Deliverable D8.5

Report on EESI2 Exascale Conference

CONTRACT NO	EESI2 312478
INSTRUMENT	CSA (Support and Collaborative Action)
THEMATIC	RESEARCH INFRASTRUCTURES

Due date of deliverable: 30 June 2015

Actual submission date: 30 June 2015

Publication date: 30 June 2015

Start date of project: 1 September 2012 Duration: 34 months

Name of lead contractor for this deliverable: Peter Michielse/PRACE

Name of reviewers for this deliverable: Corentin Lefevre/Néovia

Abstract: This document describes the organisation and the outcomes of the EESI2 exascale conference held on the 28th and 29th of May, 2015. About 100 people attended this final conference, which was aimed to present the results and recommendations of the EESI2 project.

Revision: Final draft

Project co-funded by the European Commission within the Seventh Framework Programme (FP7/2007-2013)		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Table of Contents

Glossary

Abbreviation / acronym	Description
ANR	Agence Nationale de la Recherche
BSC	Barcelona Supercomputing Center
CEA	Commissariat à l'Energie Atomique
CERFACS	Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique
CoE	Centre of Excellence
COTS	Commodity-of-the-shelf
DG	Directorate General
DoE	Department of Energy (US)
EC	European Commission
ECMWF	European Centre for Medium-range Weather Forecasts
EESI	European Exascale Software Initiative (Europe)
EPCC	Edinburgh Parallel Computing Centre
ESC	Exascale Software Centre (US)
ETP4HPC	The European Technology Platform for High Performance Computing
EU	European Union
EXDCI	EXtreme Data and Computing Initiative
FET	Future and Emerging Technologies
GENCI	Grand Equipement National de Calcul Intensif
HP	Hewlett-Packard
HPC	High Performance Computing
HPCI	High Performance Computing Infrastructure (Japan)
ICHEC	Irish Centre for High-End Computing
IDC	International Data Corporation
IESP	International Exascale Software Project
INRIA	Institut National de Recherche en Informatique et Automatique
IPR	Intellectual Property Rights
IRISA	Institut de Recherche en Informatique et Systèmes Aléatoires
ISV	Independent Software Vendor
JSC	Juelich Supercomputing Centre
KTH	Kungliga Tekniska Högskolan
NCF	National Computing Facilities Foundation (the Netherlands)
NNSA	National Nuclear Security Administration (US)

NSF	National Science Foundation (US)
PPP	Public Private Partnership
PRACE	PaRtnership for Advanced Computing in Europe
R&D	Research and Development
ROI	Return On Investment
SME	Small and Medium Enterprise
WG	Working Group (in EESI/EESI2)
X-stack	Exascale software stack

1. Executive summary

This deliverable reports on the final EESI2 Exascale Conference, which was held in Dublin (Ireland) on May 28 and 29, 2015. Originally planned in Amsterdam, but moved to Dublin, back-to-back to the PRACEdays15 to benefit from a potentially larger number of attendees, the final conference has brought together about 80 to 100 experts, scientists and policy makers in the areas of software development, performance analysis, applications knowledge, funding models and governance aspects in High Performance Computing, coming from several European countries and from academia or industrial fields.

The final conference had a two-fold aim: first, to report on the actual work done in the EESI2 project, resulting in a set of recommendations to the EC and European funding organisations, and second to discuss the foreseen needs of European industry and academia for the future and the way forward to meet these needs.

The first day, on May 28, was mainly dedicated to the recommendations of the various Work Packages of the EESI2 project. Within the organisation of the pillars, each recommendation was to be presented by an expert in the field, supported by user examples. Some of these user examples have been presented in a more thorough way through actual user presentations. A brief look forward into the challenge of software maturity and a possible way to deal with that in future was laid out in thoughts around the potential setup of a European Extreme-scale Software Centre. A round table discussion between experts and scientists concluded the first day,

The second day, on May 29, started with an analysis of the current situation on education in the area of HPC and related aspects. Interdisciplinary education becomes a must – but only becomes done hesitatingly. HPC must be seen in wider context of computational science and engineering. Through the topics of co-design and international collaboration, leading European industry representatives made clear that a common European approach to software development for exascale systems, as recommended by the EESI2 project is very important to competitiveness of these European industries. The EC presented its vision on HPC computing and its associated work programmes and funded projects. A round table discussion with stakeholders was organised to discuss the next steps towards exascale. The project coordinator Philippe Ricoux closed the conference with final conclusions, which clearly called on implementing the recommendations, collaboration between academia and industry, but most importantly, carve a real path that leads to concrete implemented efficient exascale applications.

2. Introduction

It is widely recognized that High Performance Computing (HPC) will be increasingly important to address global scientific, societal and economic challenges. Although the projected evolution of hardware is a technological challenge in itself, more and more concern is expressed on the ability of scientific software to efficiently use the future hardware architectures. In order to address the expectedly increasing gap between projected exascale-class hardware and efficient software usage, a worldwide effort has been set up at the beginning of 2009, which is known as the International Exascale Software Project (IESP), [1], [2].

IESP has issued for an international call to action, and has invited Europe and Asia to collaborate on the challenging task to develop scientific software, ranging from systems software to applications, such that exascale systems can be designed for efficient usage by these applications by the end of the decade. EESI is the European Exascale Software Initiative and is meant as a collaborative project which is expected to coordinate the European efforts in exascale challenges. EESI2 is the step following EESI. The objective of EESI2 is to build on the work done within EESI1 and to extend this role of external and independent representative of the European Exascale community.

The agenda of EESI2 is divided in 3 main phases, the phase 1 from month 2 to month 11, the phase 2 from month 11 to month 23 and the phase 3 from month 23 to month 30. Phase 3 has been extended to month 34. In this document, we will report on the final conference of the EESI2 project, with its results and recommendations.

