Initial INFN input on the update of the European Strategy for Particle Physics: software and computing

INFN National Scientific Committees of Particle Physics with accelerators, Nuclear Physics and Computing

Contact person
Donatella Lucchesi Donatella.Lucchesi@pd.infn.it

Abstract

The INFN sees two major areas of application for high energy physics computing in the time scale relevant for the Strategy Update: the exploitation of high performance computing (HPC) and the use of Quantum Computing. The first one can be realized within few years, with high energy physics experiments becoming major users of HPC. Quantum computing is still at the level of research and development, but it can be the new disruptive technology changing the computing horizon. The exploitation of the former and the research and development of the latter are described in the document.

The path toward HEP High Performance Computing

Scientific and technological contest

The computing for HEP relies today heavily on in-house custom-built computing farms, glued by the GRID middleware as developed in the previous decades, referred as High Throughput Computing (HTC).

Experiments' Computing Models are constantly evolving, with a slow but steady transition towards a smaller list of requirements, helped, in this transition, by system virtualization. As of today, even if not predominant, Cloud access is already a reality for all medium-large size experiments.

At the same time, mostly due to concerns about the sustainability of the HEP-specific infrastructure and of its middleware, larger experiments have prototyped the utilization of HPC systems. This comes from a variety of factors:

- The availability of large sized research grants from HPC centers for HEP use cases. A typical PRACE HPC[1] system deploys resources comparable to the full HEP scale.
- The push by funding agencies and countries to avoid the deployment of two parallel HPC and HTC infrastructures. While the former is seen as strategical beyond science purposes (for technological and industrial leadership), the latter has a smaller overall strategic impact, and thus is felt as less important at continental and national scale.
- HPC systems are built with the best technology available, even before that hits the shelves. HEP systems are instead built to be economically affordable, and hence usually behind the former by one, if not two, generations. The utilization of HPC systems for HEP provides access to more advanced resources, with all the relative benefits.
- HPC systems are funded via EU specific programs (like PRACE) and, in the near future, the EuroHPC Joint Undertaking (EuroHPC JU) inititiative, while HTC systems must be deployed using the standard budget of Research Institutions, with extrapolated needs beyond sustainable levels. Even recent estimates for High Luminosity LHC (HL-LHC) largely exceed current budget levels; hence the need to access new funding opportunities is felt of primary importance.

An efficient use of HPC system by High Energy Physics experiments is not trivial, since these systems are custom-built having in mind use cases largely different from HEP ones. Science related use cases range from Lattice QCD, astrophysical simulations, material sciences, simulations of nuclear systems.

Such systems feature performant node-to-node interconnections, needed for large scale MPI tasks, scarce local scratch disk, and are not meant for accessing data residing outside the facility. Recently, HPC systems include node per node accelerator cards, in order to boost total

[1] Use CINECA Marconi as an example, http://www.hpc.cineca.it/hardware/marconi.

performance and hence global ranking. Storage systems are optimized for latency and speed, and less for total size.

On the other hand, current HEP HTC systems are built using different technical solution: node-to-node connectivity is scarcely relevant, large on-node scratch areas are important, and global connectivity is needed in order to access remote datasets. The use of accelerator cards is marginal if not absent. The HEP workflows are typically data intensive, and thus current implementations deploy large storage systems close to the computing farms. Almost 100% of the experiment's software stack is designed an optimized for Intel x86_64 architecture, definitely the best choice for affordable computing in the last decade.

While the PRACE long term roadmap beyond 2020 is still not completely clear, the brand new EuroHPC JU initiative will be active up to 2023 to fund two EU level pre-exascale[2] systems, followed shortly by an actual Exascale system. Technologies are not defined yet, but at least for pre-exascale deployments will involve the utilization of system accelerators, probably GPGPU or FPGA.

