

BDEC for Europe workshop panel discussion summary

European Exascale vision and strategy on Big Data and Extreme Computing

January 28, 2015

Barcelona

The Panel was proposed the following questions to initiate the debate:

Regarding some EESI 2014/2015 recommendations

1. On one hand, legacy codes represent a very large and extremely expensive body of codes and on the other hand revolutionary approaches based on ultra-scalable algorithms may be needed to scale in the future. How do you envision priorities and approaches to embrace (or not) these two cases?
2. Software tools are usually lagging behind the hardware but are on the critical path to developed extreme scale application. Can CoE change this, or are there other approaches to be promoted?
3. Data centric approaches are very likely to be a key asset to Exascale applications. What current bottlenecks and issues should be addressed in priority?

Regarding Big Data and Extreme Computing challenges.

1. Do you envision technology or new research that may be a game changer?
2. What action would be your number one priority to be taken rapidly to ensure success?
3. How would you measure the success of the Exascale initiative from your point of view?

Panelists

- Jean-Yves Berthou, ANR
- Augusto Burgueño Arjona, EC
- Stefan Krieg, JSC
- Jean-François Lavignon, ETP4HPC
- Bryan Lawrence, NCAS
- Modesto Orozco, IRB
- Stéphane Requena, PRACE
- Philippe Ricoux, EESI

Discussion

Jean-Yves Berthou, ANR

The past years funding, at national and international level has been allocated in even proportion. The stakeholders with the support of the funding agencies have initiated many independent projects. However, there is a difficulty to integrate the results of

these projects in a global coherent roadmap (well documented). A coordination approach including the scientific and the strategic dimensions such as the one implemented by the flagships (HPB and Graphene) would help to consolidate the achievements of the Exascale roadmap and facilitate the transformation of these R&D results to products and services reaching the European HPC market and also be preparing the European Extreme-Scale demonstrator

Jean-François Lavignon, ETP4HPC

We started the EU Exascale initiative with a bottom-up approach. Now a top-down method would help to give coherency to all the effort. In other words we need to establish kind of Exascale flagship.

Augusto Burgueno-Arjona, EC

The European HPC strategy proposes an integrated approach to the development of exascale technologies, the provision of HPC resources and the involvement of user communities. Regarding the latter, the Centers of Excellence will help user communities to fully benefit from HPC technologies and resources.. The cPPP between the European Commission and the ETP4HPC is the key instrument to succeed in this endeavor. After its first year of life, all efforts should be now focused on making the cPPP deliver on its promises.

Stefan Krieg, JSC

In my field, Lattice-QCD, one of the goals we have is to develop libraries for the community. CoEs will be a very useful tool. Connections to vendors for early tests have been a very successful approach at Jülich. In future, scalability will not be an issue for LQCD but we hope that new machines will provide improved memory bandwidth. This is currently the main bottleneck for our codes.

Jean-François commenting on Stefan feedback: Memory bandwidth is expected to go in the right direction.

Bryan Lawrence, NCAS

It's clear that moving to exascale requires revolution not evolution, but such revolution requires effort, effort that is not easily available in disciplines that have relatively small teams and operational constraints such as weather prediction and climate projection. Existing codes need to be maintained and evolved, even as the community contemplates completely new codes, and our communities are currently too small to do both effectively.

The workflow and analysis tools that are necessary to both exploit current and future computers are a problem – both in the simulation phase and the analysis phase. With small teams we need to be more efficient. The weather and climate community are hopeful that some of these problems can be alleviated in the context of a proposed new centre of excellence, where three aspects of the simulation environment will be addressed: scalability of codes, usability of the software environment, and exploitability

of the storage systems. Flexibility of storage systems will be important when we reach exascale, it is likely we will need to remove Posix, and we will have to make better use of tape systems in our workflow (as opposed to as archive systems alone).

