Taking on Exascale Challenges: Key Lessons & International Collaboration Opportunities

Birds-of-a-Feather Session at SC15

Jointly organised by European Exascale Projects and ETP4HPC

Date: Thursday, November 19

Time: 3:30-5:00pm

Location: Room 13A, Austin Convention Centre

Report

Summary

The popularity of the session exceeded the expectations of the organisers. Close to a hundred people attended. The main focus of the event was a presentation of the current European HPC Technology R&D landscape, including the details of the HPC technology projects and their potential for international collaboration. In the second part of the session, a discussion took place on the areas that could benefit from relationships with partners from other regions as well as the tools needed to facilitate this process. A Handbook including the details of the European HPC technology projects was made available at the event and has been sent to all those who submitted such a request. It is also available at: http://www.etp4hpc.eu/european-exascale-projects-2/.

The fragmentation of the European HPC resources was identifies as a key issue. However, efforts should continue to leverage Europe's skills and form structures to facilitate international collaboration. As the first step, it was agreed that this discussion should continue and another workshop should be organised at the next SC.

The presentations shown at the session are at: http://www.etp4hpc.eu/european-exascale-projects-2/.

Panel discussion

The panel was moderated by Sai Narasimhamurthy from SEAGATE and it included the speakers who had presented the European projects (Filippo Mantovani, Jesus Carratero, Stefano Markidis) and also three invitees:

- Mitsuhisa Sato - Co-project leader of Post-K
- Franck Capello - INRIA Senior Computer Scientist, Co-Director of the INRIA-Illinois Joint Laboratory on PetaScale Computing
- Eric Van Hensbergen – Senior Principal Research Engineer, ARM

F. Cappello and M. Sato provided short complementary statements before the discussion itself.

There are three main challenges to be faced:

- Load balancing
- Resiliency (not only h/w failures but also bugs that might cause systematic errors)
- Plateauing of the storage bandwidth

This puts a lot of pressure on runtime systems and application developers.

Other challenges are:

- Memory performance limitations.
- The importance of node performance increase.
- There will be no silver bullet for power efficiency improvement. 50 GF/W however seems reachable (10 nm FinFET in 2018-2020).

In Japan it is considered that the co-design of HPC must optimise and maximize the benefits to cover as many applications as possible. This is not focused co-design like for embedded systems, some generality of use must be targeted (value for public money spent must be widely visible and shared for such a national flagship project). ROI and trickle down effects to more mainstream computing should be improved in relation to K and initial PostK projects. Software lives longer than hardware, it is easier and more worth sharing and more "generic".

Some examples of successful international cooperation:

- G8 ECS (Toward Exascale Climate Simulation); this was rather a one-shot project, it proved difficult to have funding agencies from different regions to synchronise themselves and it seems difficult to be repeated
- JLESC: UIUC/NCSA, Inria, ANL, BSC, JSC, Riken - Here the return of experience is that the added value is common ground of understanding and pieces of shared methodology, rather than codes or papers
- US DOE-MEXT on system software

Questions and Answers

The following is a summary of the proceedings of the session:

Question 1 – What are the biggest general challenges in the process of achieving Exa-scale capabilities?

Answers:

- **Prototyping**, i.e. putting together elements into a platform and make it work
- **Software lagging behind** hardware development on way to Exascale
- **Applications finding it hard to break** the Exa-scale barrier (the first application breaking it will be MPI – even though many people think MPI should retire; PGAS (GASPI), also task-based programming are promising)
- **EU needs to be taking risks** that other regions would not take. (this is the way we could compete as the other regions are better funded)
- **Improving the efficiency of the systems** will be key (not only peak performance but rather the whole system design)

Questions 2: - What are the technical Issues on the way to Exa-scale?

Answers:

- **Software and programming environment**
- We don't have a good idea of how **the final architecture** will look like (how the various elements such as shared and distributed memories and compute elements in hierarchies, performance tools, resilience solutions will be used)
- **Good performance analysis and profiling tools are needed** and some are available in EU (e.g. Scalasca, Allinea tools and suites)
    - Obtaining and analysing data
    - Vampire Uni Dresden is a good example of a very innovative filtering technique
    - Which data is more important than other data?
    - Strong effort needed in Europe to look at innovative and disruptive technologies
    - EU is lacking hardware companies, which makes it difficult to carry out hardware co-design? Reply: **there are a lot of companies doing hardware in Europe and it is possible to have people working on the hardware in Europe**

Questions 3 - How are those projects coming together?

Answers:

- The first set of projects will include a constant exchange between ETP4HPC and EC to filter the best ideas and streamline/leverage results from the first round of H2020 projects
- In 2018, we might have Extreme-Scale Demonstrator projects – with an aim to show-case prototypes at a certain scale and integrate the best results of 1st round projects into several larger projects. This is described in the ETP4HCP 2015 SRA (now available at http://www.etp4hpc.eu/strategic-research-agenda)

Question 4 - What are the political inclinations and roadblocks?

Answers:

- In Europe, the fragmentation hinders a bit development; in the US there is way less fragmentation (e.g. chip development costs a lot and an integrated effort is needed)
- We should focus on some key areas and let go some other areas
- Alignment of different funding across continents; there were good projects, not continued
- Hard to have common calls in Europe
- Support actions in the EU may help cover some of these fragmentation issues (eg: EXDCI)
- IP issues are intrinsically difficult as well
- On software side we would need open common framework (global, not only European)
- In terms of the operating system we have Linux which is very successful
- We might not need to agree on one standard for all the elements
- 'Tax payer' support is needed, widespread usage and value to be demonstrated (to prove the legitimacy of funding efforts)
- We still have to address the topic of data locality
- We need to think about a more heterogeneous and complex approach and we need to do this in a more open way

Question 5- Are e.g. Human Brain Project and SKA connected to those research projects?

Answers:

- There is collaboration between the technology projects and the flagship project. E.g. - Francois: the Human Brain Project is managed in by Jülich Supercomputing Centre and there is a connection through the research team in Jülich.
- ETP4HC is in contact with SKA (in terms of their technological requirements to be included in the ETP4HPC roadmap)
- ETP4HPC – BDVA (Big Data PPP) dialogue is also well-established within EU and will be continued

Question 6 - Not just technology but use case itself plays a big role; is there any specific focus on specific application domains?

Answers:

- Centres of Excellence have a domain focus and they involve the key players in the global EU landscape – we should build on that expertise