In the next future, the European HPC arena will be enriched by a number of initiatives focusing on the development of European based technology and its adoption in future HPC systems:

- the FET-HPC EU funded projects are working on the evaluation of selected technologies to implement a) low-power, high granularity, scalable systems and b) configurable, high performance computing acceleration based on FPGA. Prototypes will be soon available and open to scientific communities;
- the outputs of R&D projects will be the enabling technologies for HPC ExaScale systems to be funded in the next decade in the framework of EuroHPC JU initiative.
- the just launched European Processor Initiative (EPI), aims to define and prototype a low power high performance CPU and a computing task acceleration architecture leveraging on previous design experiences and technologies know-how from academic, research and industrial European partners.

Objectives

The panorama is quite complex. There are differences between the computing approach in HEP and in HPC, that need to be overcome. The political environment is evolving quickly towards directions that may not be suited for HEP. The major objectives are:

- current and future experiments software can be adapted to tap the enormous computing power EuroHPC JU funded systems will deploy in the next decade;
- HPC systems and HEP experiments collaborate to define a common path for the use of such systems;
- HEP experiments enter in the European arena of the HPC systems developments.

[2] http://www.etp4hpc.eu/euexascale.html

Methodology, Readiness and expected Challenges

The choice of the CPU architecture is not a real showstopper, at least in the long range. Every CPU with a supported Linux (-like) operating system, with a performant C++ compiler is in principle usable, thanks to the very wise choice of not having proprietary / licensed code in the experiment software stacks. The utilization of the on-board accelerator is more critical:

- The accelerators are assumed to account for a large part of the total computing power on those systems; not using them is simply not acceptable.
- The utilization patterns of both GPGPU and FPGA are very much divergent from standard C++ multithreaded programming; in order to be used efficiently, they need nearly full code rewriting.

Much R&D effort is needed in order to use accelerator based systems in an efficient way, one that provides a tangible benefit in terms of CPU cycles; the required competences for such optimization are beyond the standard knowledge of experimental physics users, who have been historically writing most of the experiment code bases. On the other hand, the economic return of such R&D is large when it allows to access the largest part of HPC computational power; it is hence highly suggested to invest resources on software development and user training.

From the technical point of view, given an HPC architectural design driven by the best performance, the utilization of such system for HEP, strongly depends on how flexible and general purpose the systems will be built. If the Funding Agencies want HEP to use efficiently the HPC infrastructure it is necessary that, at any level, including the European funding, few key aspects in the design are modified following the HEP requirements.

In the perspective of the effective use of future HPC systems for HEP computing, it’s clear that starting from the coming years, the HEP community should become strongly involved in this phase of future system design specification and architectural solutions evaluation. The preferred way could be a direct participation of HEP stakeholders in selected committees, like ETP4HPC (European Technology Platform for HPC) or in the various User Requirements committees of EuroHPC JU.

CERN can have a leading role to coordinate this participation as well as a suitable training program, with all the stakeholders, users and providers, aimed at raising a generation of technically savvy programmers for the most efficient exploitation of HPC environments.

Quantum Technologies for HEP

Scientific and technological contest

Quantum Technologies are emerging in research and in industry, and identify generically systems in which the peculiar characteristics of quantum systems are used to perform measurements not possible in classical systems.

At least 3 different aspects of quantum technologies can be relevant for HEP:

1. Quantum sensing, in which detectors use the peculiarities of quantum states in order to perform measurements on a physical system; even if certainly interesting, they are outside the scope of this paper;

2. Quantum simulators, which use quantum systems to simulate other quantum systems, for example imposing the same behavior of the system, at least locally; they can be used to simulate on controlled systems the behavior of particles subject to specific interactions;

3. Quantum computers, which use the peculiar capability of quantum systems to perform computations, ideally outside the possibilities of classical computers ("quantum supremacy").

