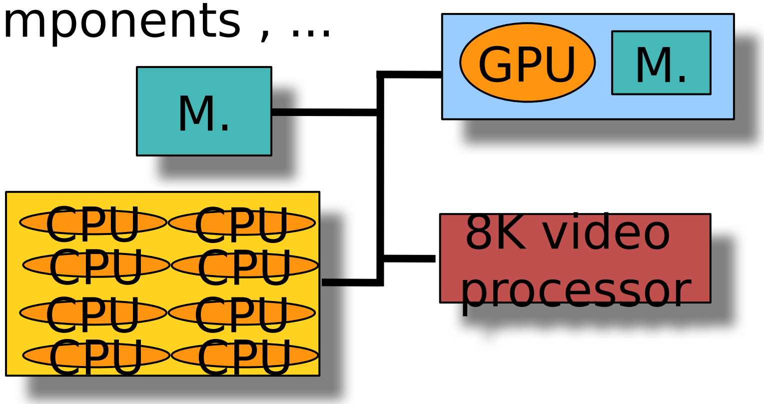

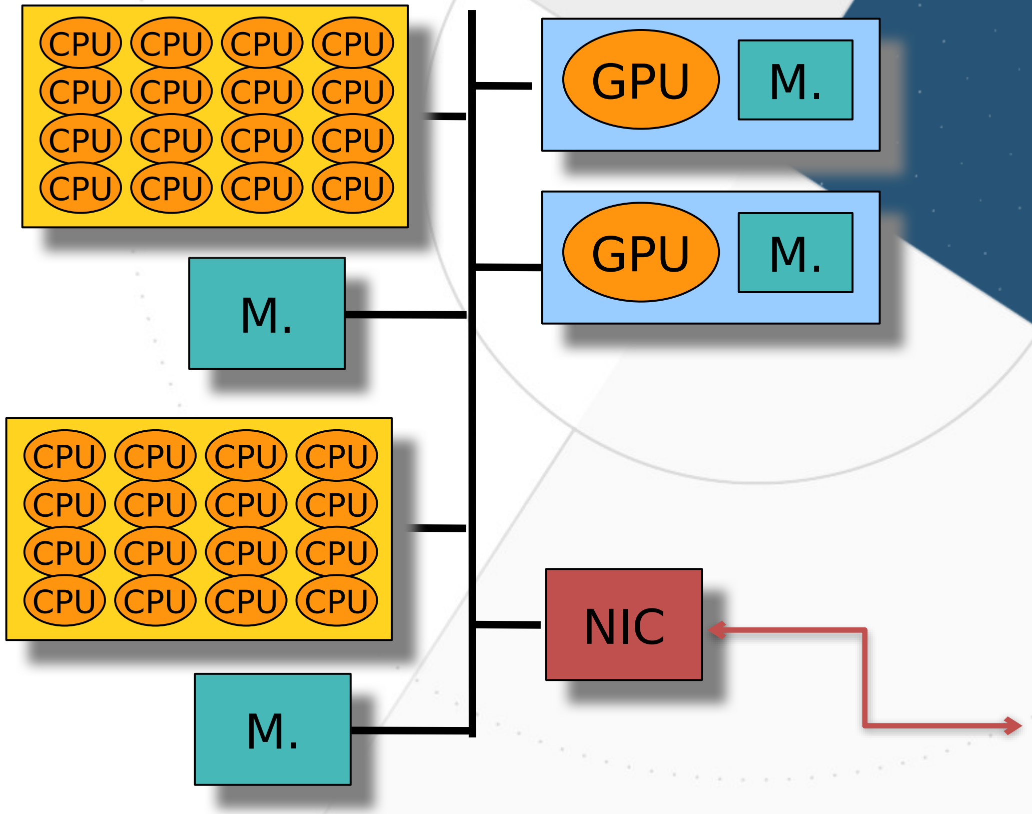

…, PC « gamer », computing servers, supercomputers.

And become more and more heterogeneous

- Accelerators : GPU, FPGA, ...
- Heterogeneous CPUs
- Unified memory architectures
  - Nvidia Grace Hopper, AMD APU MI

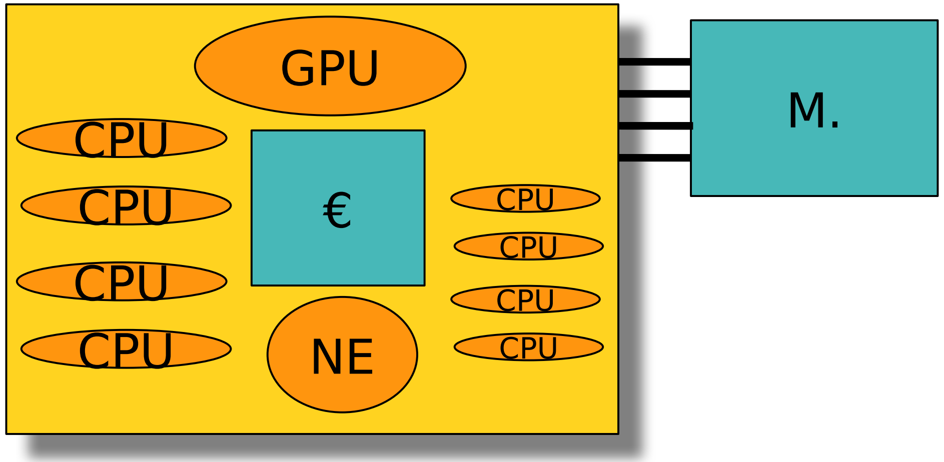

Apple M1 processor embeds :
- 4 general-purpose HP cores
- 4 low consumption cores
- 1 GPU with 8 pipelines
- 1 *Neural Engine with* 16 cores