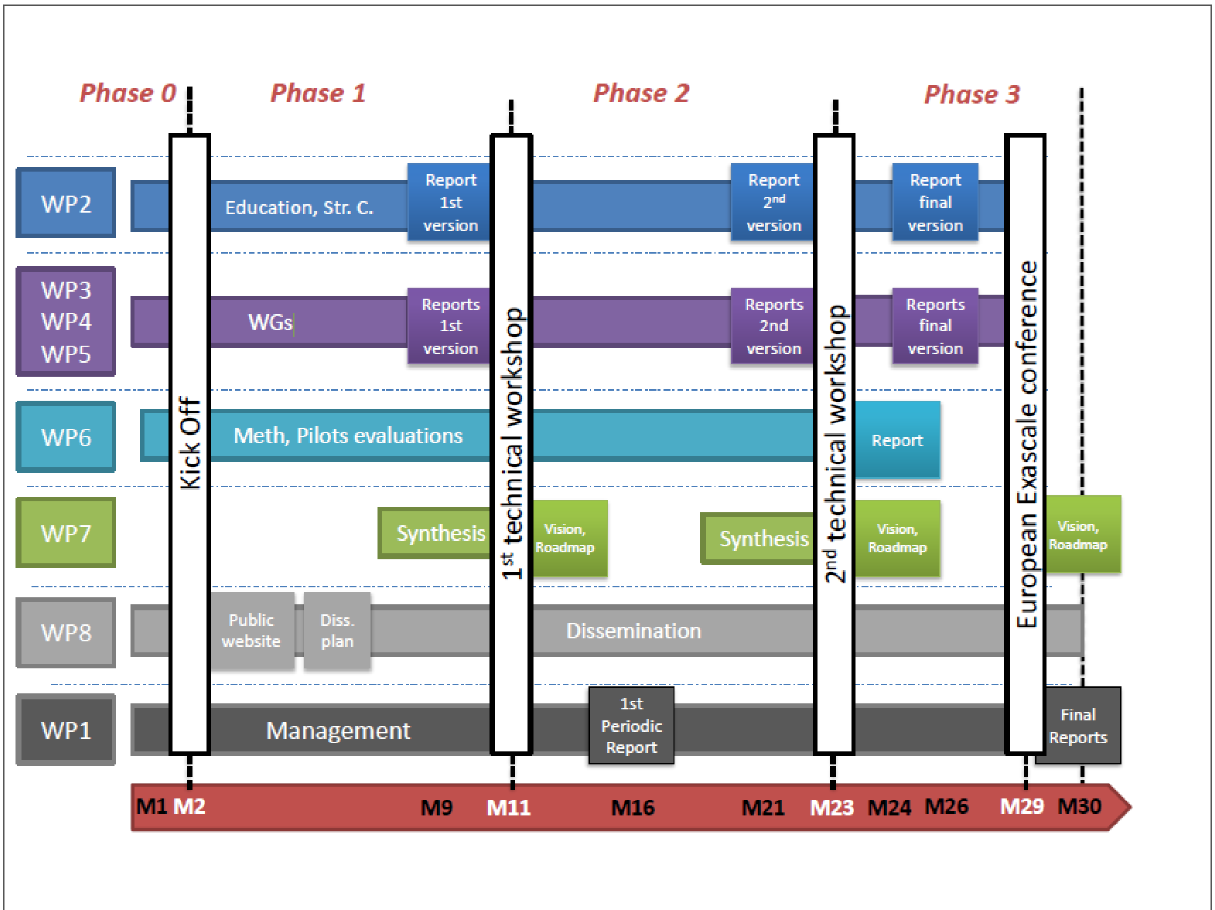

3. Setting and Agenda

The EESI2 exascale conference was organized by NEOVIA (Paris, France) and SURFsara (Amsterdam, the Netherlands) on May 28 and 29, 2015, in the Aviva Stadium in Dublin (Ireland), back to back with the PRACEdays15, in order to strengthen both meetings and to offer attendees efficient travel schedules. The number of participants to the EESI2 exascale conference was around 100. In Appendix A, the names and affiliations have been listed.

The agenda for the first day has been focussed on the recommendations of the EESI2 project in the areas of “Tools and Programming Models”, “Ultra Scalable Algorithms” and “Data Centric Approaches”. These areas are recognised in the agenda as pillars. The second day has been focussed on international collaboration and future plans.

Thursday May 28, 2015

12:30 Lunch in collaboration with PRACE

13:30 **Opening** by *Jean-Christophe Desplat, managing director ICHEC*

13:40 **EESI2, initial context, project and vision by project coordinator**
Philippe Ricoux (Total)

14:00 – 17:30 THE EESI2 RECOMMENDATIONS

14:00 – 15:00 BLOCK ON PILLAR “Tools & Programming Models”

Chair of the session: *Giovanni Erbacci (CINECA)*

14:00 High-productivity programming models for extreme computing
Rosa Badia (BSC)

14:10 Software Engineering Methods for HPC
Rosa Badia (BSC)

14:20 Holistic Approach for extreme heterogeneity management of exascale computers
Giovanni Erbacci (CINECA)

14:30 Holistic approaches to resilience
Ferad Zyulkyarov (BSC)

14:40 User Presentation on Tools & Programming Models for Extreme Computing
Camil Demetrescu (University of Rome "La Sapienza")

15:00 – 15:45 BLOCK ON PILLAR “Ultra Scalable Algorithms”

Chair of the session: *Stephane Requena (GENCI)*

15:00 Overall presentation of the 2 recommendations
Stephane Requena (GENCI)

15:05 Algorithms for Communication and data-movement Avoidance
Laura Grigori (INRIA)

15:15 Parallelisation in Time with examples on applications
Matthias Bolten (University of Wuppertal, Germany)

15:25 User presentation on Exascale Communication hiding/avoiding with examples on applications
Wim Vanroose (University of Antwerp, Belgium)

15:45 – 16:15: Break

16:15 – 17:10 BLOCK ON PILLAR "Data Centric Approaches"

Chair of the session: *Philippe Ricoux (TOTAL)*

16:15 Pillar Data Centric Approaches – recommendations overview
Philippe Ricoux (TOTAL)

16:20 Uncertainty Quantification in massively parallel codes
Vincent Bergeaud (CEA)

16:30 In-situ Data Processing in Extreme Simulations
Philippe Ricoux (TOTAL)

16:45 User presentation on Data Centric Approach in turbulent flows
Florent Duchaine (CERFACS)

17:00 Efficient Couplers for Extreme Computing
Florent Duchaine (CERFACS)

17:10 – 17:30 BLOCK ON "Software Maturity Vision"

17:10 Operational software maturity-level methodology for Exascale codes
Bernd Mohr (JSC)

17:30 ROUND TABLE on the EESI three pillars recommendations

Moderator: *Francois Bodin (IRISA)*

Panelists: *Serge Petiton (University of Lille 1), Philippe Ricoux (EESI-TOTAL), Laura Grigori (INRIA), Matthias Bolten (University of Wuppertal), Wim Vanroose (University of Antwerp), Camil Demetrescu (University of Rome) , Gabriel Staffelbach(CERFACS)*

18:30 PHOTO

18:40 End of day 1

19:30 **GALA DINNER**

Friday May 29, 2015

8:40 **Welcome & Introduction**
Philippe Ricoux (TOTAL)

8:45 – 9:45 BLOCK ON "International eco-system"

Chair of the session: *Bernd Mohr (JSC)*

8:45 Training & Education
Michael Hanke (KTH)

9:05 Co-Design for Exascale, State of Art and vision
Nick Brown (EPCC)

9:25 International collaboration and BDEC (Big Data and Extreme Computing)
Jean-Yves Berthou (ANR)

9:45 – 10:15: Break

10:15 – 11:15 BLOCK ON "Industrial Vision on EESI and Exascale Applications"

Chair of the session*: Jean-Yves Berthou (ANR)*

10:15 Three European Industrial accounts and expectations towards Exascale

AIRBUS Group: *Eric Duceau*
ENI: *Keld Nielsen*
ECMWF: *Peter Bauer*

11:00 Three Vendors / Sponsor accounts

INTEL: *Marie-Christine Sawley*
CRAY: *Harvey Richardson*
HP: *Patrick Demichel*

11:15 – 11:45 BLOCK ON "Beyond EESI"

Chair of the session*: Jean-Yves Berthou (ANR)*

11:15 EC Strategy and Plans
Panagiotis Tsarchopoulos (EC)

11:35 Next EU Initiative, EXDCI
Sergi Girona (PRACE)

11:45 ROUND TABLE: What Next towards Exascale?