Quantum simulators aim to reproduce the characteristic of a quantum system, using another quantum system as "simulator". Via the appropriate preparation of a quantum state which simulates, at least locally, the one of interest, measurements can be obtained and related to the latter. Such systems exist now for very simple potentials and interactions, but in principle the realization of emulators of systems important for HEP is just driven by technology and theoretical evolution.

The physical mechanism which makes a quantum computer potentially faster than classical counterparts is the fact that each qubit can be programmed in a coherent way with respect to the others, with a N qubit "perfect" quantum computer in principle capable of describing 2^N states at the same time. As N increases, the theoretical computing power increases exponentially.
The roadmap towards a quantum computer available and competitive for general utilization is unclear. Already now a few vendors (Microsoft[4], IBM[5], Google[6]...) provide access to small machines or emulators via Cloud like interfaces; these systems are useful for initial R&D, but too small for practical HEP applications at the moment.
The European Commission has launched the Quantum Flagship Program[7] with the ambition to drive the quantum revolution in Europe. Quantum Computing is one of the most challenging quantum technologies, still the ten years milestone foresees quantum algorithms demonstrating quantum speed-up and outperforming classical computers in operation.
It seems realistic to assume that Quantum Technologies will have a large future impact on computing at complex scale, only the current evolution profile seems outside the next decade.

Objectives and Methodology and Challenges

The objective we feel to purse is to involve HEP groups in the relevant researches, in such a way to make sure the possibilities of the systems are understood from the theoretical point of view. At the same time, test simple use cases and problems (even with zero direct advantage to HEP) on the

[4] https://www.microsoft.com/en-us/quantum/development-kit

[5] https://www.research.ibm.com/ibm-q/

[6] https://ai.googleblog.com/2018/07/announcing-cirq-open-source-framework.html

[7] https://qt.eu/

available emulators and small systems, in order to gain experience in case a major technological breakthrough happens and workable systems appear on the market. In this approach the economic investment is very moderate, but the return is, potentially, huge in case the technology panorama makes progress more quickly than expected today.

One clear use case of quantum simulator for HEP would be the realization of simulators of the particle interactions, as done today via algorithmic code (Pythia, Madgraph, Herwig, Sherpa, Alpgen, …).
While in principle the generation process, via such tools, does not necessarily scale with the luminosity of LHC, in practice the precision measurements planned in the next 10 years will need more precise simulations, with generators precise up to higher order in the resummation and perturbative orders. These are known to require larger and larger processing times on classical computing systems, and can represent a non-negligible fraction of all the HL-LHC computing resources. A quantum simulator of specific standard model processes (QCD, EM, even single diagrams) is in principle exact, and does not need the utilization of perturbations or such. While far from being possible today, it could effectively replace a part of the generation process when interfaced to standard computing resources.

As of today, the use of quantum computers in HEP can be at the level or R&D. Such computers have their strength in the simultaneous application of an algorithm of many states in parallel, in principle being very well suited[8] for applications which show combinatorial explosions. In addition, the capability to probe large portions of the accessible phase space even in multidimensional functions, theoretically allows to use quantum computers as "generic minimizers"[9].
There are many aspects specific to HEP which need more technological understanding, like the quite peculiar need by HEP algorithm to be fed with large amounts of data, of unclear applicability at the moment.

Since quantum computing requires a general and complete rethinking of the way HEP algorithms are written, beyond typical experience of HEP scientists, a program of training on quantum computer utilization and programming (being it in Cirq, Q#, QISKit or any other available or future language) is necessary and should be fostered. At the same time, a close collaboration with the experts in the field of algorithm translation / porting / design should be established, in order have guidance on the best practices and training on the best ways to use quantum computing.

In general, HEP institutes and funding agencies should ease the access to quantum emulators or early R&D systems, in order to start experimenting on real systems (although simplified) as soon as possible.

---

[8]https://indico.cern.ch/event/719844/contributions/3014804/attachments/1746826/2828452/hgray_quantumcomputing.pdf

[9] https://en.wikipedia.org/wiki/Quantum_annealing