Numerical simulation: tools and perspectives

Geoffroy Lesur (IPAG Grenoble, Chair of GENCI Astrophysics & Geophysics allocation committee)
Paris, Sep. 6 2023

The computing pyramid

Euro HPC= European High Performance Computing Joint Undertaking
eurohpc-ju.europa.eu/

GENCI=Grand équipement national de calcul intensif
https://www.genci.fr

Current national facilities

National Clusters

- Adastra (Hosted by CINES in Montpellier,)
  - 1352 AMD MI250 GPUs (128 GB memory per GPU) *(production since January 2023)*
  - 102 912 AMD EPYC Genoa cores (4 GB per core) *(production since May 2023)*
  - 74 PFlop/s (71 on GPUs)

- Jean Zay (Hosted by IDRIS in Saclay)
  - 86 344 Intel Cascade Lake CPU cores (4.8 GB memory per cpu core)
  - 2696 Nvidia (Volta & Ampere) GPUs (16-32 GB memory per gpu core)
  - 28 PFlop/s

- Joliot Curie (Hosted by TGCC in Bruyères-le-Châtel)
  - 79 488 Intel Skylake cpu cores (4 GB memory per cpu core, 6.9 PFlop/s)
  - 293 376 AMD EPYC Rome CPU cores (2 GB memory per cpu core, 11.75 PFlop/s)

Adastra @ CINES

Jean Zay @ IDRIS


Joliot Curie @ TGCC

Evolution of GENCI ressources

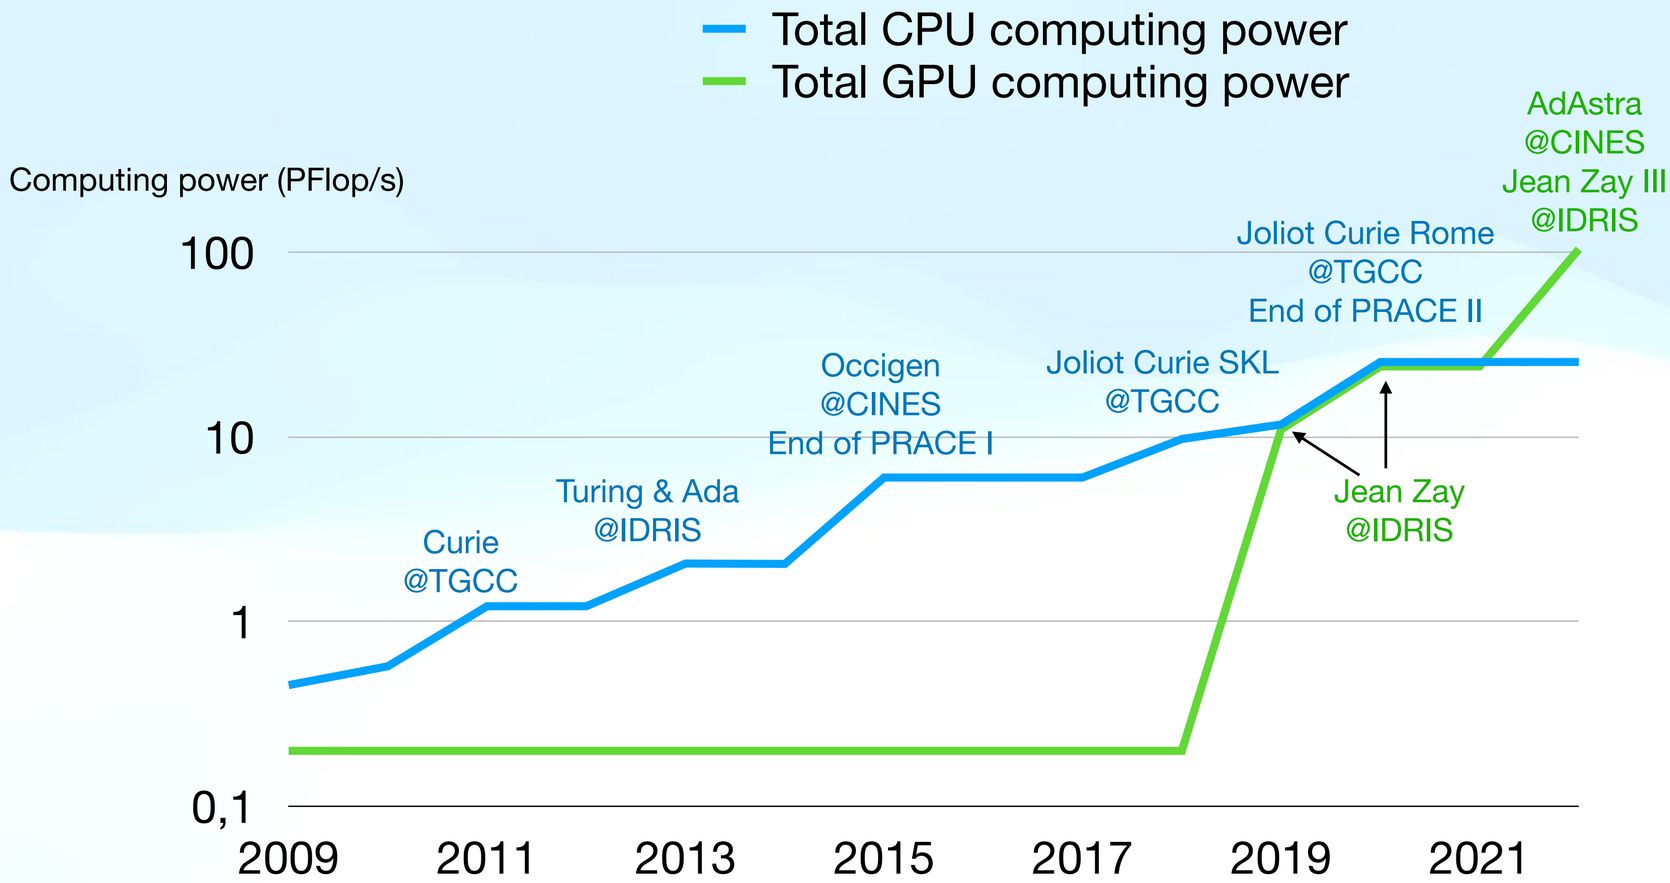

Large majority of computing power is now on GPU accelerated clusters
CPU computing power is expected to stall (at best)