Moderator: *Sergi Girona (PRACE)*
Panelists: *Panagiotis Tsarchopoulos (EC), Jean-Francois Lavignon (ETP4HPC), Philippe Ricoux (EESI), Francois Bodin (EXDCI), Eric Duceau (Airbus), Jean-Yves Berthou (ANR)*

12:35 **Conclusion by project coordinator**
Philippe Ricoux (TOTAL)

12:45 End of EESI final conference, Lunch

4. Report on Day 1

4.1 Initial Context

After the welcome speech of Jean-Christophe Desplat, director of ICHEC, the coordinator of EESI2, Philippe Ricoux, described the context of the European Exascale Software Initiative. While moving to both exabytes and exaflop/s, software in all scientific disciplines is not ready to use the full power of exascale systems. New algorithms and new implementations will need to be deviced, as has been proposed by IESP and EESI in 2008, and EESI2 since 2012. The EESI2 project has been tasked with building and consolidating a vision and roadmap at the European level, moving from petascale in 2017 to exascale in 2020-2022.

In order to get to recommendations in various fields, EESI2 has consulted around 150 experts and has held two technical workshops with 60-80 attendees to discuss those topics which are essential for Europe to remain a serious player in the developments to exascale computing and big data.

A general accepted power target for exascale systems is around 20 MW – the consequence of this is that a petascale system, as we have today, has a power envelope of only 20 kW. Not only incremental steps in technology are required for this, also disruptive ideas and methodologies will be required to get there.

The EESI2 project has formulated a number of recommendations with respect to software and applications development in which Europe is traditionally strong. These recommendations have been grouped in three pillars, which will be discussed and explained during the final conference. The three pillars are Tools & Programming Models, Ultra Scalable Algorithms and Data Centric Approaches. Given the recommendations, discussion will further focus on the future and the EC plans to achieve Europe’s competitiveness in this area.

At the end of his presentation, Philippe Ricoux outlined the programme for the rest of the conference. Not only were recommendations presented, also user presentations on the reasons and effect of certain recommendations have been scheduled, to illustrate their importance.

4.2 Recommendations on Tools and Programming Models

During the session on Tools and Programming Models, chaired by Giovanni Erbacci of CINECA, four recommendations have been presented. These have been in the specific areas of:

1. High-productivity programming models for extreme computing

Rosa Badia of BSC presented the challenges a user faces when programming peta and exascale systems: scalability, heterogeneous hardware and constraints in energy. The complexity of exascale systems is so large that it should not be exposed to the application. It is for this reason that EESI2 recommends to investigate the development and deployment of high level programmability/expression in a few lines of code, a high portability, powerful runtime environment, dynamic load balancing and domain specific languages (DSL).

2. Software engineering methods for HPC

Rosa Badia of BSC presented the process of development and deployment of scientific software applications: long life span, various contributors, commercial endorsement, applied by communities.

Software maintenance is costly, efficient tools for debugging and performance analysis are required. The recommendation is to develop predictive methods and tools to assist software re-design and co-design, to develop scalable debugging and performance analysis for exascale systems, to carry out a survey of current software engineering practices and processes and to develop new and efficient resilience tools dedicated to exascale figures.

3. Holistic approach for extreme heterogeneity management of exascale computers

Giovanni Erbacci of CINECA described the technology developments as a result of the flattening of Moore's law. New technology trends are in photonics, non-volatile memory, ionics to come. In order to benefit from these hardware developments, the design of the entire software stack has to be revisited. The recommendation is to investigate and develop hardware and software Application Programming Interfaces (API's) for the integrated management of heterogeneous systems and new upcoming technologies.

4. Holistic approaches to resilience

Ferad Zyulkyarov of BSC presented some insights in resilience aspects of ever-growing systems. Due to their complexity and scale, future exascale systems are expected to see faults within sub-second time intervals. Not only hardware failures but also software errors will potentially cause many down-time intervals. Current resilience techniques do not scale on exascale-sized systems, and will therefore have to be re-considered and re-designed. The recommendation is to investigate and develop API's with respect to resiliency and to investigate disruptive checkpoint/restart techniques to enable applications to cope with hardware and software faults.

Camil Demetrescu (University of Rome "La Sapienza") presented a big data analytics application to illustrate the importance of high-productivity programming tools. In the area of big data analytics, most final users are not HPC experts, and so can not go in details with respect to efficient programming, This means that the full software ecosystem need to enable these scientists to focus on their application, not on the implementation. This will lower the barrier for data scientists. The recommendations of EESI2 in the are of Tools & Programming Models are key to less experienced HPC users, as shown in the figure below.

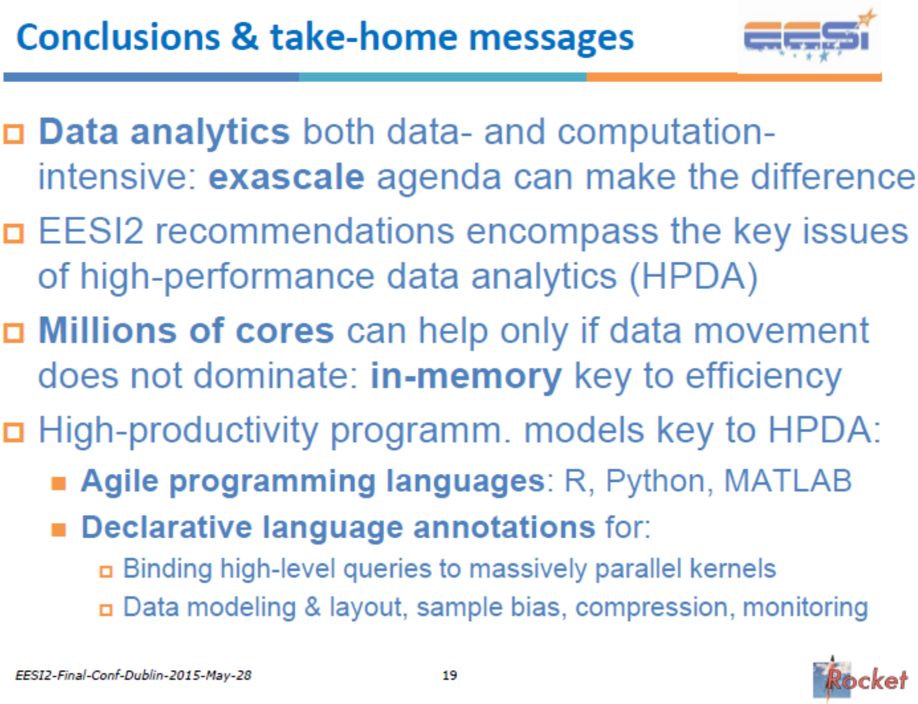

4.3 Recommendations on Ultra Scalable Algorithms

During the session on Ultra Scalable Algorithms, chaired by Stephane Requena of GENCI, two recommendations have been presented. Getting to exascale requires improvement in all areas of a computer system: interconnect, nodes, cores and energy efficiency. Data movement is both energy consuming and prohibiting scalability of algorithms. With the increasing amount of cores and threads towards millions or even billions, new methods of parallelisation of applications will need to be investigated. The recommendations on Ultra Scalable Algorithms focus on the areas of:

1. Algorithms for avoidance of communication and data movement

Laura Grigori of INRIA noted that the time to move data is much larger than the time per flop. With growing systems, in order to reach sufficient scalability, overlapping communication with computation is a key aspect. It means that already in numerical algorithm design, aspects on parallel communication must be integrated into the design process. The recommendation therefore is to design communication avoiding algorithms and concepts within the entire application.

