Past, present and future of HPC at Météo-France

F. Bouyssel, B. Devillers, R. El Khatib, F. Jac, P. Marguinaud, N. Oger, E. Sevault

21st ECMWF workshop on high performance computing in meteorology
Bologna, Italy | 15-19 September 2025

Plan

- A timeline of supercomputers and NWP models
- Previous and current reshaping of the data structure of our models
- Progress and plans on GPU code adaptation
- Progress and plans on AI for NWP
- Next HPC procurement: overview of the Benchmark 2025

Introduction

Météo-France operates a HPC to perform :

- Operational numerical weather prediction
- R&D on numerical weather prediction
- Climate studies
- Process studies
- Other activities (external partners)

2 HPC, 2 implementations

In operations since February 2021
No upgrade during the 4 year contract

Each HPC: ATOS BULL Sequana XH2000
2292 computing nodes

10.39 PFlops peak performance

2 AMD Epyc Rome processors with 64 cores at 2.25 Ghz

=> Five fold increase in performance than the previous HPC

- In 1988, ECMWF and Météo-France launched collab on **IFS/Arpège** global modelling project -> VAR DA
- In 1990, **Aladin**, extension of the Arpège-IFS project aimed at modelling the atmosphere over limited areas
- 37 years later, the IFS/Arpege/LAM code is still alive, **cooperation between ECMWF, Météo-France and ACCORD partners continues, a ‘cycle 50’ of the common code released**.