Astrophysics & geophysics projects by themes

CPU hours requested (total: >200Mh/year)

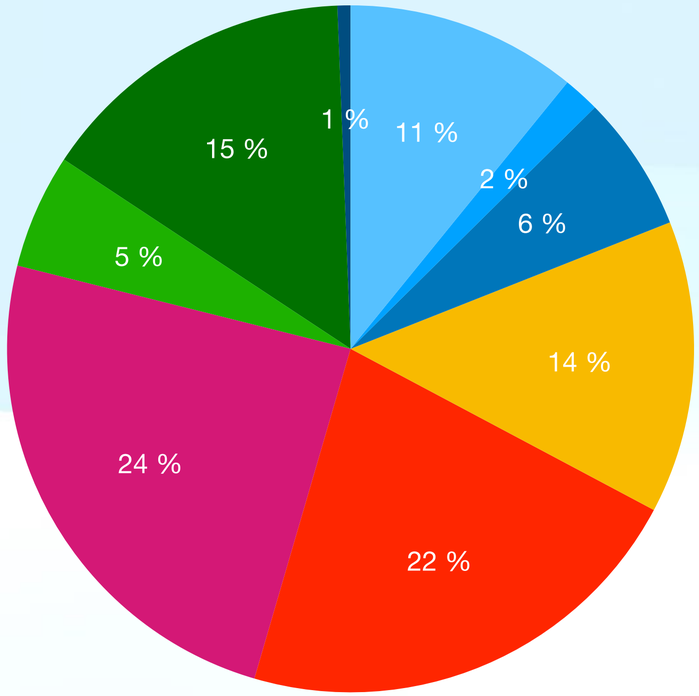

#projects

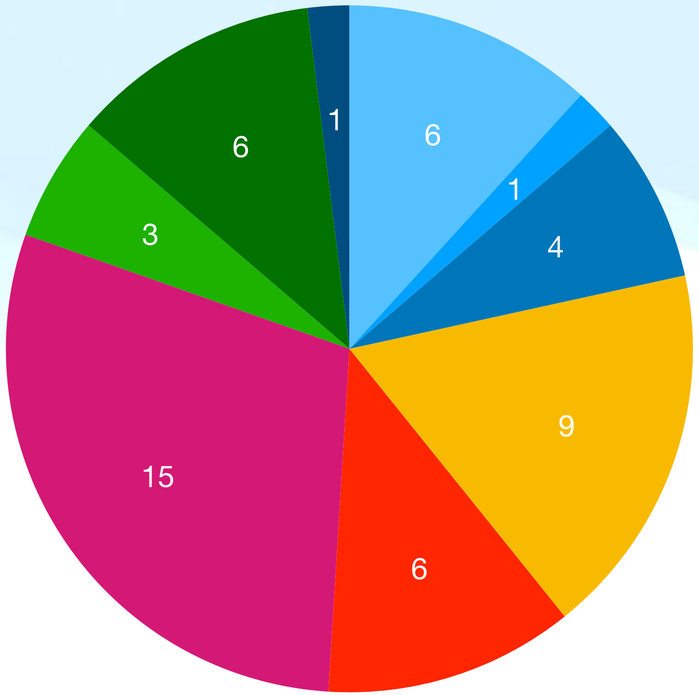

Sismologie
Galactique
Astrochimie
Etude des sols
Plasmas
Planetologie
Stellaire
Cosmologie
Etoiles jeunes

~25% are connected to PNHE (mostly Plasmas)
Some disciplines request much more hours than others