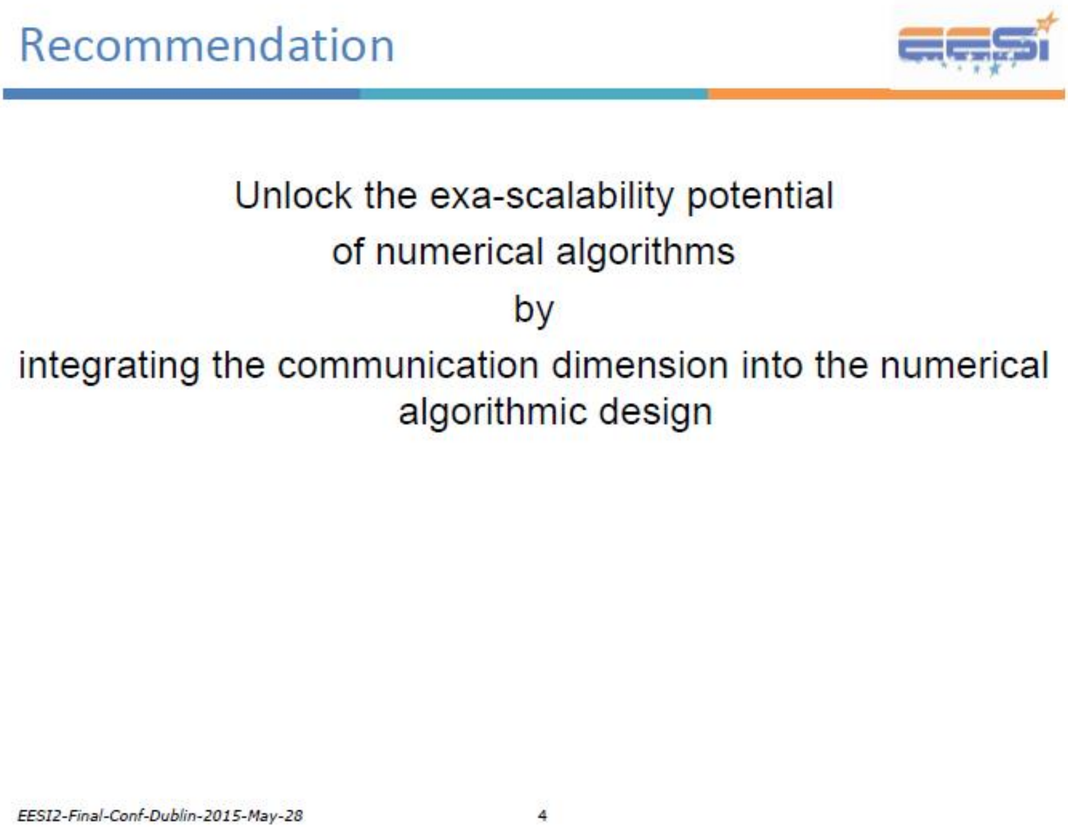

2. Parallelisation in time

Matthias Bolten of the University of Wuppertal pointed out that spatial scalability is limited by the actual grid size used by the application. So far, parallelisation is done in the spatial directions, within a single time step. The idea of parallelisation in time, so execution of more time steps in parallel and combining the numerical results in some way, will provide another direction of parallelisation, which may push the limits of scalability. Correctness aspects of the results will need to be considered when experimenting with parallelisation in time. The recommendation is to investigate and develop the concept of parallelisation in time.

Ultra Scalable Algorithms - Recommandation

« Parallel-in-Time: a fundamental step forward in Exascale Simulations »

- Rationale
  - Spatial decomposition strategies are not enough for exploiting all the massive amount of concurrency of Exascale systems
    - Use of the parallelision accros the time dimension
    - Potential application areas include: climate research, CFD, life sciences, materials science, nuclear engineering, etc
    - European researchers are leading Parallel-in-Time developments

- Goals
  - Establish of multi-disciplinary consortia to co design the deployment of Parallel-in-Time methods, encapsulated in reusable scalable libraries
  - Establish a series of benchmarks for Pro/Cons of Parallel-in-Time methods
  - Fund 2 to 4 internatonal projects between €2M and €4M

EESI2-Final-Conf-Dublin-2015-May-28 7

3. User application on communication hiding/avoiding

Wim Vanroose of the University of Antwerp started with a classic implementation of the Conjugate Gradient method, and showed in which way communication could be avoided and hided in the computational work of the algorithm. The order of calculations changes, which may give rise to a different behaviour and propagation of rounding errors, which will need to be monitored carefully.

4.4 Recommendations on Data Centric Approaches

The session on Data Centric Approaches was chaired by the EESI2 project coordinator Philippe Ricoux of Total. The importance of data centric approaches merely lies in the fact that data movement is expensive with respect to energy usage. The recommendations on Data Centric Approaches focus on the areas of:

1. Uncertainty quantification in massively parallel codes

Vincent Bergeaud of CEA pointed out the increasing importance of methods to validate the numerical results of simulations. Peta- and exascale systems enable large simulations in critical areas such as nuclear safety modelling. In order to understand and trust the results, mechanisms should be in place to determine the uncertainty of results and verification methods to validate them. With exascale, the need for these tools will only grow. The recommendation therefore is to develop verification and validation tools, accessibility of such software and to stimulate methodological progress.

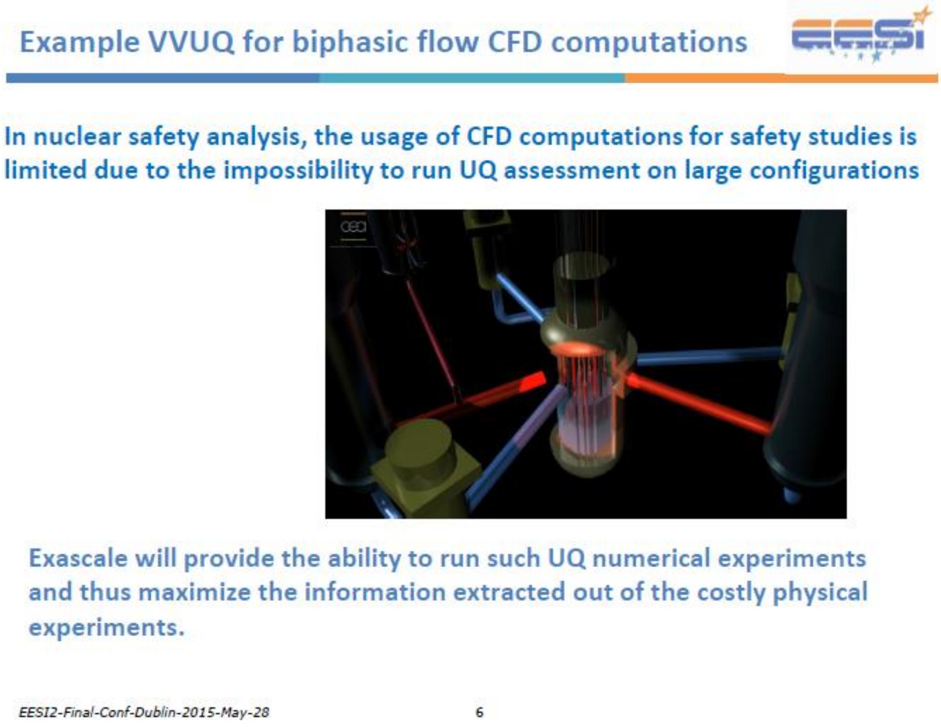

2. In-situ data processing in extreme simulations

Philippe Ricoux of Total presented the challenges in example areas as oil reservoir engineering and combustion to illustrate the need of using data once they are generated, avoiding excessive storage and data movement. At the moment, PBytes and soon EBytes of data will be generated (see figure below), for which it will be impossible to store in the traditional ways. Data will need to be processed and visualised during the actual application running, either on the same system or on systems with direct access to the data. The recommendation therefore is to investigate and develop real time Extreme Data Processing and better science through I/O avoidance in High-Performance Computing systems.