Operational NWP systems

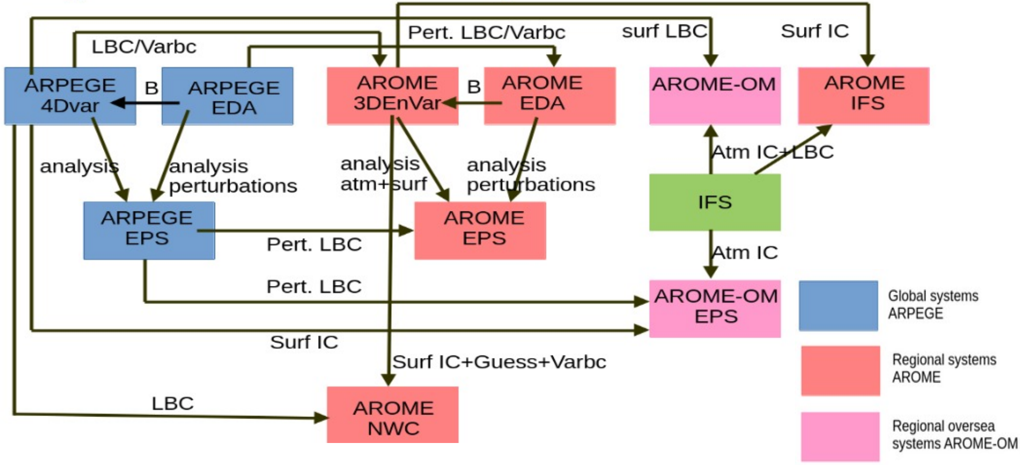

Computing power and storage capacity : a timeline

CRAY C98-4 > 98-8 FUJITSU VPP700E > 5000-31 > 5000-128 NEC SX8R > SX9 BULL DLC IVYBRIDGE > BROADWELL ATOS Sequana XH200

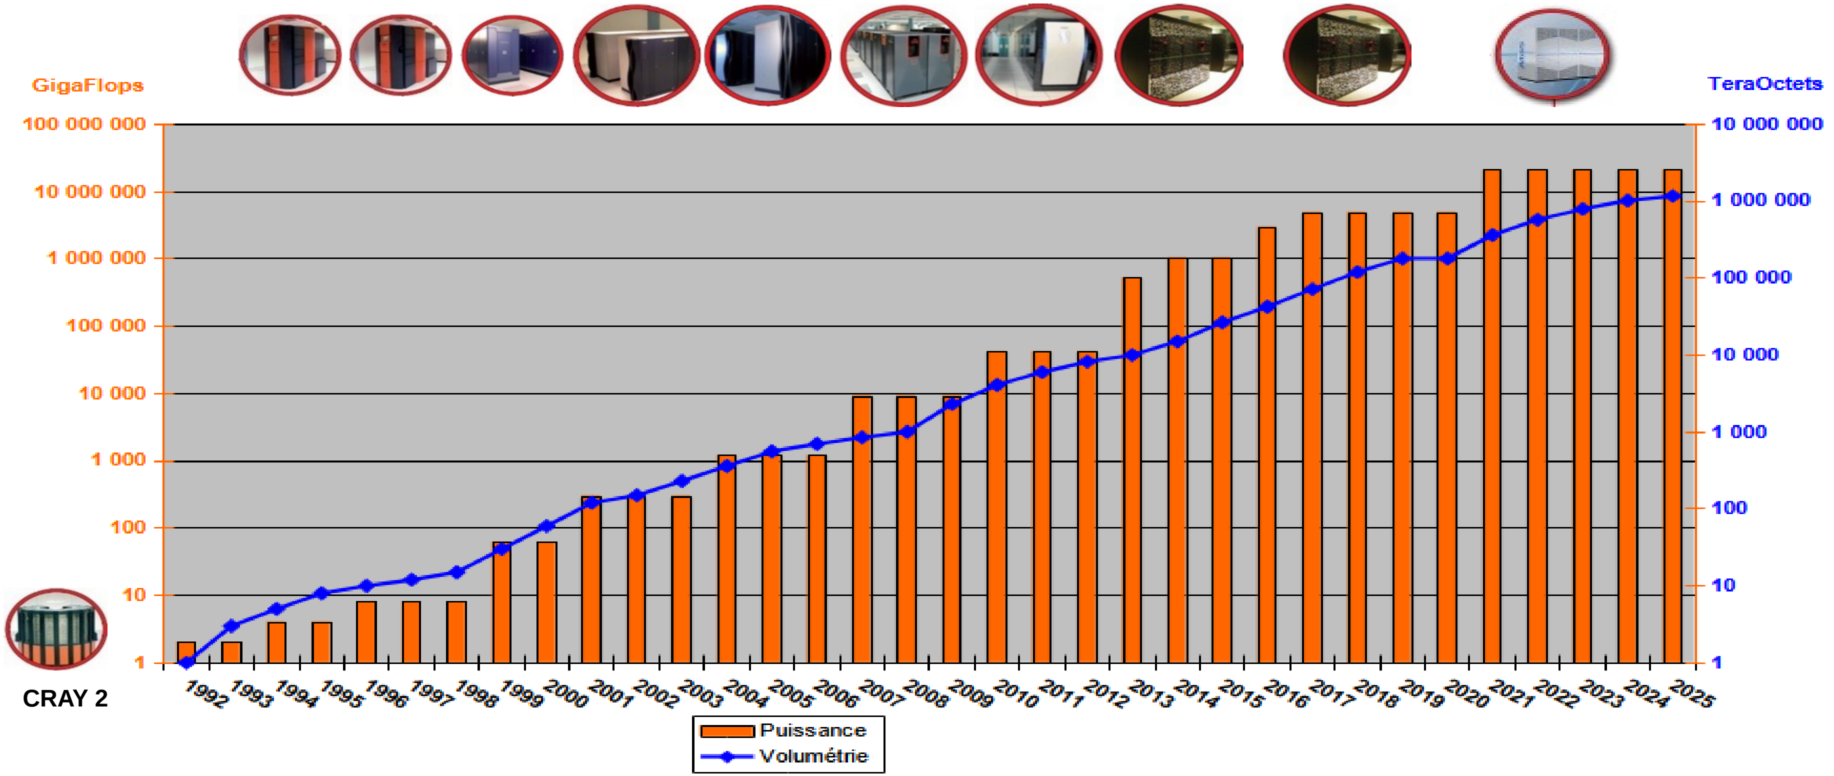

Evolution of NWP models,

but also data assimilation systems, global and regional climate and seasonal models (not shown)