Future facilities

The next step: the « French » exascale machine

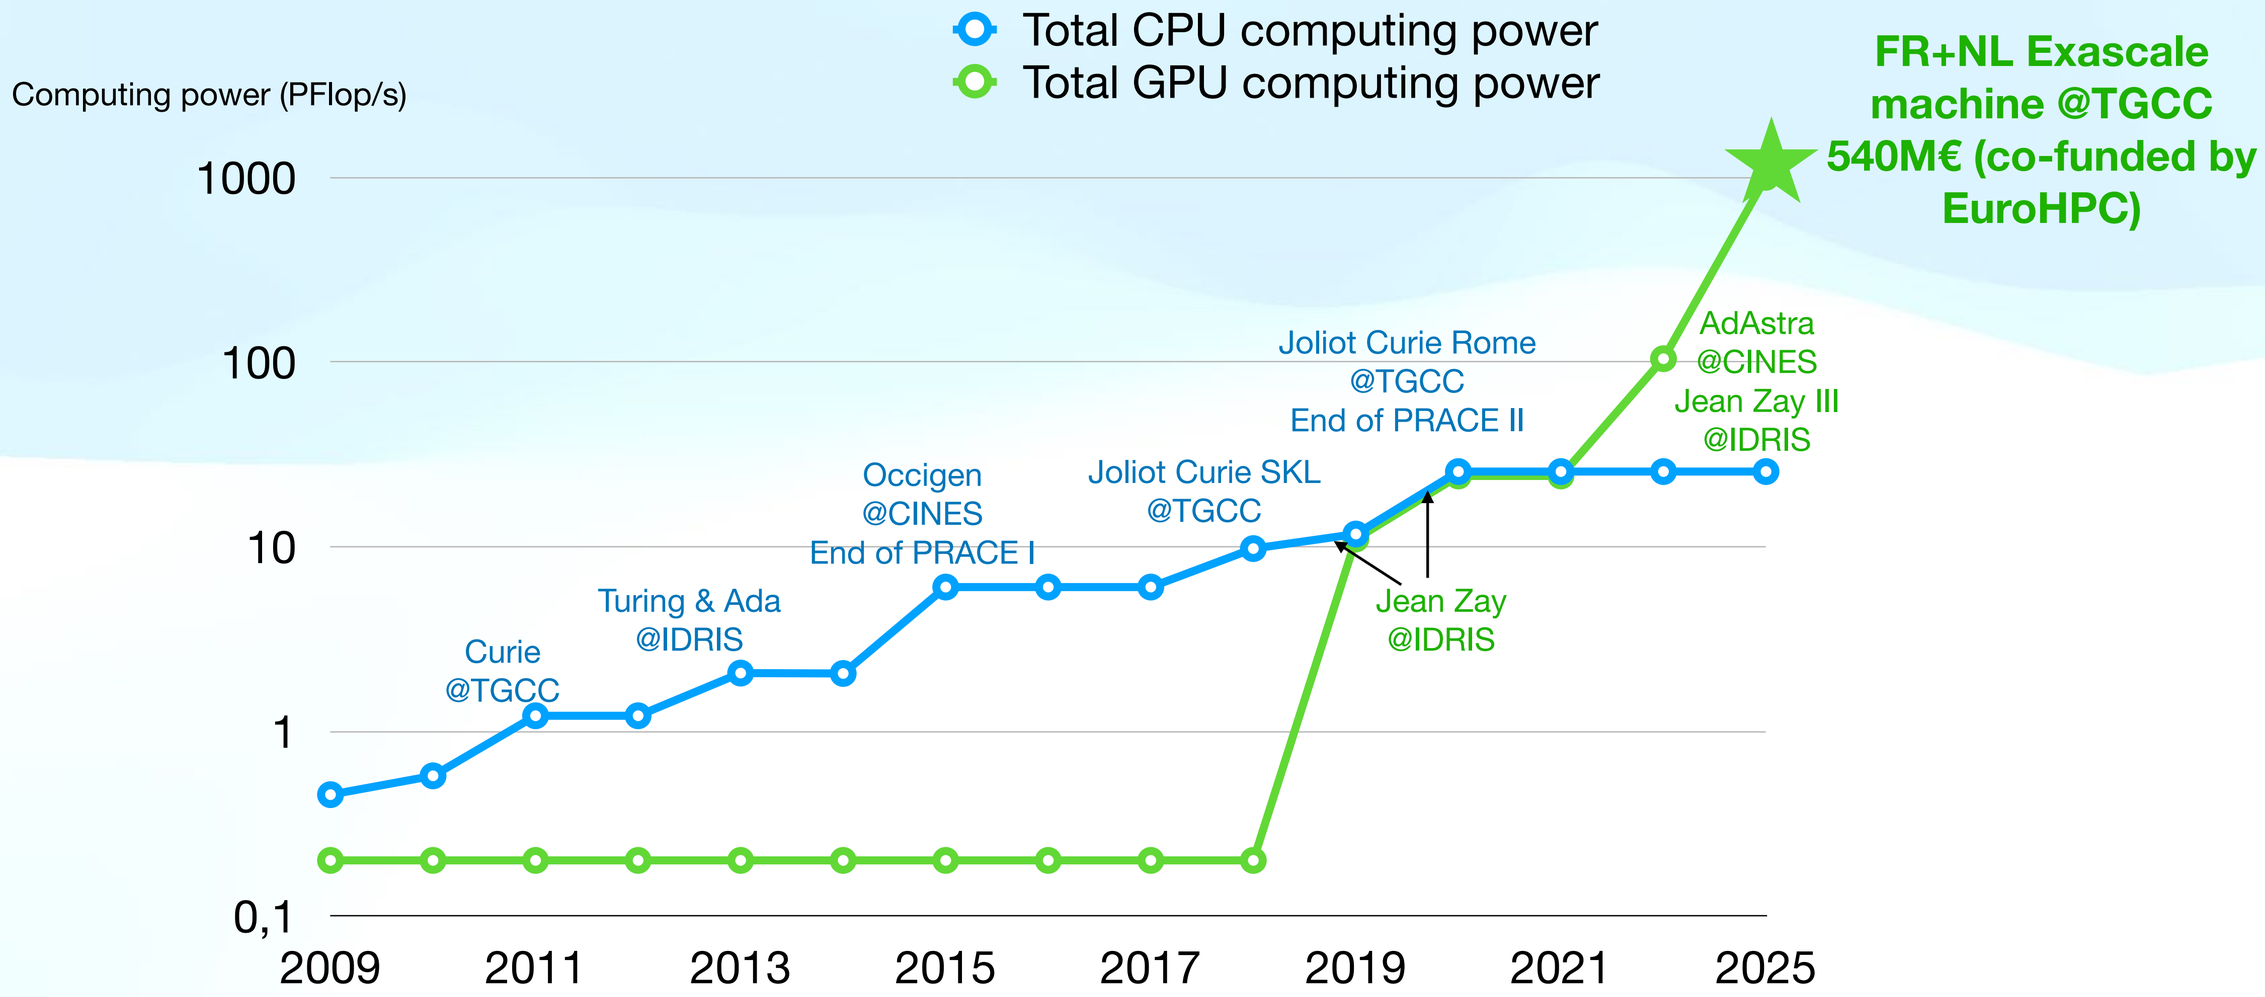

Exascale machine: what for?