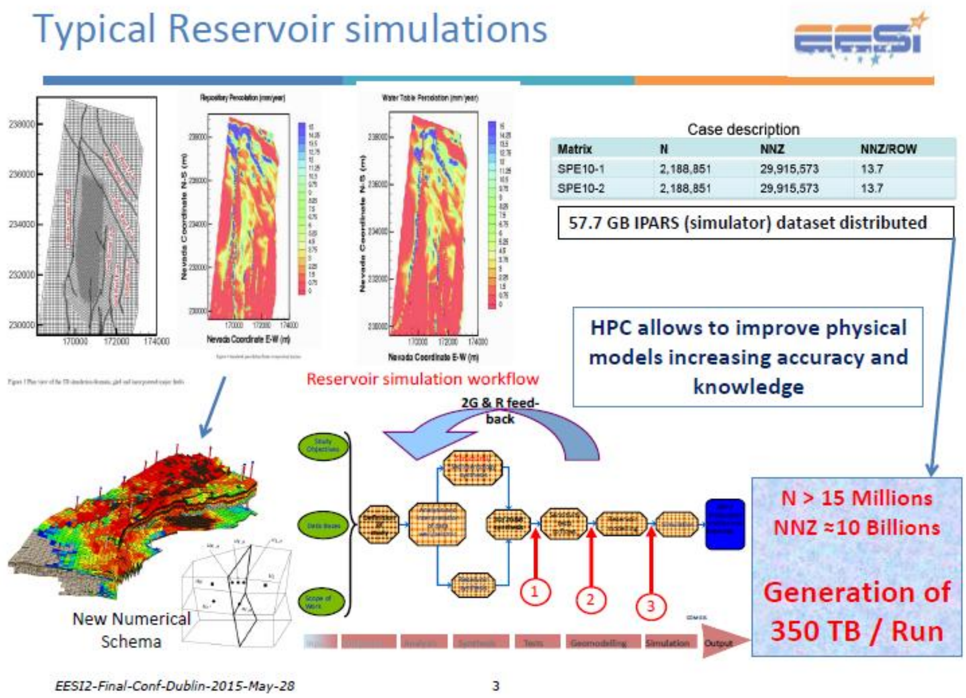

3. Data centric approaches in turbulent flows

Florent Duchaine of CERFACS discussed the challenges in modelling turbulent flows on petascale and exascale systems. Due to the tremendous size of the simulations, traditional ways of analysing data are not sufficient and are hardly achievable. Hence, new intelligent ways of extracting data are required. In turbulent flows, this may be done by tools to extract and follow automatically pertinent flow structures. The recommendation is to investigate and develop a complete toolbox which is able to extract the interesting data on the fly from applications, not only for turbulence but also applicable in other scientific disciplines.

4. Efficient couplers for extreme computing

Florent Duchaine of CERFACS presented the challenges of multiscale, multiphysics applications as one of the main drivers for exascale systems. Simulation software packages have been developed separately, while now the need for coupling various models has become obvious, in order to significantly improve simulations and hence to increase understanding of many complex processes in nature and industry. Software couplers between the various model components are required for this. The recommendation therefore is to investigate and develop all aspects of software coupling, which include assessment of best ways of coupling, especially avoiding data movement (see figure below), development of standards and integration of coupling tools in pre/postprocessing stacks.

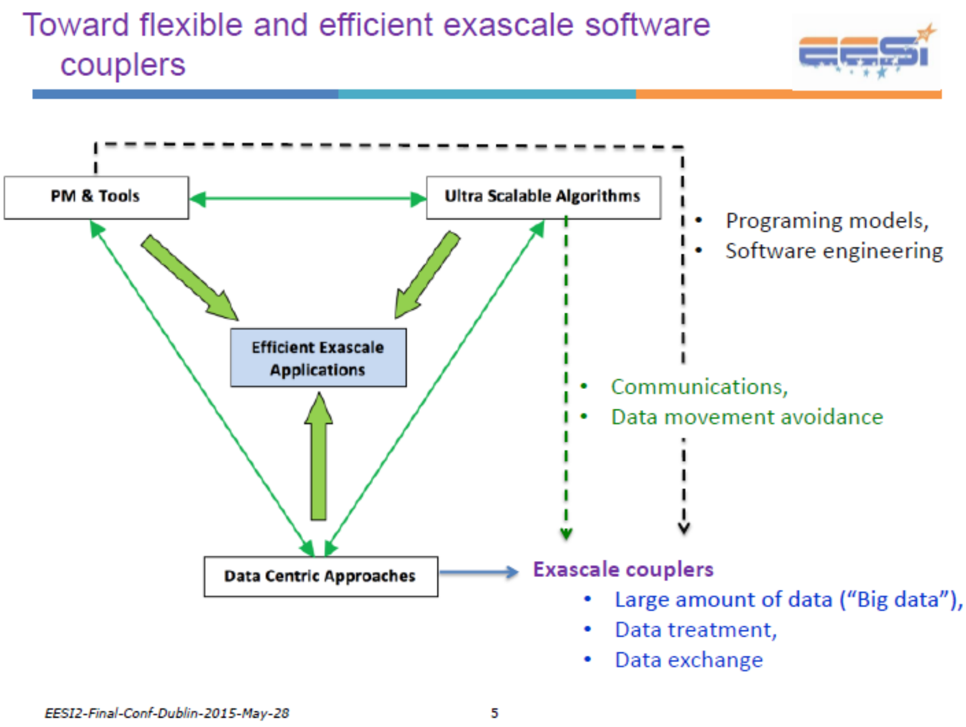

4.5 Software Maturity

Bernd Mohr of Julich Supercomputing Centre discussed the need for a European Extreme-scale Software Centre (EESC), in order to cope with the development and sustainability of software developed to cope with the EESI2-project recommendations. As the EC funds many software projects within the HPC area, it is natural to evaluate the software production and further maintenance and development. This could be reached by creating an independent European Extreme-scale Software Centre, initially funded by the EC, with sustainable funding through selling services to companies, HPC Centres and other EC-funded projects. The PPP mechanism of the EC could play a role in the creation, initial funding and monitoring of the EESC (through ETP4HPC).

Main tasks of the proposed EESC would be on coordination aspects of EU Open Source HPC Extreme-scale software ecosystem components and modules, maintenance and collection of software, and linking to vendors and ISV's. Aspects as structure, governance, access models, ambitions and further tasks will need to be worked out, there are certainly ideas on directions.