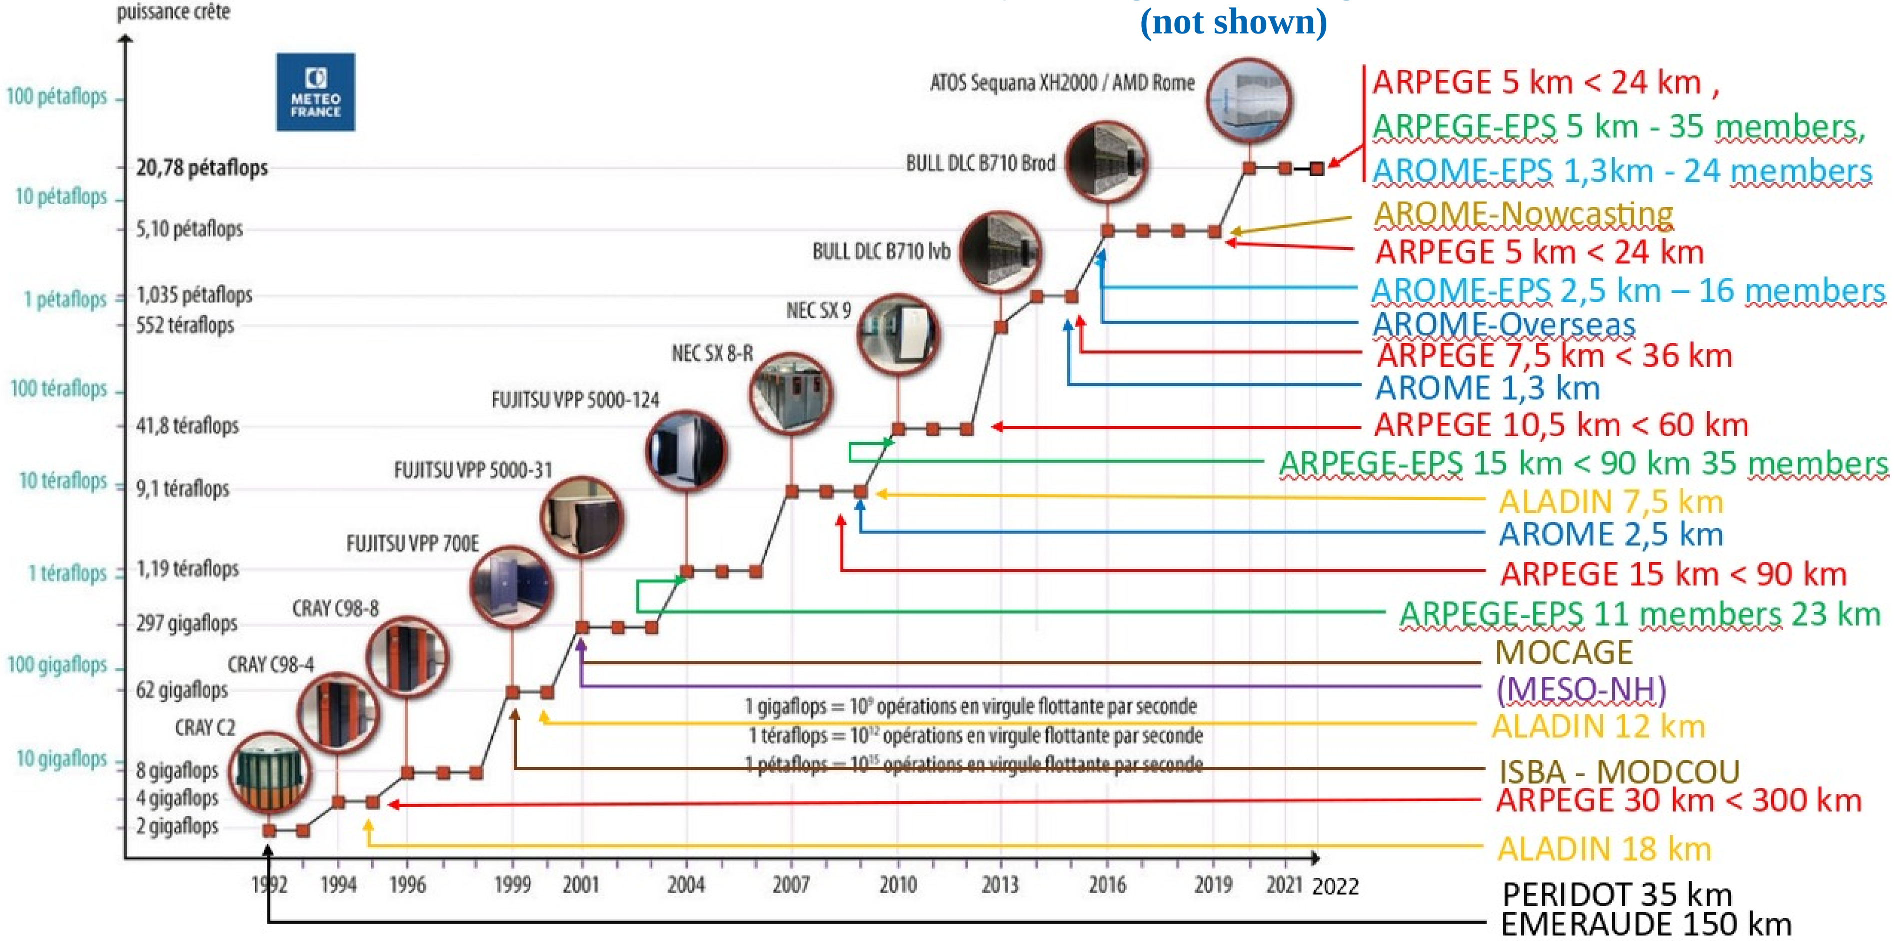

Programming for sustained performance since early Cray

Getting performance had not only lead to code tuning but from time to time large reshaping of the data structure of our models.

- Vectorization
  - Compiler directives
  - Large computations in row
- Moderate Parallelism and Vectorization
  - Still, some compiler directives
  - MPI library
  - One or two levels in the transposition strategy
- Massive Parallelism and Multi-Threading
  - Transposition strategy in every dimension
  - OpenMP directives
- GPU Programming
  - OpenACC
  - Dedicated libraries and languages
  - Object-oriented encapsulation of fields

The code itself has grown considerably since the early days of the launch of the ARPEGE-IFS common framework, up to more than 4M lines of code. Both the language and the overval organisation of the source code (modules, versioning, etc.) had evolved :

- FORTRAN for ever !
  - Starting FORTRAN 77 + DOCTOR norm
  - No more implicite types
  - FORTRAN 90, 95, 2003, 2008
- Some associated languages
  - C -- for example with ODB
  - C++ for OOPS control level
  - ksh, bash, Perl and now Python for scripting
- Version control
  - ClearCase >> Perforce >> Git
- Not to mention the never ending nightmare of compilation… from command line to Cmake.

ARPEGE forecast prototype on accelerators

Goals :

- Increase computing power
- Cap power consumption & TCO
- Increase competition between HPC & processor manufacturers
  - → Important for Météo-France call for tender
  - → Start porting ARPEGE global model in 2021

How ?

- Refactor the code
- Manage field data with Field API
- Apply coding norms
- Use source to source translation tools (fxtran/loki)
- Keep the code vendor agnostic
- Work with our partners (ECMWF, ACCORD, NVIDIA)