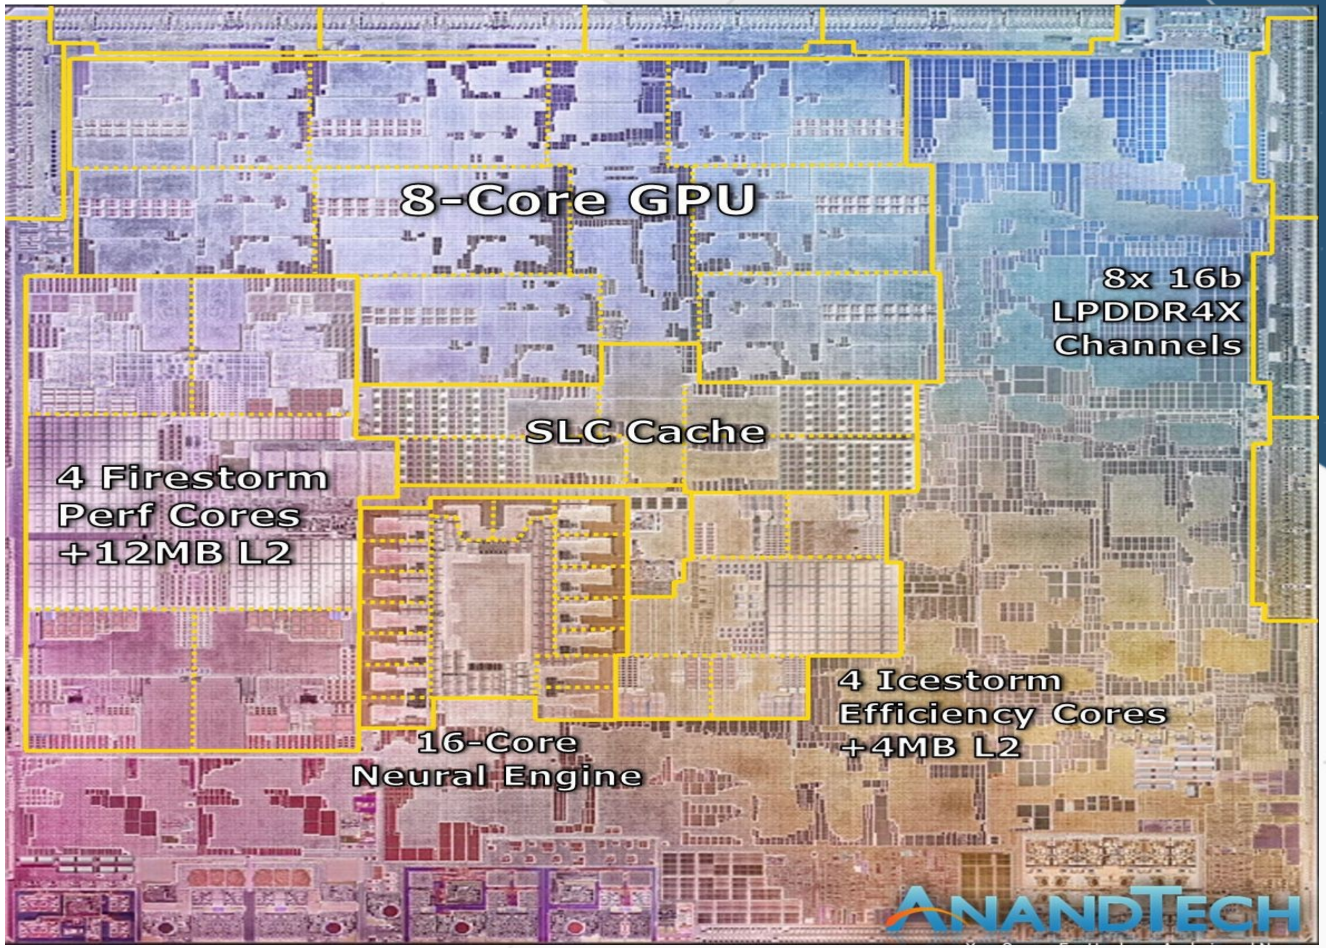

M1 processor

Why high performance parallel computing ?

- Handling large problems (scientific simulation, big data)
  - Gain in precision (simulation, video games, learning, etc.)
  - Perform more tests (cryptanalysis)
  - Make faster decisions than others (finance)

High-performance simulation has become

- The third pivot of science:
  - experimentation / mathematical modeling / simulation
- A must-have for major industrial groups:
  - Simulation of physical/chemical phenomena
  - Digital mock-up for design and manufacturing