4.6 Round Table Discussion on the EESI pillars Recommendations

Francois Bodin of IRISA has acted as the moderator for the round table discussion on the EESI pillars recommendations. The panelists were:

Matthias Bolten (University of Wuppertal)
Camil Demetrescu (University of Rome)
Florent Duchaine (CERFACS)
Laura Grigori (INRIA)
Serge Petiton (University of Lille 1)
Philippe Ricoux (EESI-TOTAL)
Wim Vanroose (University of Antwerp)

Questions posed to the panel:

1. How to implement the recommendations?
2. What are the roadblocks that should be overcome to make these recommendations effective?

3. How can they boost EU competitiveness (and benefit from the EU ecosystem)?
4. Can applications achieve a high efficiency to ensure a sufficient enough ROI?

Opinions of the panel:

1. The recommendations mentioned in the three pillars are not standing on itself, but are linked. Hence, implementing the recommendations requires multi-disciplinary teams to work on this. Bringing the people together in a kind of special research area could be a way to enable implementation work on this, both by academia and industry. The opportunity to rethink the full simulation chain, from pre-processing until post-processing should be facilitated. Teams could work on the various aspects, with some way to oversee what is going on. It also would make sense to find a way to let industry better use publicly funded systems, as seems to be the case in the US and Japan.

2. Apart from funding that needs to be raised, there are more things. Bringing together multi-disciplinary teams, from academia and industry, requires processes which facilitate this. Skilled people are required to make this kind of work successful, which in turn requires HPC skills to be taught at universities and applied universities. The dream is to have HPC being used like most people use smartphones. From a technical point of view, trust in results is required. This can be approached by code verification principles. Definition of algorithm first, next is implementation on systems.

3. The panel recognized that a side effect of having a 20 MW exascale system is that there is a 20 kW petascale system – so a supercomputer in a box. SME's and industry should first have access to this type of building blocks, and then, secondly, must be able to use it as easy as possible. This requires training of people, but also usage of software which is easy to use and does not require to be an HPC expert. The ETP4HPC strategic research agenda should be leading the way here.

4. Efficient applications is key for industry. There will be no company that invests in exascale hardware if it is not clear what the benefit is, how costs are reduced and how income is increased. However, industry alone most likely cannot achieve desired efficiency levels of applications. This is another reason for collaboration between industry and academia. Aspects as hiding communication and asynchronous computing will need to be investigated and adopted. New future codes which try to overcome the usability aspect may be possible, however, will most likely not perform optimally.

4.7 Social Event at end of Day 1

Dinner was organised in be held in the Thomas Prior Hall of the Bewleys Hotel Ballsbridge.

A photograph of the attendees was taken at the end of Day 1.


AVIVA STADIUM
EESI

5. Report on Day 2

5.1 Introduction

After a successful first day the EESI2 Exascale conference will continue with its scope on international efforts and the way beyond EESI2, including insights from industry and the EC. Philippe Ricoux kicks off the second day.

5.2 International eco-system

Bernd Mohr of JSC acts as the chair for the session on International Ecosystem. The following topics have been presented in this session:

1. Training and Education

Michael Hanke of KTH addresses the issue of continuous education and teaching, certainly in the area of HPC. Not only straightforward training on HPC aspects is required, but more and more also computational science and engineering and interdisciplinary education. In the recommendations, it has become clear that combination of implementation, numerical algorithms and applications is required to have a chance of success to get to exascale performance. Industry should come in earlier in the curriculum as well. Focus must be shifted to the complexity of parallel algorithms and the real-life cost to solve a computational problem. Three levels of users are identified, ranging from required low to high knowledge: method users, method developers and core researchers.

2. Co-design for exascale: State of art and vision

Nick Brown of EPCC brought us in the world of co-design – what do people understand it is, how is it being used? Through questionnaires, interviews and project descriptions, it became clear that co-design is not well-defined but centred around collaborations and feedback. It furthermore requires communications between people coming from various different domains. These challenges have become clear in one of the US co-design centres (ExaCT on combustion), but have been solved. The Centers of Excellence, as funded now by the EC, may is some way be viewed as co-design centres in the areas of software and applications in specific scientific disciplines, in which science is the driving force.

3. International collaboration and BDEC (Big Data and Extreme Computing)

Jean-Yves Berthou of ANR described the history and activities of the exascale (software) initiatives. Through IESP, EESI and EESI2, the current global collaboration is BDEC: Big Data and Extreme Computing. This collaboration has the intent to renew the roadmap as made by IESP in 2011, and keep this up-to-date as long as possible. Currently, a series of workshop is run by BDEC to gather updated information since IESP. Apparently, there is not a lot of news from most countries.

5.3 Industrial Vision on EESI2 and Exascale Applications

Apart from the views of experts and key users, the vision and requirements of European industry are an important factor in reaching efficient usage of exascale level systems, for the benefit of European research and competitiveness.

1. Eric Duceau of AIRBUS Group explained the challenges faced by his company to remain competitive in the global aircraft industry. Simulation and the need for large HPC systems are key in coping with that challenge. As many companies, AIRBUS Group faces a challenge with respect to legacy codes, already in place for more than 25 years. Aerodynamics simulations target for full multidisciplinary aircraft simulation, including noise and targeting for real-time computational fluid dynamics in-flight simulations. The picture below shows the aerodynamics/aeroacoustics roadmap. In the area of structural engineering many ISV's are active, meshing is computationally and labour-intensive. Scalability issues and programming models will need to be addressed when moving to exascale.