- Multi-scale physics (e.g. coupling small-scale kinetic models to large-scale fluid flows)
- Coupling with AI models (e.g. sub-grid models)
- Secular evolution (do more than a few dynamical times of the system)
- Multi-physics, multi-code problems (on-the-fly multi-group radiative transfer, multi-fluid evolution, particle-fluid coupling)

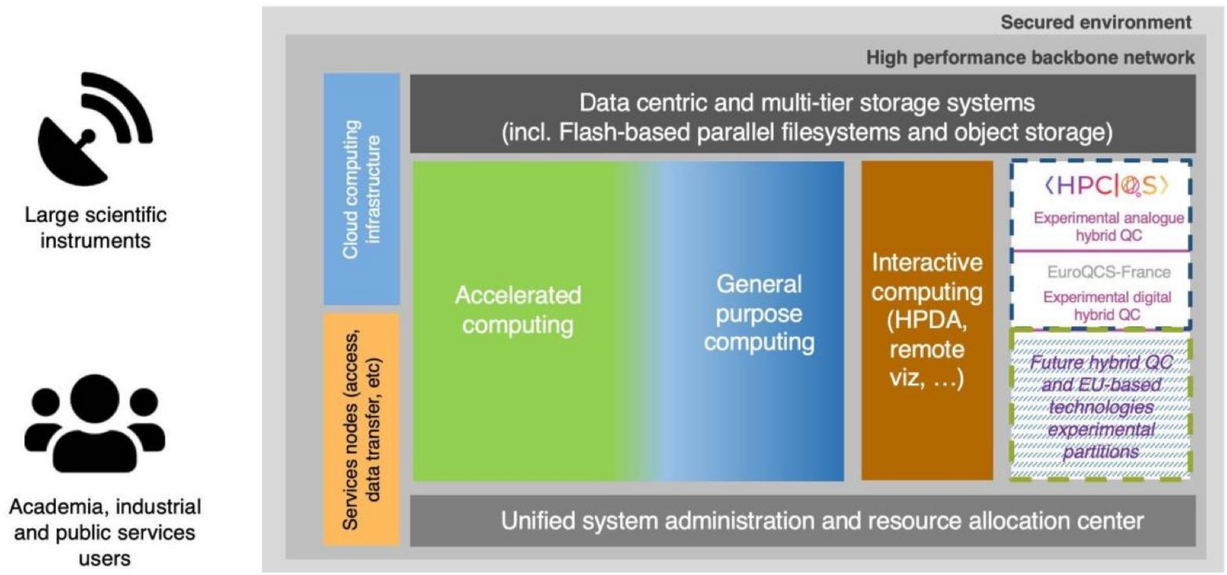

Image : architecture générale du supercalculateur Exascale proposé

Towards Exascale Machines

An electrical power issue

=?

- The current average power consumption of a x86 processor (your laptop!) is ~2Gflop/s/W
- an Exaflop machine would need about 500 MW of power=1 french nuclear reactor
- The « socially acceptable » power consumption is 30 MW → reduce the energy footprint by a factor 6

Accelerated clusters (GPU-like) are mandatory

The tools

Exascale-ready codes

- Not possible to directly use your favorite code on an accelerated machine
- One needs to port codes, or to write them from scratch (in both cases, years of development)
- To make things worse: each manufacturer tend to promote its own technology (e.g. CUDA@NVidia, HIP@AMD, SYSCL@INTEL…).
- Current trend is to write performance-portable codes relying on meta-programming library (heavy C++, e.g. Kokkos)
- A few open-source performance-portable codes for PNHE applications:
  - Dyablo (AMR framework, Maison de la simulation, soon public)
  - Idefix (Multi-fluid physics, IPAG Grenoble, available on GitHub)
  - Parthenon (AMR framework, Los Alamos National lab., available on GitHub)