The democratization of high-performance simulation is one of the keys to our industrial future.

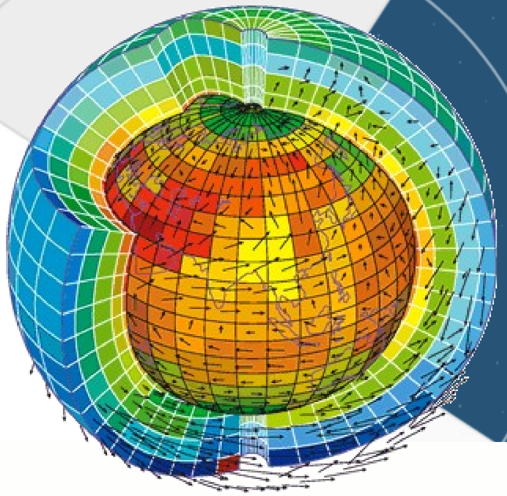

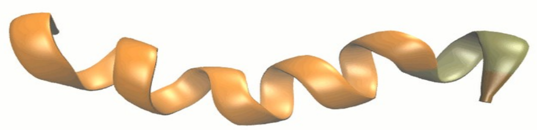

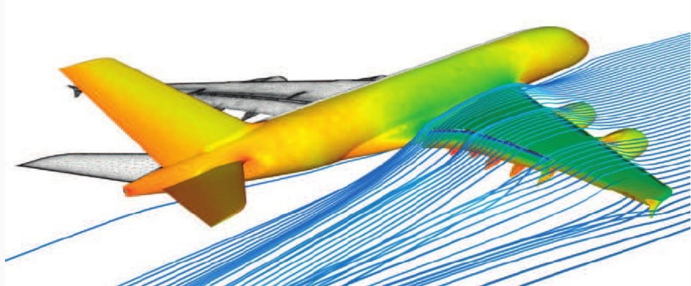

TOP 500

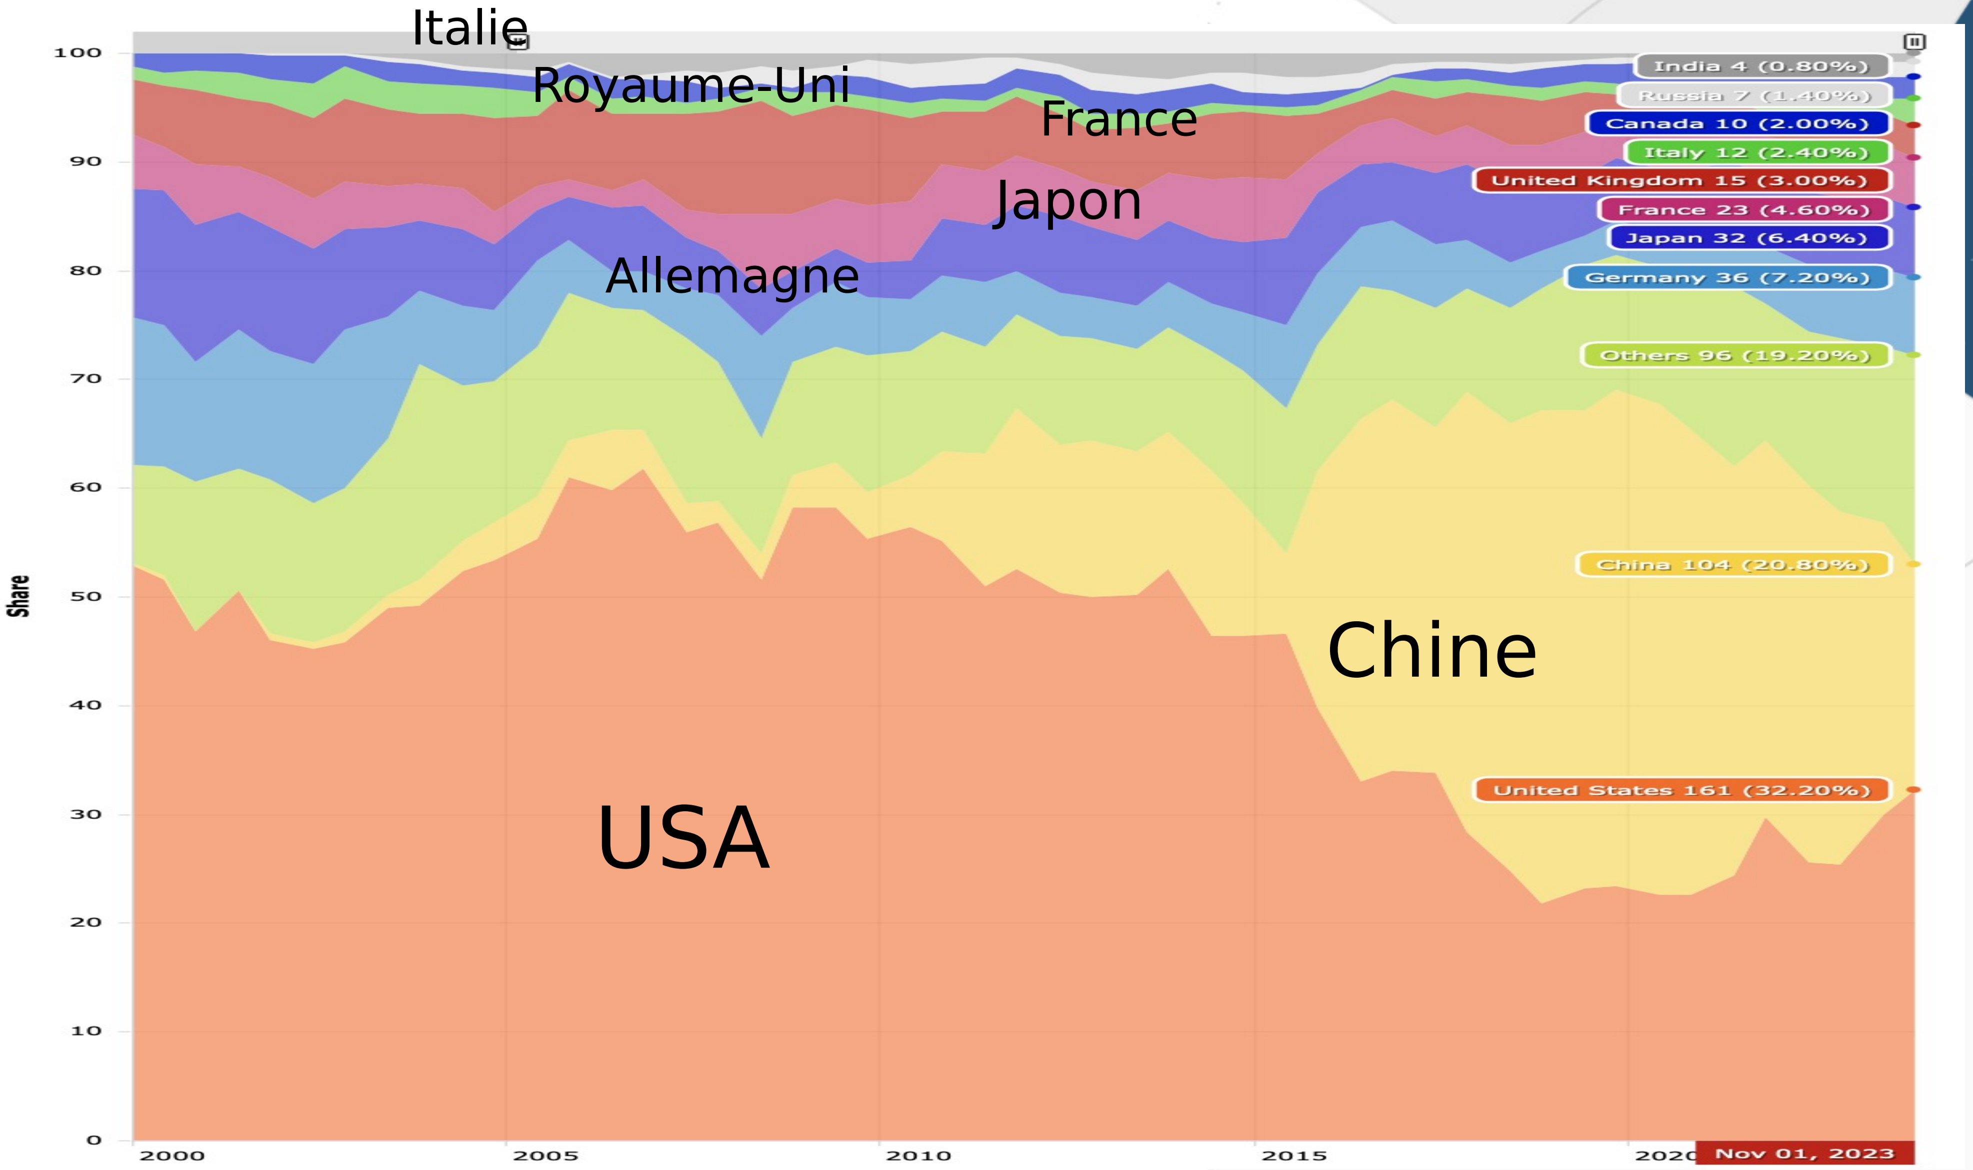
Italie
Royaume-Uni
France
Japon
Allemagne
Chine
USA
Share
100
90
80
70
60
50
40
30
20
10
0
2000
2005
2010
2015
2020
Nov 01, 2023
India 4 (0.80%)
Russia 7 (1.40%)
Canada 10 (2.00%)
Italy 12 (2.40%)
United Kingdom 15 (3.00%)
France 23 (4.60%)
Japan 32 (6.40%)
Germany 36 (7.20%)
Others 96 (19.20%)
China 104 (20.80%)
United States 161 (32.20%)

