T2.1

State of the art on funding agencies

T2.4

Towards a public collaboration

Final activity status report

Jean-Yves Berthou, Thierry Bidot, Franck Cappello, Sergi Girona, Philippe Ricoux

July 31th 2015

Introduction

These tasks are dedicated to the participation and the organization of International Exascale meetings between public and private stakeholder's organizations. The different actions envisaged initially were:

- Investigate funding structures and strategies of public bodies
- Maintain a global network of expertise and funding bodies
- Participate and organize international Exascale meetings between public and private stakeholders
- Constitute a small group of recognized worldwide Exascale leaders
- Propose governance structure, rules and missions for an International Exascale Software Initiative
- Look for potential collaboration and or coordination between Exascale centres
- Look for potential collaboration and or coordination on Exascale research activities, development projects and educational programs
- Establish and maintain a specific link with EC

Survey

A meeting happened in March 2013 to work on funding agencies cartography. A set of criterion to allow the review of what is planned and on-going at worldwide level has been defined. A questionnaire has been sent to some funding agencies from spring 2013 to spring 2014. However the collected information was too much heterogeneous and not as valuable as the one collected during the BDEC workshops we organized and presented on the EESI2 and BDEC web sites. Thus, they are not presented in this report.

International collaboration between agencies, what we learnt from the G8 call "*Interdisciplinary Program on Application Software towards Exascale Computing for Global Scale Issues*" experience

Jean-Yves Berthou organized an international workshop: Preparing HPC Codes and Software for Exascale Computing - Early results of the G8 Exascale Projects - Workshop - November 12, 2012 - Salt Lake City, USA. This meeting allowed getting some international contacts with key people. A second and final conference was organized in June 2014 by DFG in Princeton/USA with the collaboration of ANR, NSF, EPSRC, JSPS, and RFBR in order to evaluate the major results of the call regarding scientific outcomes as well as added values of the multilateral call.

The conclusions of the final conference were the following. The call topic was chosen for its potential to bring together research teams from different countries. Since challenges in high performance computing demand teams from different disciplines, countries need to join forces to tackle them. The multilateral call was issued in 2010. A highly competitive review led to the selection of six projects, which started their work in spring 2011. The projects reported on the results in a mid-term review meeting in November 2012 and in a final review meeting in June 2014. The participating members of the external experts judged the scientific outcomes with respect to the number and quality of peer-reviewed publications as well as the added value of the projects with respect to scientific networking aspects to be excellent. This reflects the highly competitive review process, which only outstanding research consortia successfully passed. Due to the high profile international nature of the call, the projects even attracted non-funded participants from other countries, thus further adding value to the overall project outcomes and stressing the importance of such multilateral initiatives.

The scientific results in most projects were very much associated with the multidisciplinary and multilateral nature of the call, which brought together top scientists who otherwise would hardly have had an opportunity to do research as a group. The initiative was a unique opportunity to enable

collaborations between 3 or more different countries (4 projects even included 5 or 6 countries). This exceeds known bilateral or regional international collaborations.

For some research areas this multidisciplinary and multilateral approach can be a prerequisite for successful projects. For example, the Ingenious project on the modelling of dynamics of large biomolecular systems needed both the hardware expertise of a Japanese team and the domain knowledge of the UK and Russian teams. This is a particularly pertinent example as the problem being considered was very unlikely to be tractable without specialized hardware being developed. The multinational nature of the teams facilitated mutual exposure to new software tools and ways of working. Project ECS demonstrated this particularly well with an EU-developed software tool ending up in the development process in a US lab. In turn, the new US users now drive improvements to the tools as well, thus helping the wider community that uses these tools.

In addition, a sufficiently wide international collaboration of key users can help drive standards (e.g. for data formats) that can be of wide importance for the whole domain community. Comparison of modelling approaches was a common feature of many projects. With many new codes being created to address the Exascale challenge, understanding the strengths and weaknesses of different approaches at an early stage is beneficial to all. Doing this with real software codes on common benchmark problems gives more insight than a paper-based comparison. Different cultural backgrounds introduced new ideas and broke down established methodologies to gain better insights and helped reducing silo effects. In addition to enabling scientists to network and collaborate, access to top HPC infrastructures was a further benefit. Through their team members, projects had access to several of the leading worldwide HPC systems that would not necessarily have been available to the team members without the collaboration.

The external experts and project leaders gave helpful feedback on the overall structure and administration of the call. Apart from the general praise of the initiative enabling collaborations that otherwise would not have been possible, some remarks will be useful to improve future initiatives. For example, ensuring synchronous project starts in all participating countries and providing a framework for the handling of intellectual property would have eased the project starts. Most projects asked for opportunities to formally engage partners outside the G8 countries to enhance multi-national interdisciplinary collaborations (e.g., with the US National Laboratories, with Chinese or further European partners).

Based on the feedback of the project leaders, the external experts and the agency representatives with respect to the 1st G8 call, the G8 initiative in general and the Exascale call in particular can be regarded as highly successful, leading to results that would not have otherwise been achieved.

A strong point was made by the external experts that there is an obvious need to continue multilateral joint efforts (and funding) to meet HPC challenges at the Exascale. The projects of the 1st G8 call in particular would serve as excellent starting points for continued research and putting in place a follow up call. Against this background, several funding agencies (ANR, DFG, DOE, EPSRC, JSPS, NSF, and RFBR) set up the call Interdisciplinary Program on Application Software towards Big Data and Extreme Computing for Global Scale Issues (BDEC). With all agreement papers ready to sign, this initiative is now suspended since March 2014 because of the G8 political situation (Ukrainian crisis).