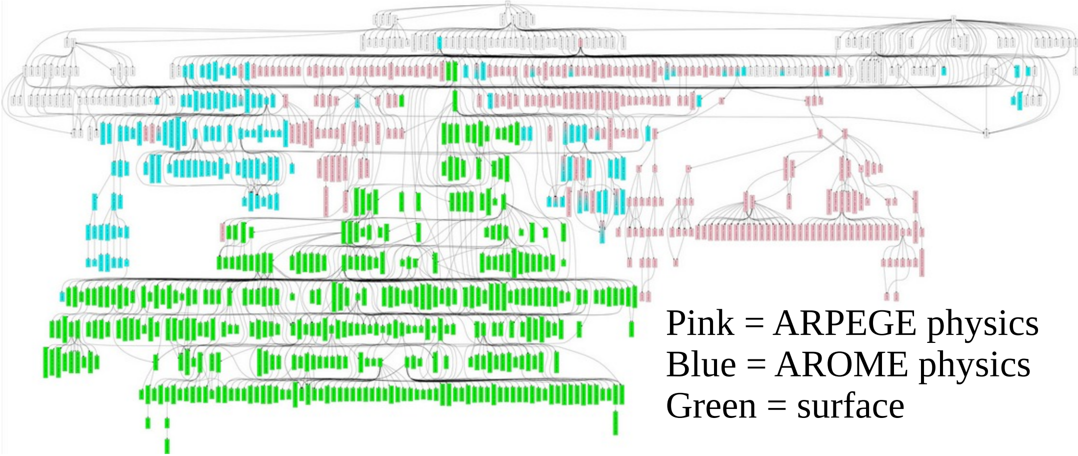

Pink = ARPEGE physics
Blue = AROME physics
Green = surface

ARPEGE forecast prototype on accelerators

ARPEGE global model, informative benchmark used in MF call for tender:

- Full time step on device
  - Grid point calculations refactored and transformed using source to source software
  - Use ECTRANS & ECRAD, both ported by ECMWF
- Communications using MPI CUDA aware
  - GPU ↔ GPU (ECTRANS, Semi-Lagrangian, Semi-implicit)
  - GPU → CPU (IO)
- No transfer of field data
- Computations on device = 90 % of elapsed time

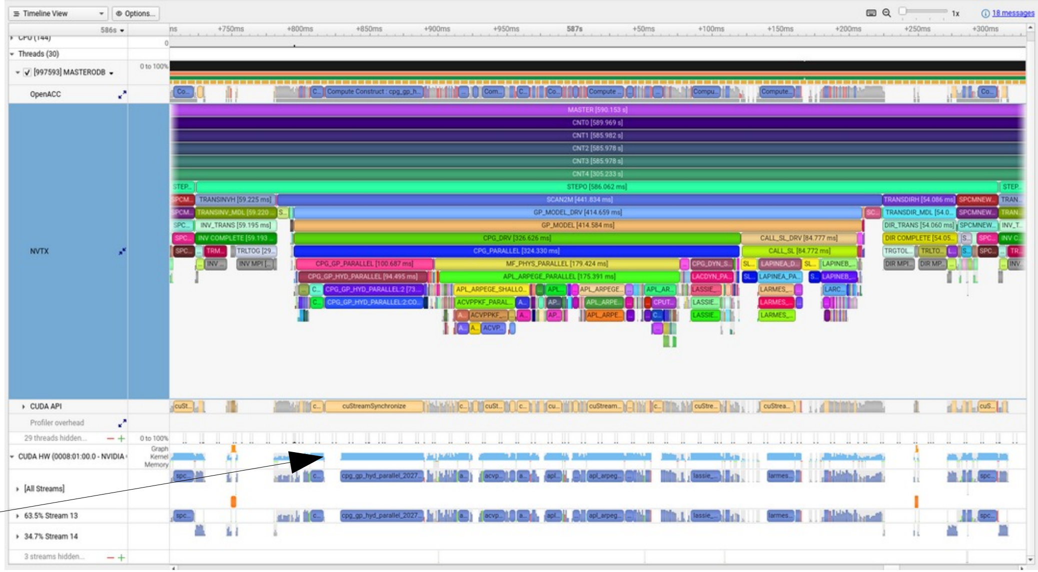

Blue = the GPU is working

(Credits: NVIDIA)