TOP 500 – 2023

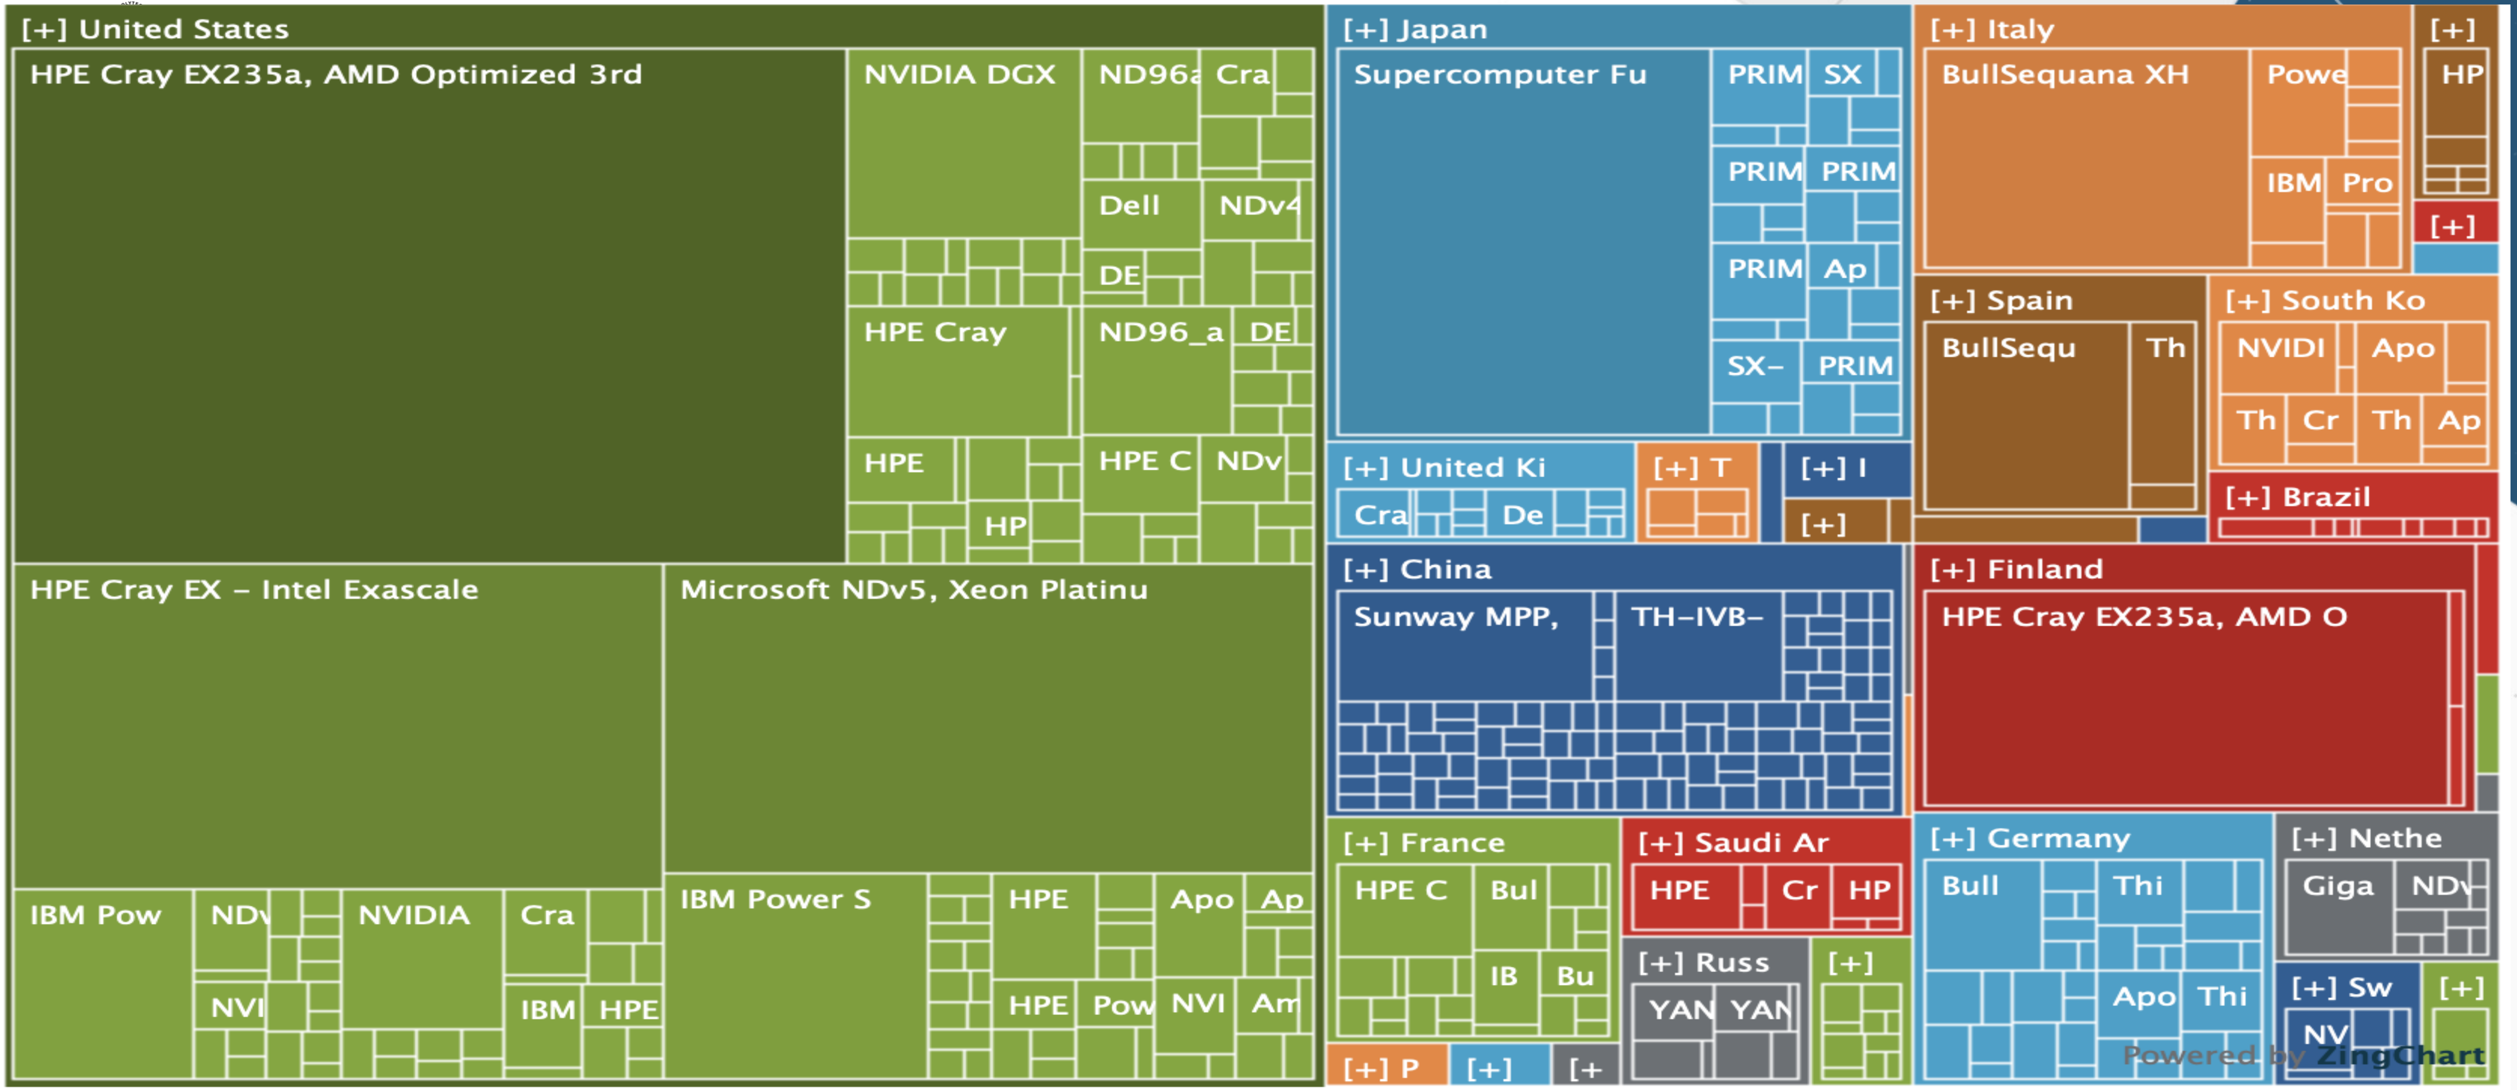

[+] United States
HPE Cray EX235a, AMD Optimized 3rd
NVIDIA DGX
HPE Cray EX – Intel Exascale
Microsoft NDv5, Xeon Platinu
IBM Power S
[+] Japan
Supercomputer Fu
[+] Italy
BullSequana XH
[+] Spain
[+] South Ko
[+] Brazil
[+] United Ki
[+] China
Sunway MPP,
TH-IVB-
[+] Finland
HPE Cray EX235a, AMD O
[+] France
[+] Saudi Ar
[+] Germany
[+] Nethe
[+] Russ
[+] Sw
Powered by ZingChart

Frontier (Oak Ridge, USA) #1

~40 000 GPU + 8 730 112 AMD cores → 1,68 exaflops (1,68 10^{18} flops)
22,7 MW (= 2 TGV Duplex)

https://www.olcf.ornl.gov/frontier/

Aurora (Argonne DOE) #2

9,264,128 CPU → 1,06 exaflops - 25 MW (> 2 TGV Duplex)

CEA-HE(CEA, France) #17 in France#1

389,232 Grace-Hopper CPU-GPU → 112.56 petaflops - 1,207GW

Weather Modeling
An example of numerical simulation

Objective : compute humidity, pressure, temperature and wind speed relatively to x,y,z and t.
Solving a dynamic system formally impossible

Deterministic numerical simulation :

- Space discretization
  - Illustration : 1 point represents $2km^3$ by 20 km of atmosphère => 5.10^9 points
- Model initialization
  - Satellite and other data sensor
  - Interpolation of missing values required

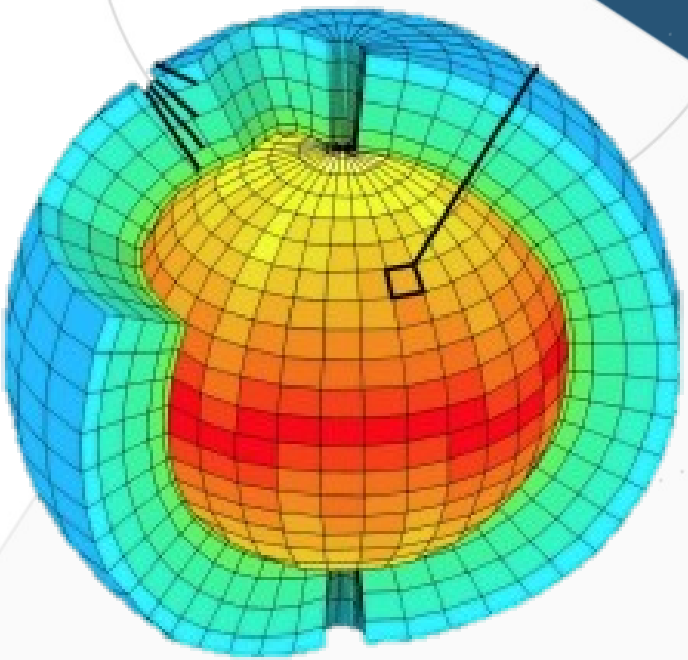

Weather Modeling
An example of numerical simulation

- Model running
  - Time discretization
    - Ex: compute by step of 60sec
- Solving a system of equations for each mesh
  - calculate the next state of a mesh as a function of its neighborhood, solar gain and the earth's rotation
  - Illustrative cost : 100 FLOP/point