BDEC overview

The main achievement is the organization of a series of three International workshops entitled Big Data and Extreme Computing (BDEC) that have been organized in USA, Japan and Europe. The general focus of these workshops is the post-Peta path forward to extreme scale scientific computing, taking into account the increasing prevalence of extreme quantities of data.

BDEC workshop series motivations, goal and expected impact

In the past three years, the United States, the European Union, and Japan have each moved aggressively to develop their own plans for achieving Exascale computing in the next decade. Such concerted planning by the traditional leaders of HPC speaks eloquently about both the substantial rewards that await the success of such efforts, and about the unprecedented technical obstacles that apparently block the path upward to get there. But while these Exascale initiatives have understandably focused on the big challenges of Exascale for hardware and software architecture, the relatively recent emergence of the phenomena of Big Data in a wide variety of scientific fields represents a tectonic shift that is transforming the entire research landscape on which all plans for Exascale computing must play out. This series of workshops on Big Data and Extreme-scale Computing (BDEC) is premised on the idea that we must begin to systematically map out and account for the ways in which the major issues associated with Big Data intersect with, impinge upon, and potentially change the national (and international) plans that are now being laid for achieving Exascale computing.

In the spirit of the workshops organized by the International Exascale Software Project (www.exascale.org/iesp) and the European Exascale Software Initiative (www.eesi-project.eu), the overarching goal and expected impact of the BDEC workshop series is to recreate a community process making an international effort to create the software infrastructure for extreme scale scientific computing, which involves both extreme levels of computation and extreme amounts of data.

The BDEC executive committee is composed of:

- Pete Beckman, Argonne National Laboratory (USA)
- Jack Dongarra, University of Tennessee (USA)
- Satochi Matsuoka, Tokyo Institute of Technology (Japan)
- Yutaka Ishikawa, University of Tokyo, (Japan)
- Jean-Yves Berthou Agence National de la Recherche (Europe/EESI)
- Philippe Ricoux, TOTAL (Europe/EESI)

Attendance at the workshops is by invitation only.

Information and materials on all meetings are accessible at http://www.exascale.org/bdec/

A good international community momentum has been maintained with the right level of people in term of HPC. The Big data community is now invited and is growing inside the overall community. Connection between Big data and HPC people is not always easy to implement.

New international people are participating such as Singapore and China. A follow-up meeting will take place in China in May. Africa, Australia, South Africa and Brazil are still missing.

First BDEC workshop in USA April 30th – May 1st 2013

The first BDEC workshop has been organized in USA, the April 30th - May 1st and took place in Charleston, South Carolina. 50 people attended the meeting, half from USA, half from Europe and Japan. The local logistic expenses of the organisation of the meeting were covered by NSF. The workshop was also partly sponsored by private companies. EESI2 took in charge in charge the travel expenses of the European participants.

The main focus for this BDEC workshop was put on Big Data and Natural Sciences.

An important output of this workshop has been the proposition to build an international Exascale mini-apps library, shared between USA, Japan and Europe. This is already an ongoing effort between USA (DOE) and Japan. EESI2 decided to engage preliminary work for building such mini-apps library in Europe that then could be shared with USA and Japan. An EESI2 recommendation is related to this proposition.

Second BDEC workshop in Japan February 27th and 28th 2014

The second BDEC workshop has been organized on February 27th and 28th 2014 in Japan and took place in Fukuoka. 94 people attended the meeting, half from Japan, and half from Asia, Europe and USA. The local logistic expenses of the organisation of the meeting were covered by Japanese Public Research organizations (Riken, Kyushu University, AIST, Kyoto University, CCS/University of Tsukuba, Tokyo Institute of Technology) and supported by Japanese and US industries (see BDEC web site). The workshop was also partly sponsored by private companies. EESI2 took in charge in charge the travel expenses of the European participants.

The main focus for this BDEC workshop was put on the post-Peta path forward to extreme scale scientific computing, taking into account the full range of complicated issues that the community confronts in the era of **Big Data** (http://www.exascale.org/bdec/agenda/fukuoka-japan). The workshop started with the presentations of US (DOE and NSF), Japanese (MEXT, Univ. of Tokyo) and European presentation on future regional work plans and BDEC infrastructure plans.

The Japanese side organized a workshop on February 26th, entitled the Japanese Extreme Big Data Projects Workshop which happens a day before BDEC as an unofficial "pre-BDEC" event to introduce the status of Japanese "BDEC"-associated projects. This event was public. Information and materials on the meeting are accessible through the BDEC web site.

Third BDEC workshop in Europe January 29th to January 30th 2015

The third BDEC workshop has been organized on January 29th to January 30th 2015 in Europe and took place in Barcelona. 105 people attended the meeting, half from Europe, half from Asia and USA. EESI2 took in charge the local logistic expenses of the organisation of the meeting. The workshop was partly sponsored by private companies.

The main focus for this BDEC workshop was put on HPC software and infrastructure for BDEC, HPC and Big Data applications and applied mathematics/algorithms for BDEC.

The European side organized a workshop on January 28th, entitled the “BDEC for Europe workshop” which happens a day before BDEC as an unofficial "pre-BDEC" event to introduce the status of European "BDEC"-associated projects. This event was public. Information and materials on the meeting are accessible through the BDEC web site.

Collaboration with ETP4HPC and PRACE

ETP4HPC, PRACE and EESI share many experts and members. As a consequence natural links and relations always existed between the two organisations. A more formal collaboration has been established in 2014 with the goal to elaborate jointly the official ETP4HPC roadmap in the areas of expertise of EESI2. As a result, most of the 2014 EESI2 recommendations had been included in the ETP4HPC draft Work programme 2016/2017 proposed to the European Commission for funding.

ANNEXE

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, …. Its 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 teraflopic 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 Exaflopic 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: Today is Europe but tomorrow will be world-wide. It is important to follow this trend.

Stéphane Requena: 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.