Mid-term plans

- Optimize ARPEGE global model
- Finalize refactoring
- Progress on AROME regional model with our ACCORD partners
- Integrate source to source transformations in our build system
- Port our surface scheme (SURFEX)
- Produce an ARPEGE configuration for operational use in two years

AI Weather Prediction with anemoi

Experimental Production

AIFS	BRIS	2025 Arome-IA	2026 Arome-IA
ECMWF	Norwegian Meteorological Institute	METEO FRANCE	METEO FRANCE
Global	Stretched Deterministic	Stretched Deterministic **Titan**	Stretched Probabilistic **ARRA/ERA5**

Downscaling methods
Arpege

AIWP: mid- and long-term Objectives

- Setup and develop Operational AI Forecasts
- Using multiple sources: include observations
- Predict more variables, with a better temporal resolution
- AI for nowcasting
- Comparison: anemoi vs. Weather Generator

Next HPC procurement

- Current procurement with Atos (now Eviden) since 2020 until 2028

- The goals were defined in 2021 with a x6 throughput increase in mind to have finer horizontal resolutions and more vertical levels on physically based models

- The main focus of the next HPCs will still be “traditional” NWP, but there will be more AI workloads than today

- Purpose of a benchmark:

  - Calibrate with application performance measurements the relative power of the future machine compared to the one of today

  - Reference applications should at the same time : represent what will run on the next machine, be efficient on the current machine

Challenges of the benchmark

We want CPUs for NWP, GPU for NWP, GPU for AI, but …

- How to benchmark a NWP model on GPU if it cannot compete with CPU ?
- How to estimate the performance of a GPU against a CPU ?
- What if AI applications prevail over NWP on GPU in a few years ?
- Are we experienced enough to benchmark AI applications today ? (performance, network, IO specifications)
- How to deal with the differenciation between GPU for HPC and GPU for AI ?
- How to stimulate or maintain the competition between GPU vendors ?

  Etc ...

Strategy:

- A portability benchmark for NWP on GPU (ARPEGE)
- + a performance benchmark for GPU with AI applications
- Renewal of HPC in two phases :

1) A large part in CPU + a smaller part in GPU

2) Purchase order to extend the clusters, proportions CPU/GPU to be defined during the first phase

Benchmark content for CPU

AROME forecasting model without post-processing

- Requested elapse time : 26 min/24h. [current HPC : 100 nodes]

ARPEGE forecasting model without post-processing

- Requested elapse time : 30 min/102h. [current HPC : 165 nodes]

Hybrid 4DVAR ARPEGE (under OOPS)

- Same geometry as the forecasting model
- In double précision
- 2 minimisations en troncation T499, cubic grid
- Same observations (amounts and kinds) than in operations today
- 200 members
- Requested elapse time : 35 min. [current : 44 min/224 nodes]

Benchmark content for GPU

ARPEGE forecast on GPU (state of the art) for portability testing

- Same high geometry than the version for CPU
- Source-code ready for NVIDIA GPU chips
- Requested elapse time : 45 min/24h (Large value because it must not be a difficult target)

AIFS AI application (learning)

- Based on ECMWF AIFS model
- With a relatively small dataset (175 GB)
- Designed to test intra-node GPU performance and communications

Py4Cast AI application (learning)

- Based on Météo-France Py4Cast model
- With a large dataset (15 TB)
- Designed to test multi-node GPU performance and I/O

Timeline of the HPC renewal

2025 — 2026 — 2027 — 2028 — 2029

2025	2026	2027	2027	2028	2028	2029
Competitive dialog (3 rounds) during 2025	Notification to the winner (May 2026)	Installation of porting machine (Jan 2027 March 2027)	Installation, Acceptance cluster #1 (May 2027 Sept 2027)	Installation, Acceptance cluster #2 (Feb 2028 May 2028)	Extension, Acceptance cluster #2 (Nov 2028 Feb 2029)	Extension, acceptance cluster #1 (July 2029 Sept 2029)

< FIRST PHASE >< SECOND PHASE >

Purchase order for second phase
(Sept 2027 - Nov 2027)

Thank you for your attention