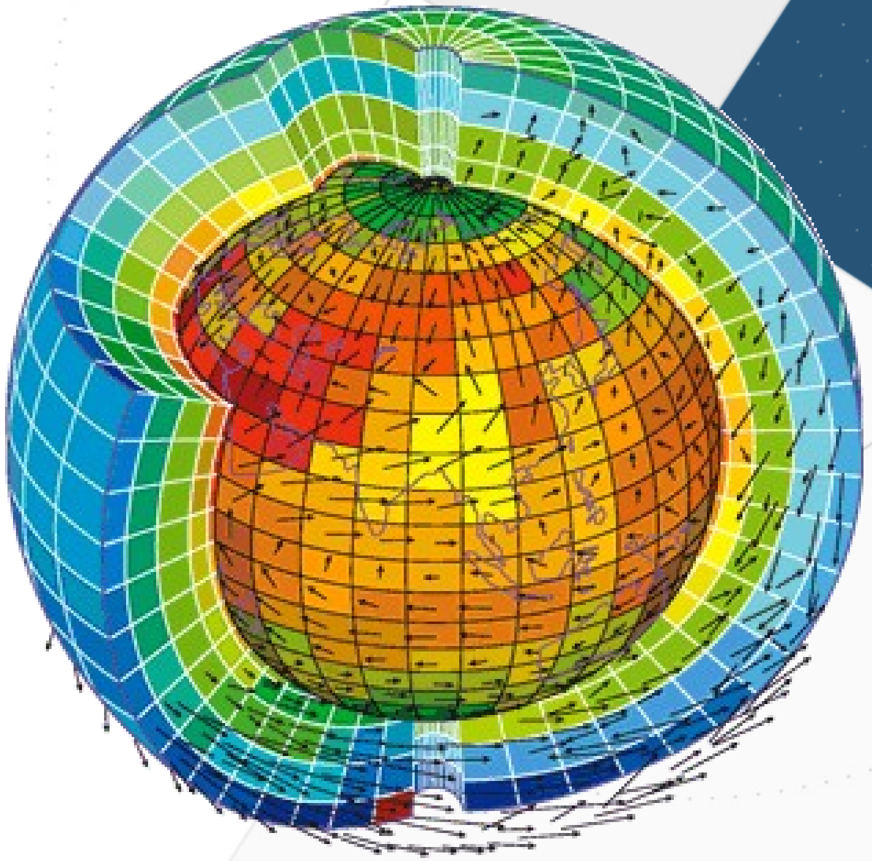

Weather Modeling
An example of numerical simulation

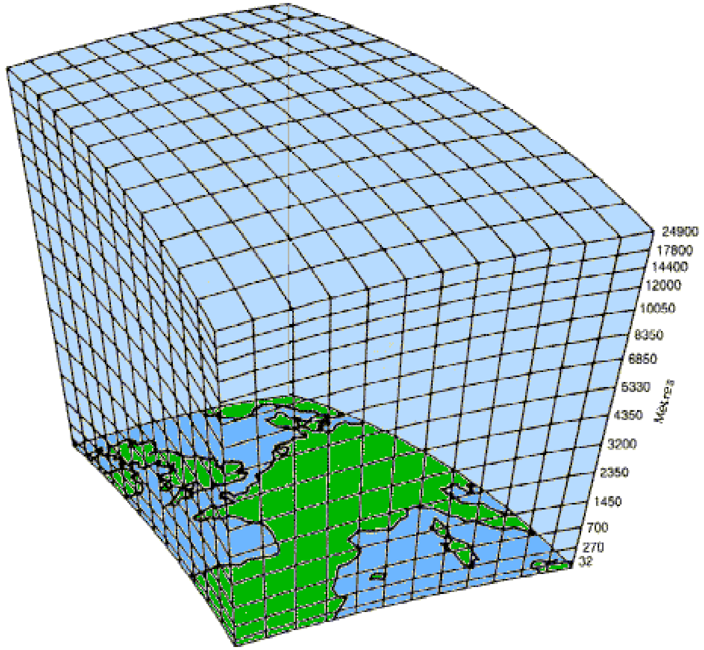

- 1 step requires ***5.10^{11} flop***

Computing power needed :

- Real time – 1 step in 60s : *5.10^{11} / 60 =* ***8,33 Giga flop/s***
- Forecast - 7 days in 1h *: 5.10^{11} * 7j * 24h * 60m / 3600 =* **1400 Gflop/s**
- Climatology - 50 years in 1 day : 1 060 000 Gflop/s **= 1,06 Pflop/s**

Hardware of a boosted PC

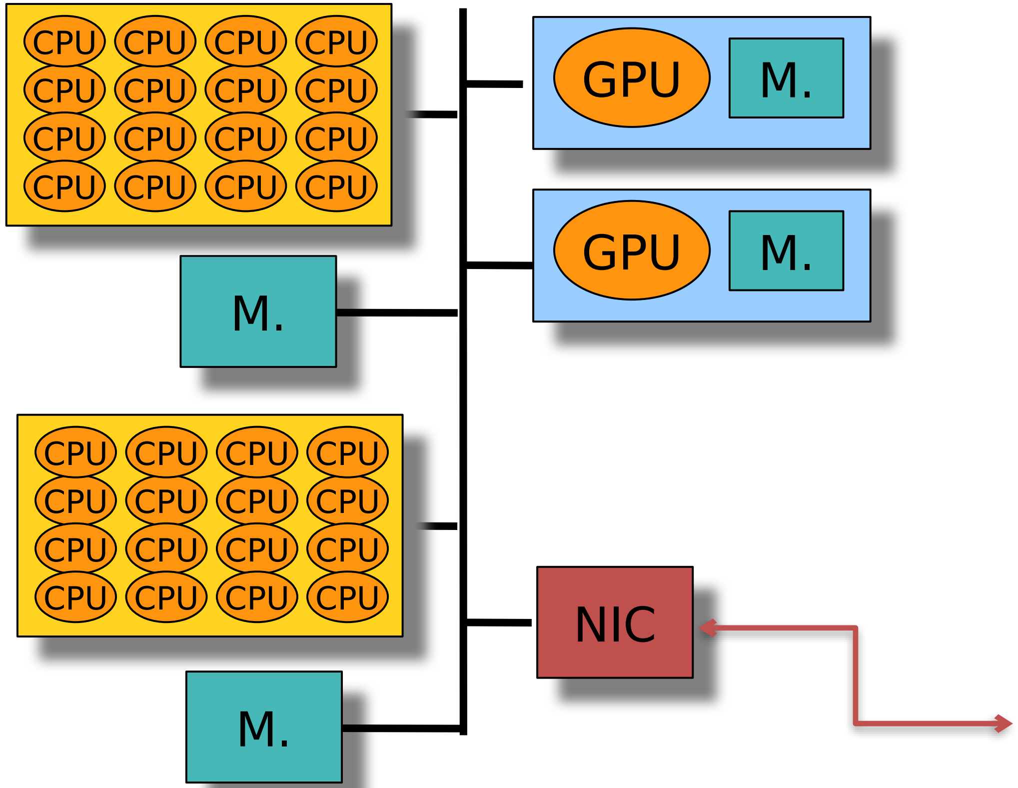

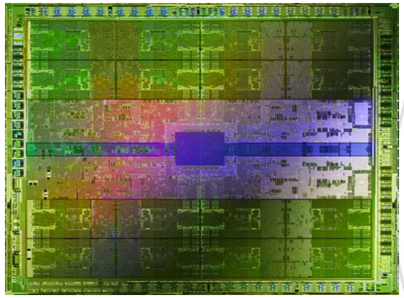

Nvidia A100
128 SM
9,2 Tflops
(10k€)