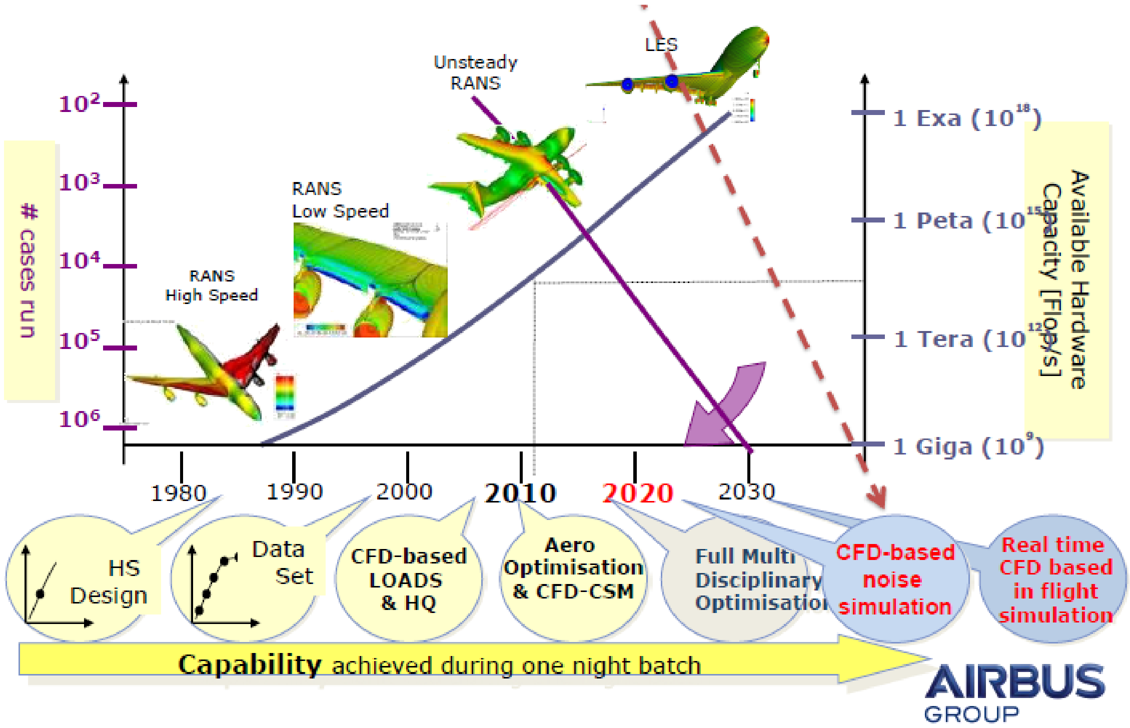

2. ENI is a large Italian-based energy company, committed to growth in the activities of finding, producing, transporting, transforming and marketing oil and gas. Keld Nielsen of ENI explained the challenges in subsea system engineering: find the field, then model the reservoir and then design and build production facilities. The figure below describes the various aspects which ENI cones across with respect to the design of subsea oil recovery systems. In each step, HPC plays a crucial role. In the design process, challenges will be in numerical methods, handling of datasets and development of multi-physics and multi-domain simulations. This requires scalable, verifiable and maintainable software and algorithms. The recommendations of EESI2 in these areas are certainly expressing industry needs.

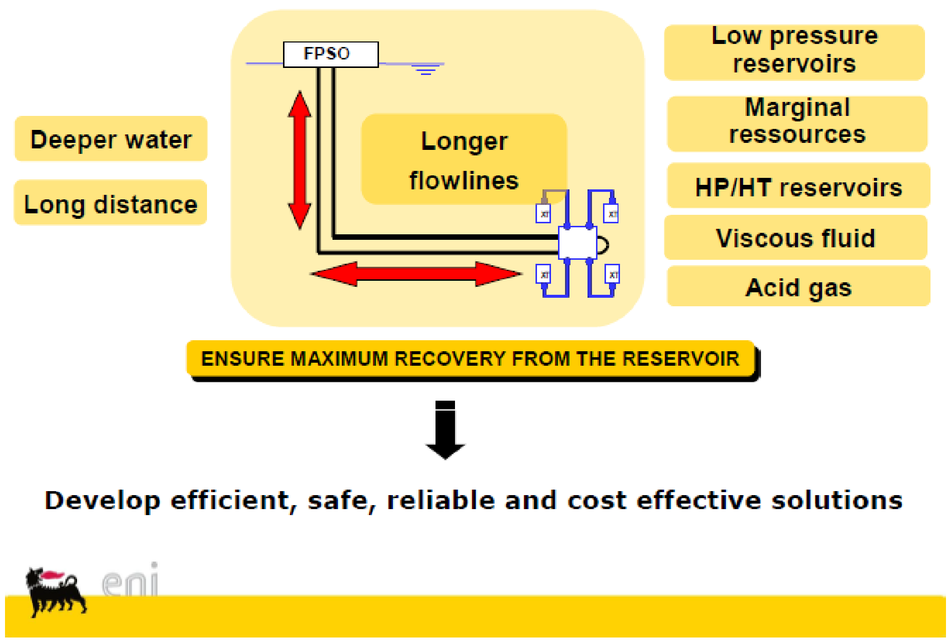

3. ECMWF is the European Centre for Medium-range Weather Forecasts. Peter Bauer indicated the current computational grid for global forecasts as 5 million grid points, to grow to 500 million grid points and smaller time steps, leading to a factor of 2000 more computational effort per time step, as is shown in the figure below. Operational requirements stay within one hour, apart from that ensemble runs are needed for uncertainty quantification. An interesting observation is that exabytes are reached before exaflops. Current algorithms, mathematics and hardware are by far not sufficient for the future plans of ECMWF. Starting from scratch to develop new codes for climate research and numerical weather prediction is unrealistic and not feasible. Implementation of the recommendations of EESI2 is crucial for Europe.

		Observations	Models
Today:	Volume	20 million = **2 x 10⁷**	5 million grid points 100 levels 10 prognostic variables = **5 x 10⁹**
	Type	98% from 60 different satellite instruments	physical parameters of atmosphere, waves, ocean

		Observations	Models
Tomorrow:	Volume	200 million = **2×10^8**	500 million grid points 200 levels 100 prognostic variables = **1×10^{13}**
	Type	98% from 80 different satellite instruments	physical and chemical parameters of atmosphere, waves, ocean, ice, vegetation

Factor 10 per day **Factor 2000 per time step**

EESI-2015 PB 052015 © ECMWF

4. Three vendors (sponsoring the conference) presented their activities within the exascale framework, focussing on exascale labs and collaborations in Europe. The three vendors were Cray, Hewlett-Packard and Intel.

5.4 Beyond EESI2

This session featured the High Performance Computing Strategy and Outlook of the European Commission and the next granted EC initiative on EXDCI.

1. Panagiotis Tsarchopoulos (Future and Emerging Technologies, DG CONNECT, European Commission) presented the three elements of the EC's HPC strategy, which are:

- HPC in FET, towards exascale HPC;
- E-infrastructures, achieving excellence in HPC applications through Centres of Excellence, and
- E-infrastructures, providing access to the best supercomputing facilities and services through PRACE.

The figure below shows the interrelation between the three elements:

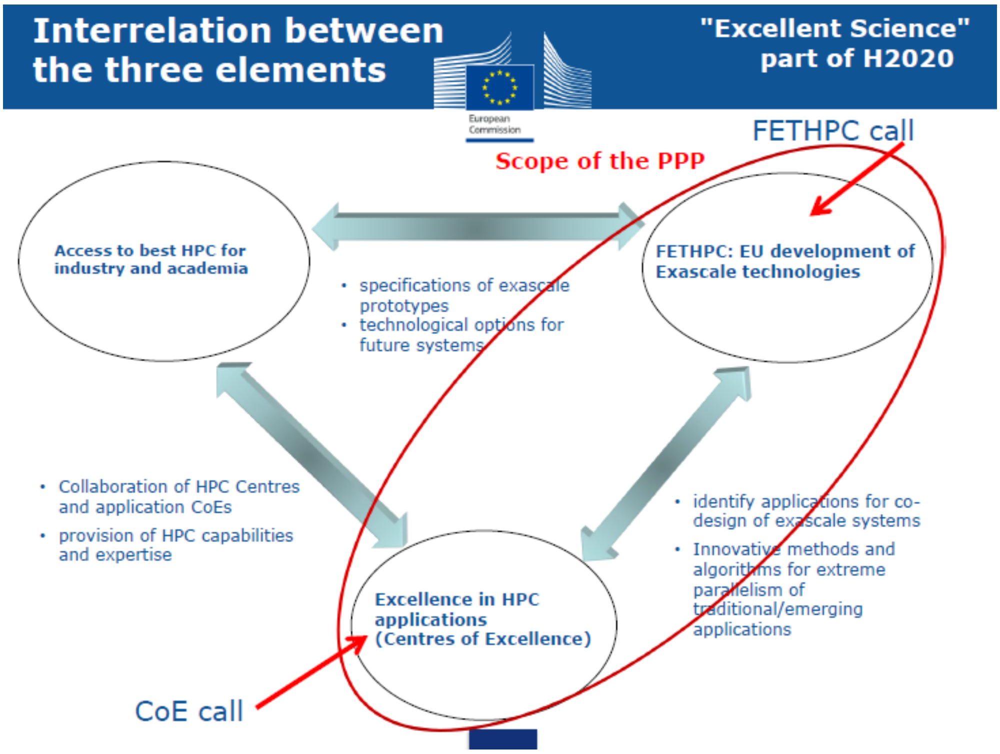

The CoE call and the FETHPC call have been closed and proposals have been retained, expected to start in the course of 2015. Both calls fit in the HPC overall strategy within Horizon 2020, as shown in the figure below.