Parallel performances

The Idefix code running on the 3 national HPC centers

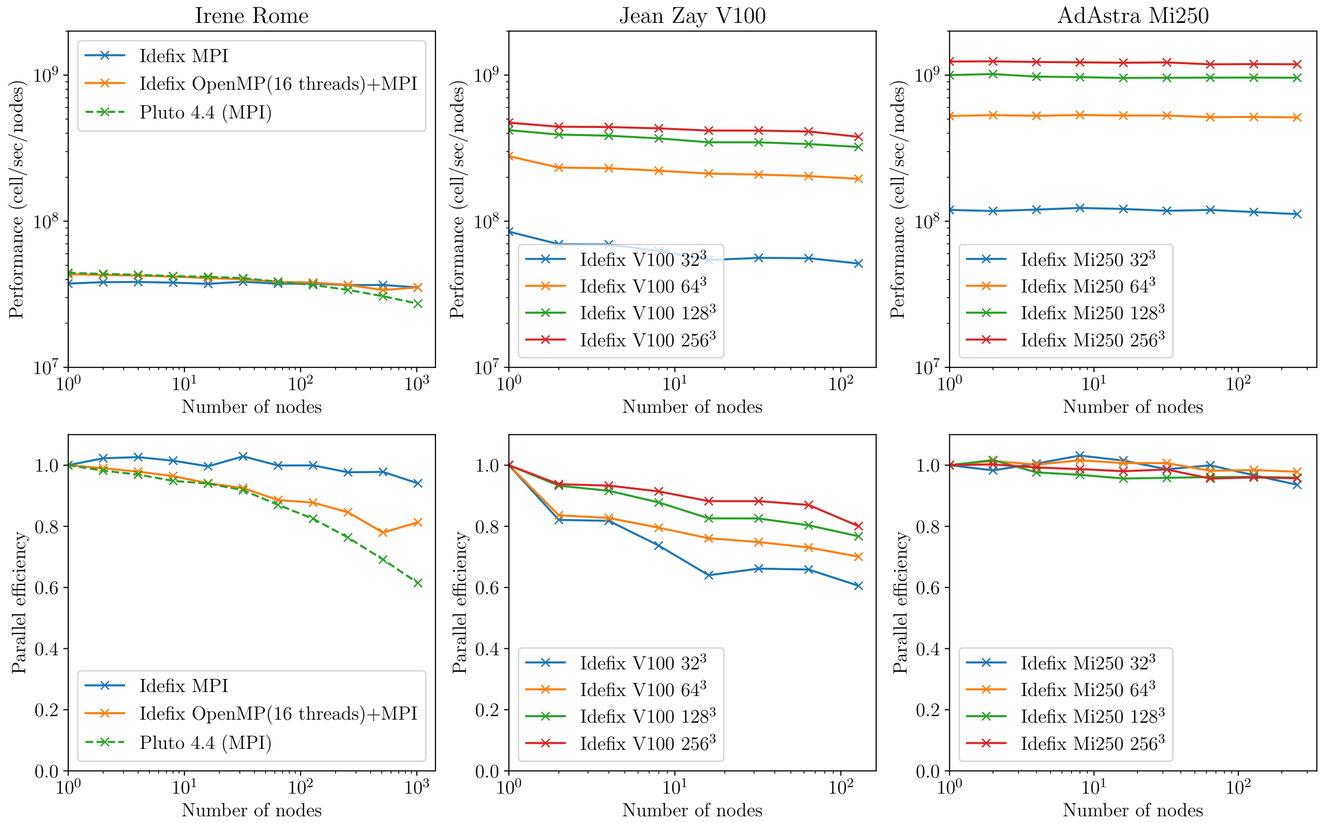

[Lesur+2023]

- Similar performance to Pluto 4.3 on CPUs → performance portability is a success!
- 1Mi250 AdAstra node~30 Rome nodes
- Fairly good scaling (75% efficiency on 131072 CPUs, 80% efficiency on 512 V100s, 96% on 1024 Mi250)

Energy efficiency

Idefix code

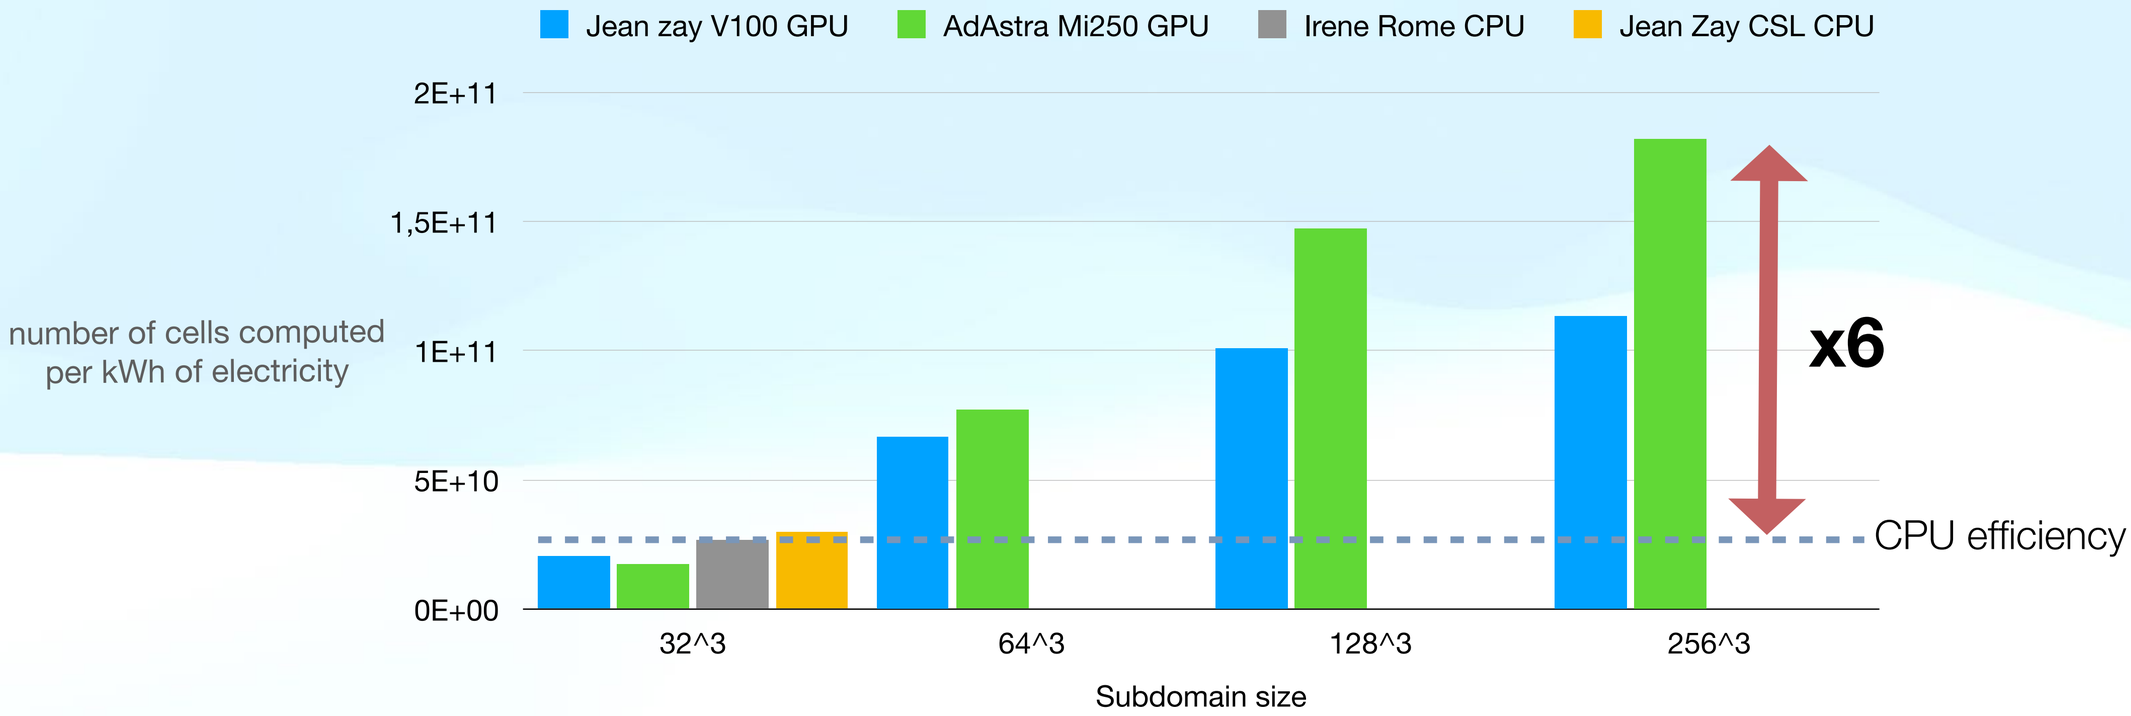

Running the same problem on a GPU **is up to 6 times** more energy efficient on real applications

but

Strong dependence on the domain size!