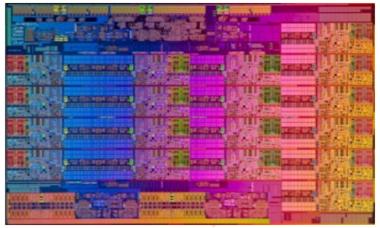

INTEL® XEON®
PLATINUM 9282
56 cores
2,4 TFlops

High speed
interconnect
200 Gb/s

Performances of a boosted PC

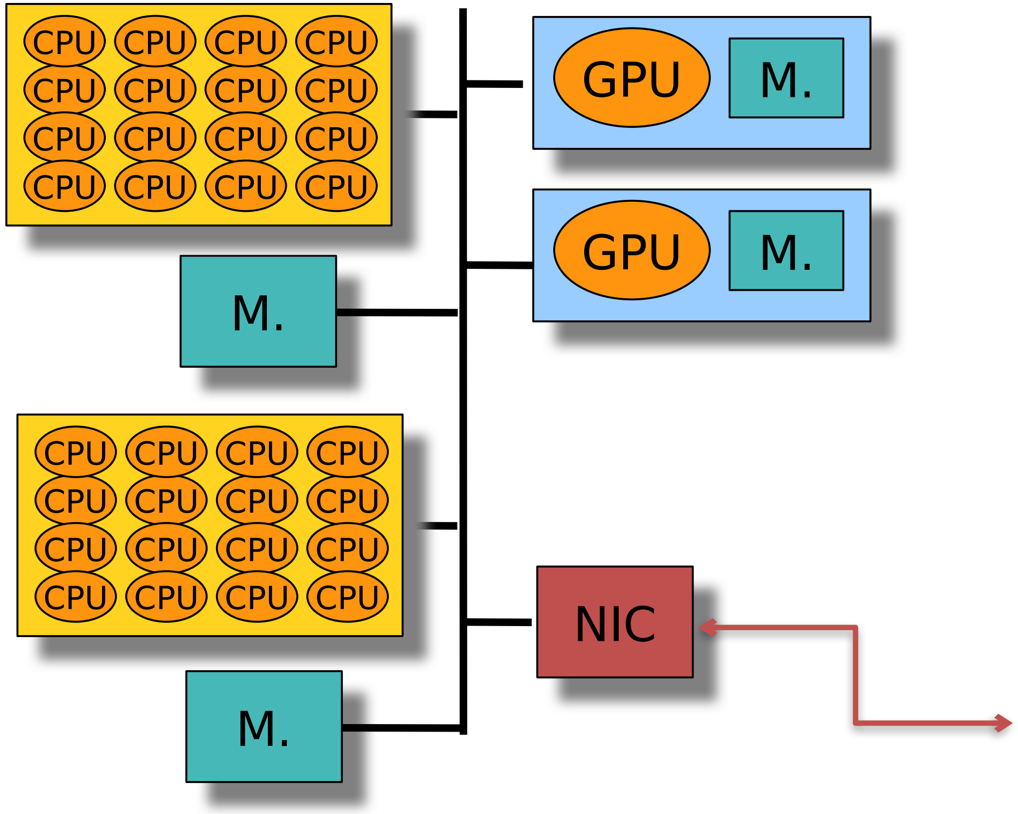

2 GPU + 2 CPUs = 17,6 Tera flop/s

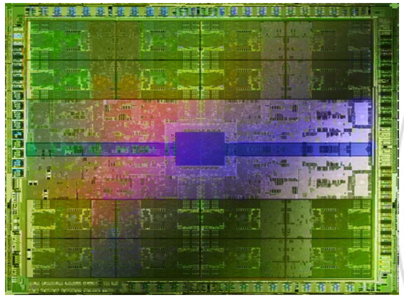

Nvidia A100
128 SM
9,2 Tflops
(10k€)

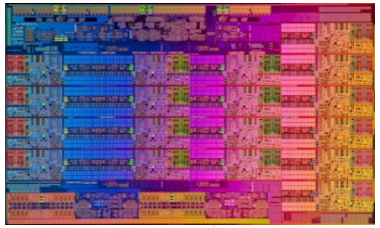

INTEL® XEON® PLATINUM 9282
56 cores
2,4 TFlops

High speed interconnect
200 Gb/s

Forecast compute time

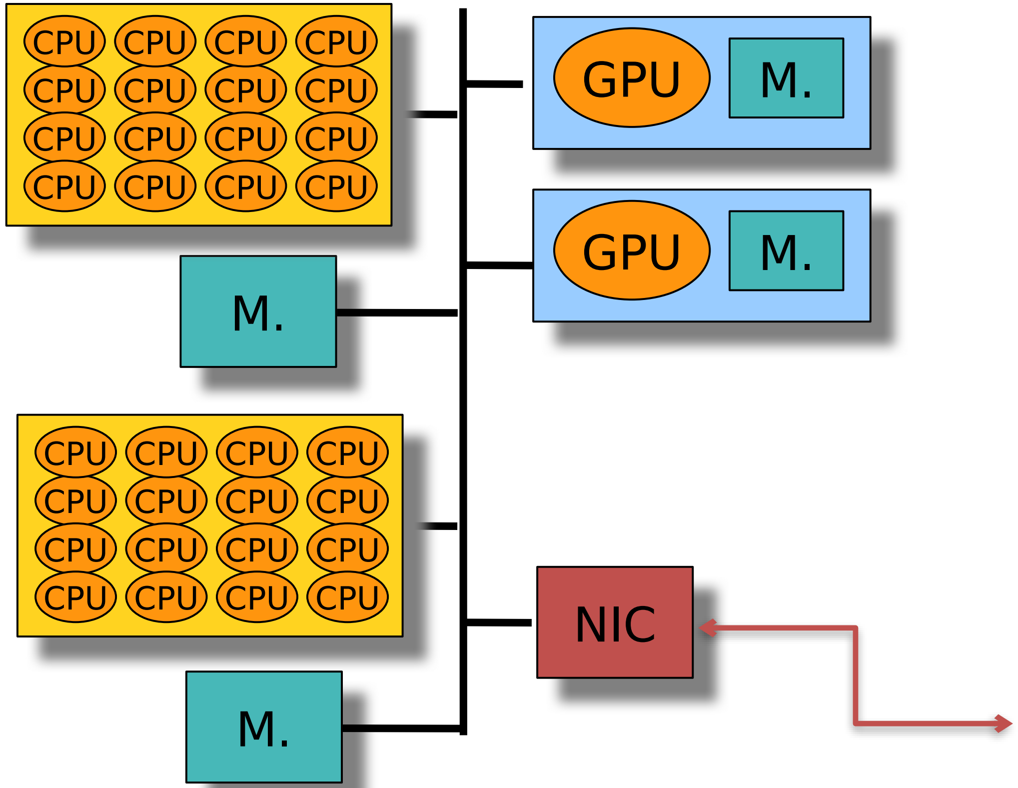

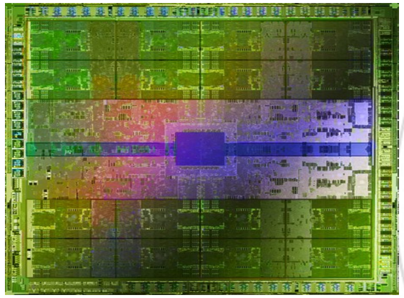

Nvidia A100
128 SM
9,2 Tflops
(10k€)

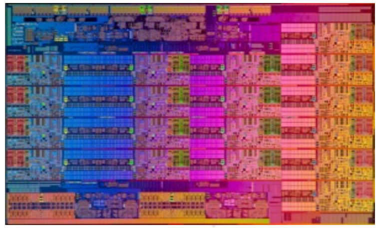

INTEL® XEON®
PLATINUM 9282
56 cores
2,4 TFlops

High speed
interconnect
200 Gb/s

2 GPU + 2 CPU = 17,6 Tera flop/s
7 days weather forecast
-> *5.10*^{11} * 7j * 24h * 60m = 5 040 Teraflop
computes in 4 min 45

Forecast compute time

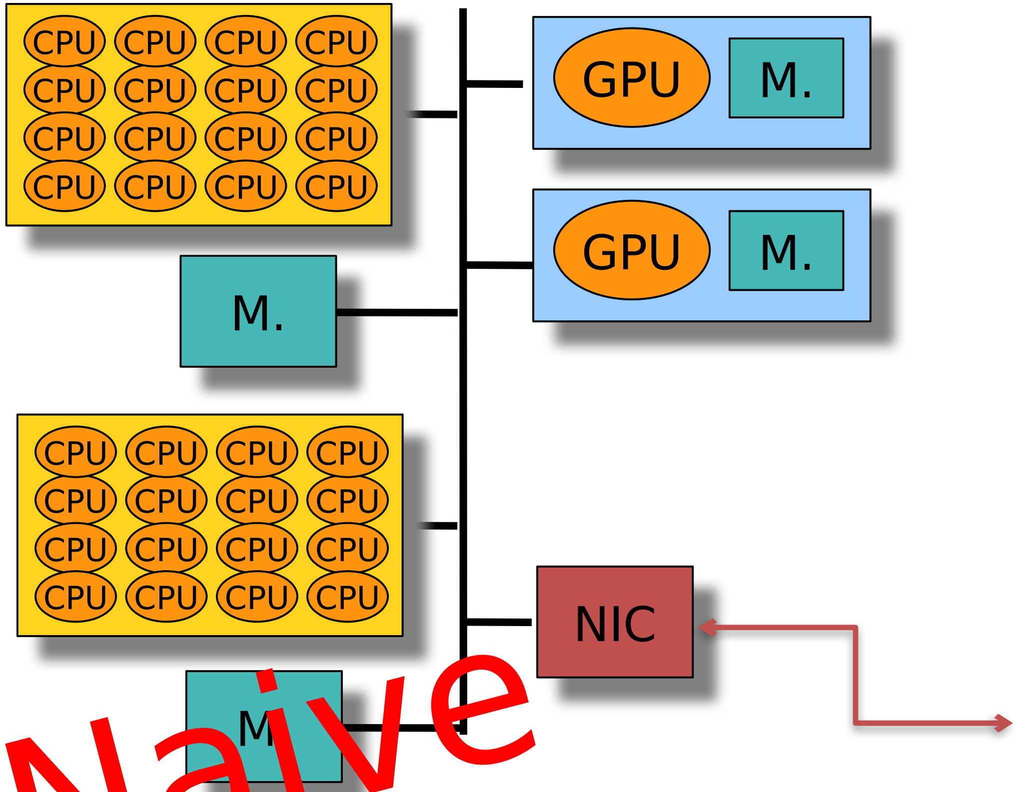

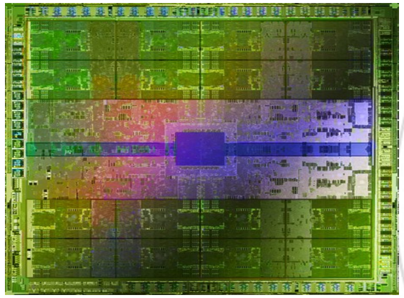

Nvidia A100
128 SM
9,2 Tflops
(10k€)