Modesto Orozco, IRB

My community (i.e. bioinformatics) is divided into two sets of people: the one with HPC experience and the one without. There is a lack of communication between these two sets. Legacy codes are not an issue, codes are changing rapidly (e.g. every week) but this fast evolution does not allow wasting a lot of effort on performance optimization. We have huge data and we lack standard to properly handle the I/O. Data and computing issues cannot be dealt with separately and there is a lack of integrated solutions. CoE must approach the community and vice-versa. We also need to elaborate more creative protocols to access to HPC resources.

Stéphane Requena, PRACE

In the future if we consider that complexity of future architectures is managed by the software stack, interoperability of the tools is mandatory otherwise we will not use them. Vendors must help bringing open standards and tools and avoid proprietary software that increase the heterogeneity and the cost of the landscape. EU is one of the largest contributors in terms of system software tools and applications and the European scientific community has been developing many proofs of concept but they are rarely industrialized. The EU should help providing a long-term effort to develop, industrialize, disseminate and support open-source robust tools such as DoE is doing in US with BLAS, Lapack, PETSc, Hypre, It is also the only way to see these tools used by industry. CoE as well as PRACE should accompany this. From an European perspective, the success of the Exascale effort will depend on the fact that the 3 EU pillars (excellence in science, infrastructure and applications) find efficient interactions (which requires a balanced funding strategy and good tools like EXDCI to stimulate the ecosystem) and it is a good moment to think about it and make it happen.

Philippe Ricoux, EESI

HPC and Exascale are two different things. HPC is a lot about operational issues: for instance teraflop systems are mature and affordable by many groups including SMEs, while Exascale is disruption oriented and has to be invented. We don't know yet if we will be able to have many efficient codes for Exaflop machines. So approaches should differ: if classic HPC processor could be perhaps examined by a cost driven, Exascale is not yet market driven.

So, it is now necessary to well take in account these large differences in the funding decisions (by public agencies or by private).

CoE should address the R&D aspects and so R&D programs to tackle Exascale hardware and software issues together. CoE should deal with vertical approaches but also with horizontal ones, for instance to work on fundamentally new numerical methods (e.g. time discretization). Furthermore, CoE need to be multidisciplinary and to gather critical mass for each of them with the normal consequence in term of budget.

From an industrial point of view legacy codes last many years (typically they are 30 years old). Because of their uses in operations they cannot be frequently modified (and the operational people are very conservative), so we need to cohabit with the old codes while we are developing the new ones. This is an expensive and complex strategy where moving from research to operation is very difficult and validation step is key. This strategy is not only a scientific & technological one but also human resources and politic (including training).

Discussion with the Room

Augusto: We need to stick to the strategy and make it work. At this stage, debates on amounts and instruments distract attention from the common goal.

Mateo Valero: In the past we have developed processors such as the Transputers. Now EU lacks processor hardware development. The MontBlanc project offers the opportunity to develop a full software stack. HPC overall has more funding than the flagships but we need to direction and coordination. Overall we have presently a unique opportunity to succeed.

Jean-Yves: Effectively, the planets are aligned, strategy-money-wish-strengths-tools. We need to coordinate the overall, progress and make it happen.

Mark Asch: The need to exchange at international level is important. We are on the good direction. The French Ministry is participating in building an international coordination.

Jean Gonnord: Shares many of the ideas presented so far but, at EU level, we are missing the implication from the private sector and the supporting industry.

Mateo: Many ideas that have been issued in Europe are industrialized in Europe. EU supercomputer technology is possible today.

Wolfgang Nagel: is strongly against increasing the investment in hardware technology, creating software capable of exploiting all type of hardware is more important. Anyway, we don't have the needed ecosystem.

Augusto: The industry role debate is a complicated one. Presently, a set of organizational tools is in place (i.e. PPP for the HPC). We need to focus on them and with the involvement of all stakeholders get them to work efficiently.

Philippe: There is not point in caring of the origin of the hardware or software. Only, the ROI matters through efficient software capable of fully using the hardware resources.

Jean-Yves Berthou (ANR): Today is Europe but tomorrow will be world-wide. It is important to follow this trend.

Stéphane Requena (PRACE) : one of the major issues in Europe is the fact that we could have budgets for Exascale Technologies but the funding for the calls is too spread among too many projects, leading to excessive fragmentation. We need to be more selective.