Conclusions

- Pre-exascale machines are available, Exascale machines are coming (~3 years)
- Opportunities to tackle long-standing problems involving large scale separation and/or multi-physics
- Support for accelerated architecture will become mandatory (energy constraints)
- New performance-portable codes start to be available on the market (use them!)
- **Strong need** for a long-term support to maintain & develop these new codes, train students & create communities around them

GENCI’s comités thématique (CT)

CT1 : Environnement
Président : M. Jean-Louis Dufresne

CT2A : Écoulements non réactifs
Président : M. Guillaume Balarac

CT2B : Écoulements réactifs ou/et multiphasiques
Président : Mme Pascale Domingo

CT3 : Biologie et santé
Président : Mme Yolanda Prezado

CT4 : Astrophysique et géophysique
Président : M. Geoffroy Lesur

Keywords: Cosmologie. Formation des galaxies, des étoiles et des systèmes planétaires. Dynamique des systèmes gravitationnels. Modélisation d'objets astrophysiques. Plasmas géophysiques et planétaires. Géophysique interne. Hydrologie des sols. Géomatériaux.

CT5 : Physique théorique et physique des plasmas
Président : Mme Virginie Grandgirard

Keywords: Electromagnétisme. Physique sur réseau dont QCD. Chaos quantique. Propriétés électroniques des solides. Physique nucléaire. Interactions ondes électromagnétiques avec la matière. Plasmas chauds. Sciences de la fusion magnétique ou inertielle.

CT6 : Informatique, algorithmique et mathématiques
Président : Mme Hélène Barucq

CT7 : Modélisation moléculaire appliquée à la biologie
Président : M. Patrick Fuchs

CT8 : Chimie quantique et modélisation moléculaire
Président : M. David Loffreda

Keywords: Propriétés électroniques des molécules. Structures. Réactivité. Calculs ab initio. Calculs semi-empiriques. Dynamique quantique (Car-Parinello). Calculs Monte-Carlo quantique (Méthodes QMC). Etat liquide. Solvation. Diffusion moléculaire. Collisions (molécules-ions, électrons). Dynamique quantique. Evolution d'un paquet d'ondes.

CT9 : Physique, chimie et propriétés des matériaux
Président : M. Thierry Deutsch

CT10 : Intelligence artificielle et applications transversales du calcul
Président : M. Michaël Krajecki

Usage of GPU-accelerated cluster in Astro & Geophysics

Share of Astro & Geophysics in the total time allocated by GENCI (national machines)