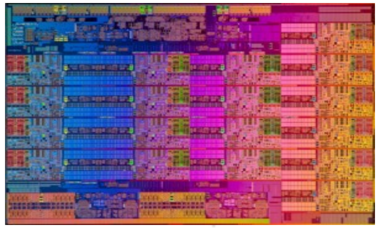

INTEL® XEON®
PLATINUM 9282
56 cores
2,4 TFlops

High speed
interconnect
200 Gb/s

2 GPU + 2 CPU = 17,6 Tera flop/s
7 days weather forecast
-> $5\ 10^{11}$ * 7j * 24h * 60m = 5 040 Teraflop
computes in 4 min 45

Utopian peak performance

- Really hard to reach peak performance, even in easy cases
  - Matrix multiplication with Intel MKL et CudaBLAS librairies
  - E5-2680 v3 - 2.5 MHZ
  - Kepler 40 (1430 GF/s -> 2013)

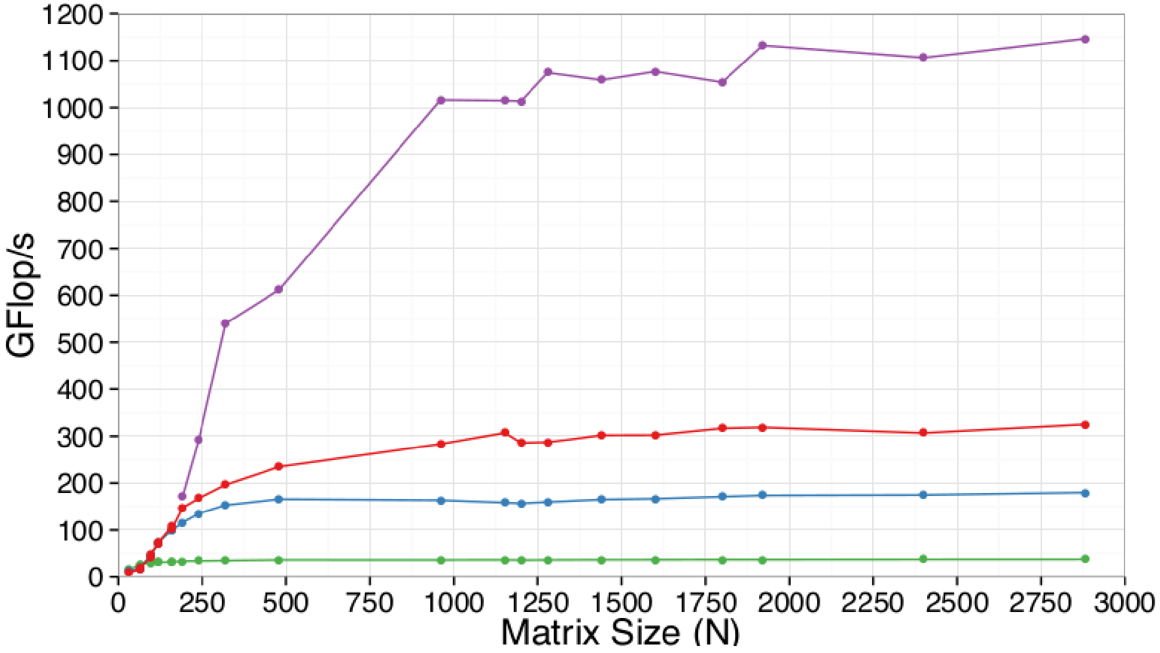

Efficiency
95% 1 core
93% 5 cores
87% 10 cores
81% 1 GPU 2013
[9? % GPU202x ?]

Utopian peak performance

- Complex matrix computation

  Factorisation de Cholesky

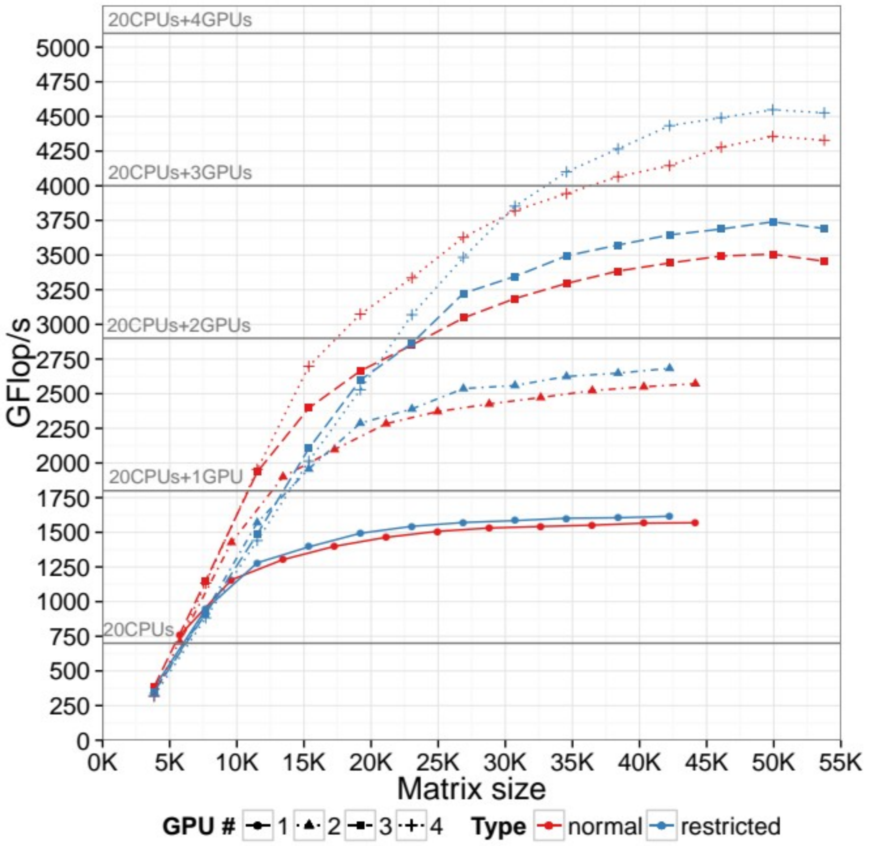

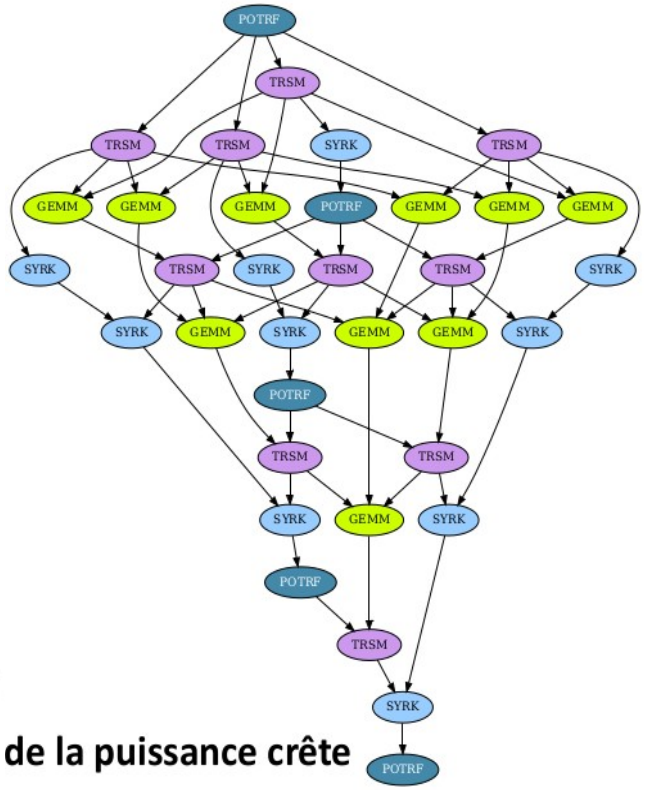

~90% of the theoritical perf
80% of the peak performance

Intuitively, acceleration is bounded
by the execution time of the critical path

Amdahl Law

Speedup = Sequential computation time / Parallel computation time

- *S,* the sequential part of the program
- *1 – S,* the parallelizable part of the program
- P, number of processors

