



Scientific preparation of the Jules Verne project

for the Alice Recoque supercomputer

Laurent Duflot

Co-coordinator for WLCG activities in France

Introduction

- Jules Verne project
 - Around 2020, France has decided to be candidate to host one of the exascale EuroHPC machines, this lead to the current consortium called Jules Verne
 - Lead by GENCI (PRACE partner that was coordinating HPC sites/machines in FR), CEA which will host the machine and Netherlands SURF
 - The pre-projet was lead by research agencies (CEA, CNRS, INRIA and university consortium), GENCI and the Ministry for Education and Research
 - In particular at that time a sub-project for the “applications” was created “App WG”
- To reach the exascale, the machine will be mostly GPU-based
- The App WG mandate was
 - to identify the main applications / libraries for each scientific domain which
 - Will be put forward in the proposal to EuroHPC
 - Would receive support for porting
 - Would be used as benchmarks for the machine design
 - Would in the end be the main applications to run at exascale-level on the machine
 - To understand the application needs to be taken into account by the team that work on the machine design
- *DISCLAIMER: a lot of the intermediate documents are no longer accessible, so this is from memory and the few docs I could find*

The App WG

- Indeed, most applications currently running on HPC were either not ported at all to GPU or not able to exploit the expected ratio CPU/GPU of the exascale machine
- The WG had a series of meeting to explain the goals and organise the contribution, then a series of meeting by scientific domain
 - Each candidate library / application would have to fill a form with informations like (from memory):
 - Who were the developers, are they French, from a foreign or international team, is it open source or freely accessible
 - What kind of algorithm are used, on which kind of HPC machine it is currently used
 - Does it make use of GPU? Would it run on a GPU-based machine (e.g. 4 high-end GPUs for a mid-range processor)? Does it run on ARM CPU ? which parallelism (MPI, OpenMP)
 - If not what you be the porting effort and does the team need help

IN2P3 contributions

- The IN2P3 scientific domains are similar to JENA: particle physics, nuclear physics, astro-particle and gravitational waves
- Few applications use HPC, and most at a relatively small scale compared to an HPC Tier-1 machine apart from plasma simulations and lattice QCD...
- As a consequence we and our needs are not well known in the HPC user communities (and political sphere around it)
 - These discussions were an occasion to discuss with the HPC communities
 - Invitation afterward to present LHC computing in the national HPC forum

IN2P3 contributions

- The applications we proposed
 - QDC:
 - CHROMA-WM and openQCD-saveurs (Janos, DynQCD)
 - Nuclear:
 - FUSION and ABINS
 - Laser and plasma for accelerators:
 - SMILEI (also used for astrophysical plasma)
 - Particle:
 - LHC as a general use case for data intensive usage
 - LHCb: potential use of exascale machine for second level of trigger processing (need ~ Tb/s with CERN)
 - OuroborosBEM: simulation for gaseous detectors
 - CMB:
 - MAPPRaiser
- Only one of the code was retained as a benchmark for the machine design (confidential at some point, I don't know as of now so I won't name it)

Final steps of the WG and personal view

- It took quite a long time between the deadline for contribution and the final report (published [here](#) in French)
- After several months, the next step in the exascale project was a funded project NUMPEX of ~40M€ that focuses on research and developments of the software stack for exascale and beyond and mutualized developments to help porting applications
 - WG were created but I did not hear anything from the one (on data access) I registered with
 - SMILEI and LHCb teams were contacted to participate to the application WG however none of them were considered further because too far to be used to test or use these new fundamental software brick they intend to developp
- **Personal remarks:**
 - We were not part of the HPC community and I felt we were crashing their party... they did not turn to us for aspects where we would have expertise (e.g. data distribution and access)
 - Our use cases (in particular for HEP) are very foreign to them
 - E.g. Tb/s connectivity
 - Even the vocabulary is not really common
 - There is a lot of politics and power balance involved



Questions?