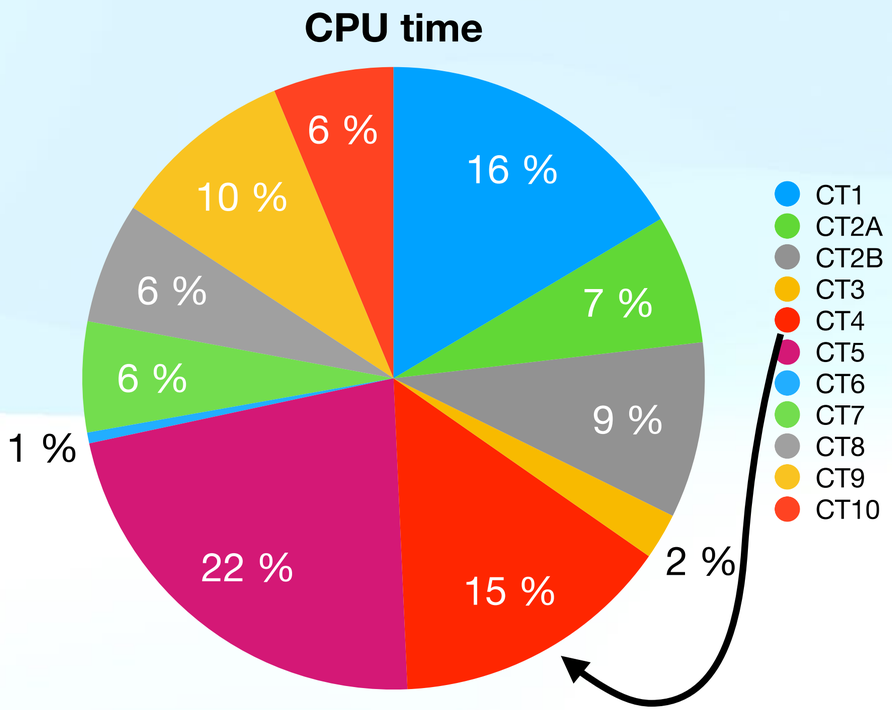

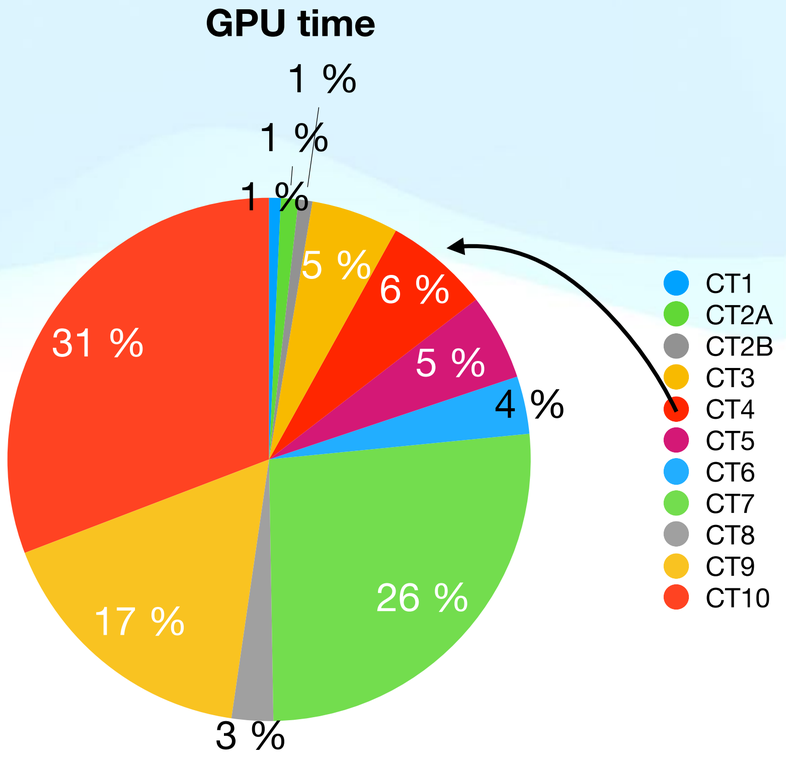

While Astrophysics represents 15% of the available CPU time, it reaches *only* 6% of the available GPU time