Potentiel parallel version makespan ≥ (S+ (1-S)/P) * sequential makespan

Speedup expected on a P-core machine bounded by

***Speedup* ≤ *1*/ (S+ (1-S)/P) < 1/S**

i.e
« if 1% of the application is sequential, no hope to reach an 100 time faster computation»

Amdahl Law

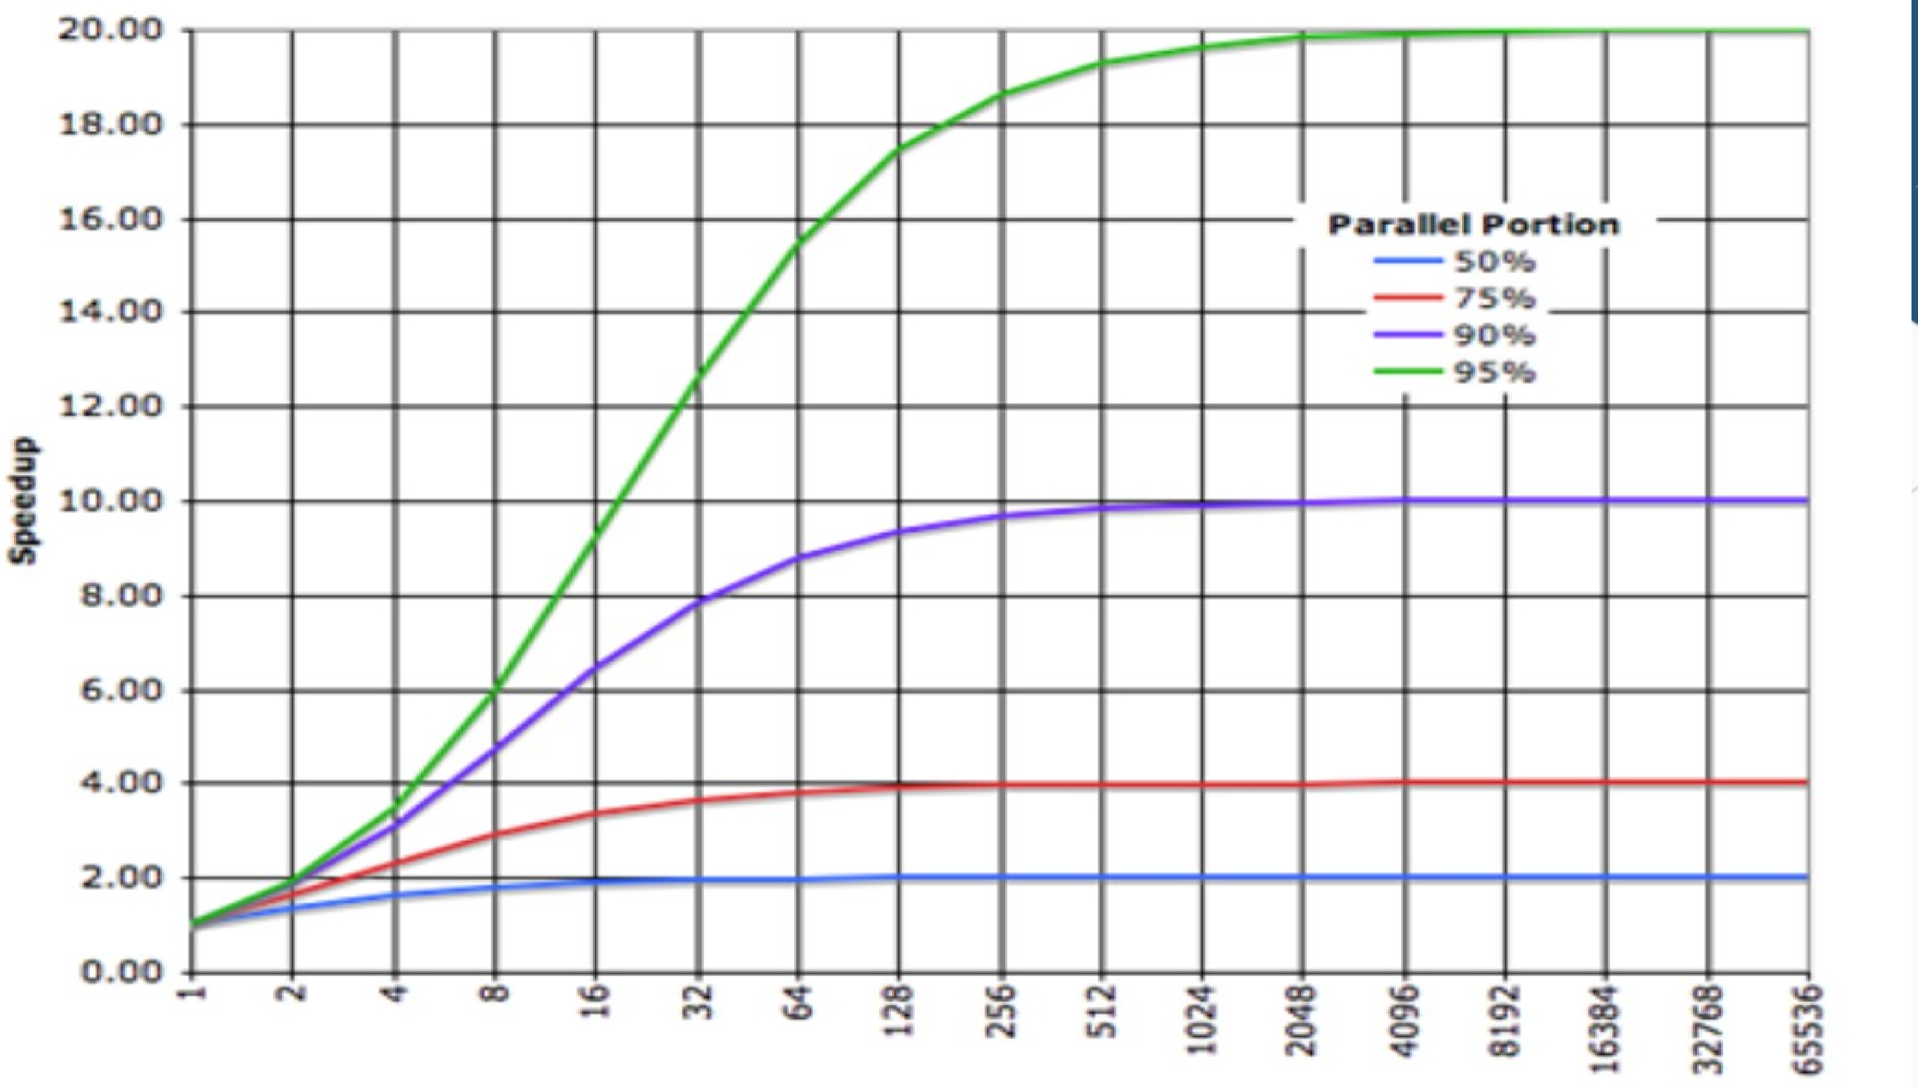

20.00
18.00
16.00
14.00
12.00
10.00
8.00
6.00
4.00
2.00
0.00
Speedup
Parallel Portion
50%
75%
90%
95%
1
2
4
8
16
32
64
128
256
512
1024
2048
4096
8192
16384
32768
65536