CT4 « standard » project workflow

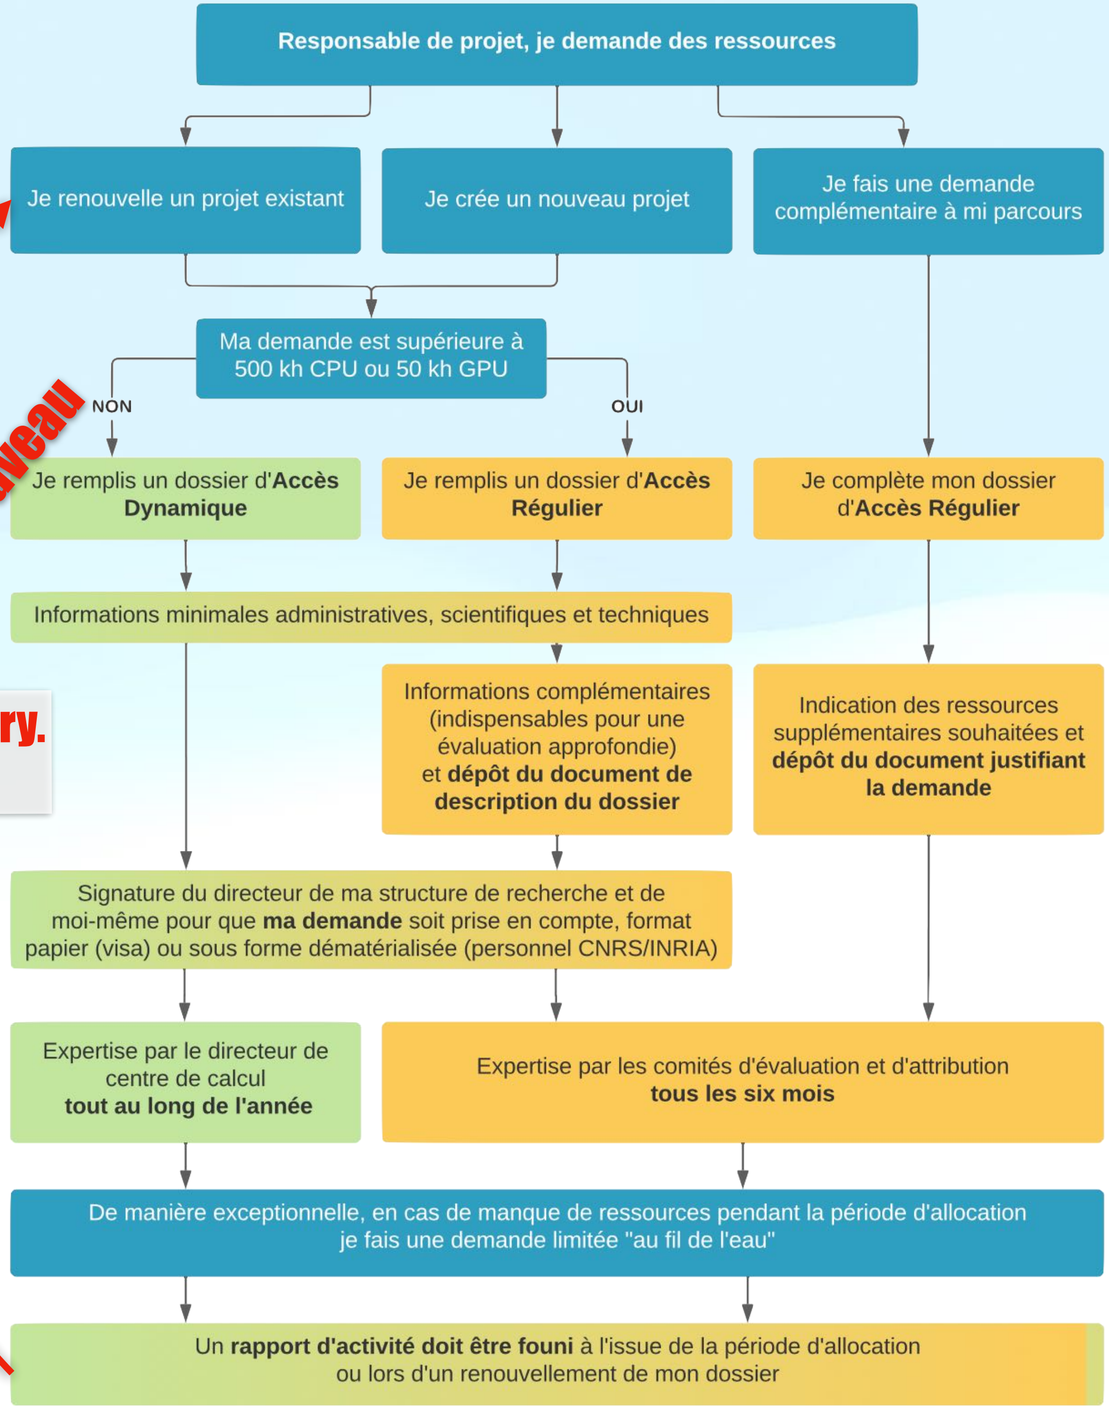

Project report is mandatory.
No report=no renewal

Pressure on ressources

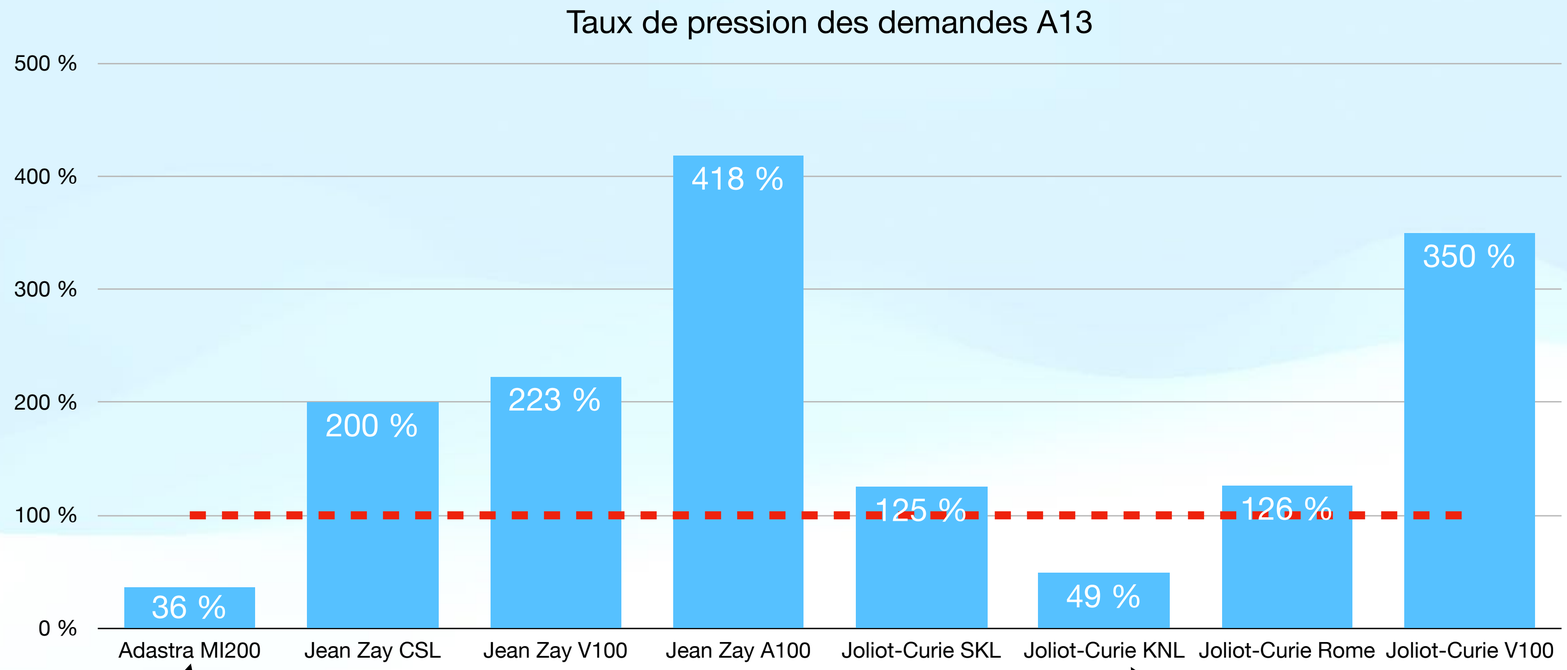

New machine
(in prod. early 2023)

Early shutdown (dec. 2022)

Pressure can be high on some specific clusters. If possible, tell us whether you can move your code/tools to another cluster to optimise allocations.

Tips: how to write a good GENCI proposal

- **Rule I**: people who read your proposal are *also* researchers, and are *also* busy
  - Keep it short: 4 pages to explain the scientific context should be sufficient
  - Do not copy-paste your last ERC/ANR proposal
  - Make it readable (sentences have a verb in English & French…)
  - Do not list all of your projects over the last 10 years ending with « we don’t request anything this year on this part »…

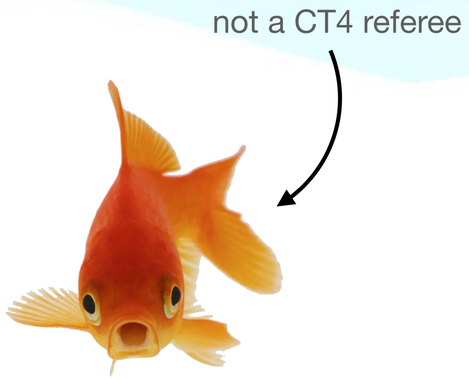

- **Rule II**: people who read your proposal are not goldfishes (i.e. they have memory)
  - Do not copy paste your proposal from last year without changing a word
  - Do not try to submit a renewed proposal as « new » without uploading a report, hoping we will not notice it (or expect a phone call from me)

Tips: how to write a good GENCI proposal

- **Rule III**: be honest
  - You did not manage to use all of your previous allocation: explain why.
  - Note that it’s possible to give back a part of your allocation on e-dari (section « restituer des heures »).
  - Do not try to cheat us on the hours you have used: we have access to project accounting on every GENCI machines…
- **Rule IV**: justify
  - Justify the job profiles (resolution, integration time, physics…).
    « I want the biggest » is not considered a valid justification
  - Show a scaling curve (with the same code, the same physics and a similar machine…)
  - If doing a parametric survey: justify why you need 100 runs!
  - Check the configuration of the machine (do not apply to a GPU machine if your code is CPU only)
  - Storage has a cost: if you ask for 200TB, we need a justification