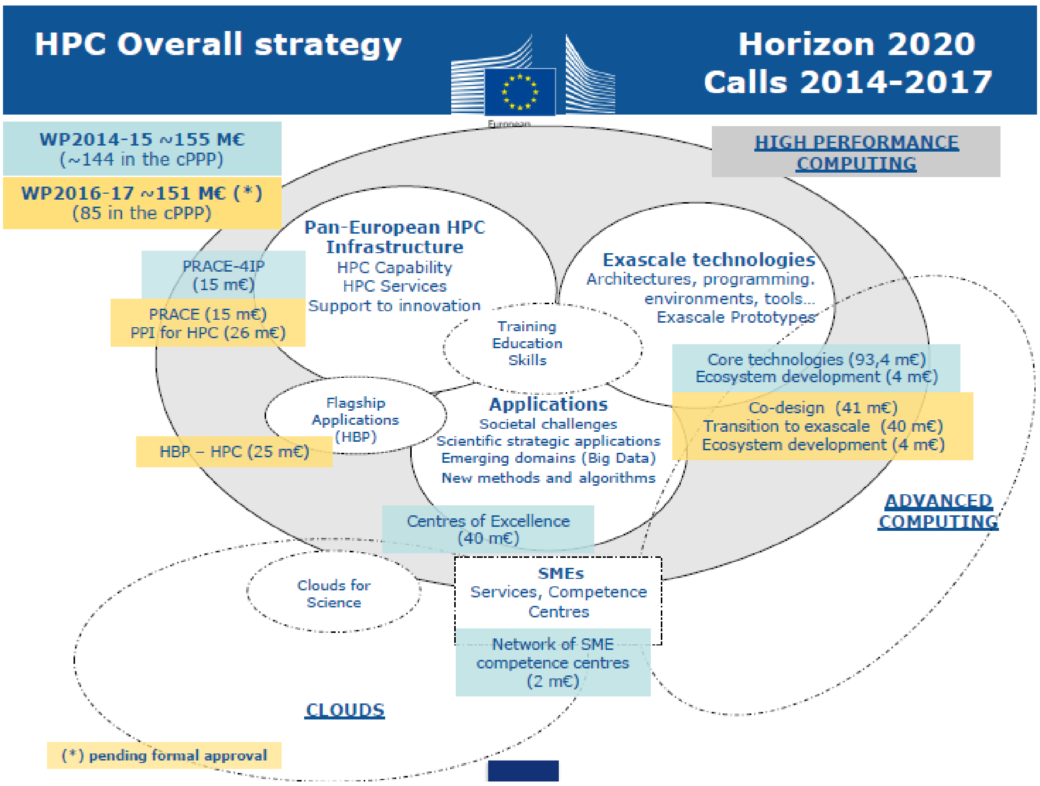

An indicative timeline for the contractual Public-Private Partnerships (cPPP) in HPC, under which the EC will bring together technology providers and users in the European Technology Platform for HPC (ETP4HPC) for (a) developing the next generation of HPC technologies, applications and systems towards exascale and (b) achieving excellence in HPC applications, is shown below.

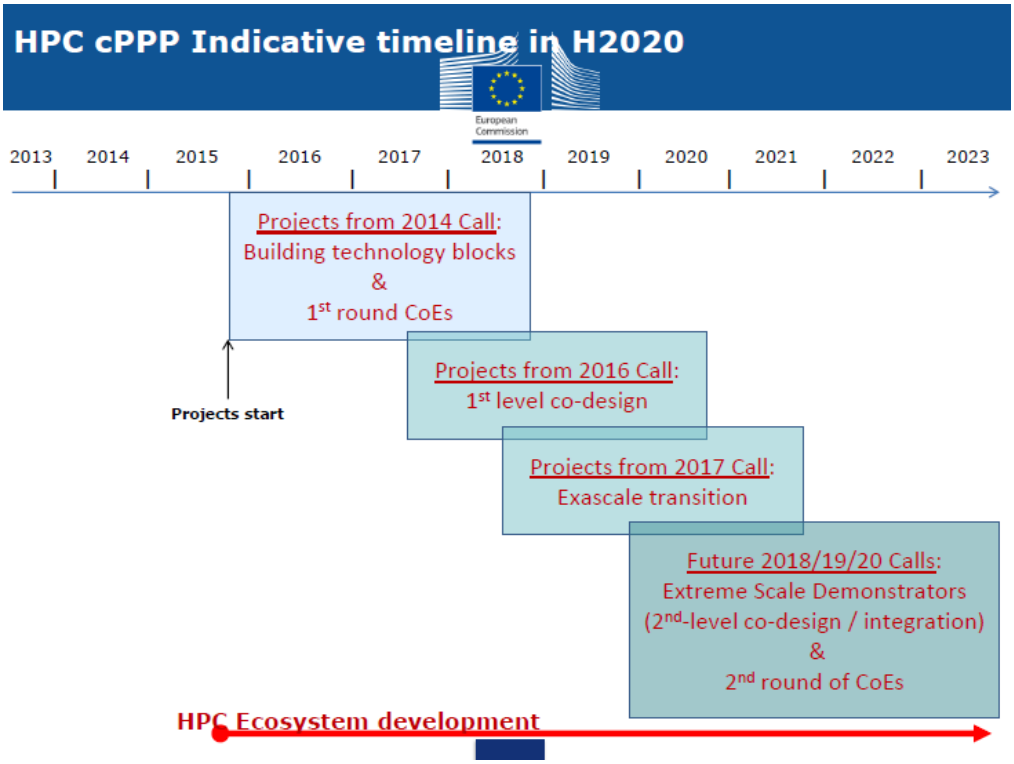

2. Sergi Girona, representing PRACE, presented the EXDCI (Extreme Data and Computing Initiative) project, which was awarded by the EC. It is targeted to the coordination of the development and implementation of a common strategy for the European HPC ecosystem, including working towards a synchronised European HPC community. An initial workshop with all current HPC projects will be organised in the early phase of EXDCI.

5.5 Round Table on what's next towards exascale

Sergi Girona of PRACE has acted as the moderator for the round table discussion on what's next towards exascale. The panelists were:

Panagiotis Tsarchopoulos (European Commission – DG Connect)
Jean-François Lavignon (ETP4HPC)
Philippe Ricoux (EESI – TOTAL)
François Bodin (EXDCI - IRISA)
Eric Duceau (Airbus)
Jean-Yves Berthou (ANR)

Questions posed to the panel:

1. What are you planning to transform the future into present?
2. What are the damages/costs if the plan to implement exascale does not succeed?

3. What are the contingencies plan?
4. How can we link the EU R&D programs?

Opinions of the panel:

1. Effectiveness, both in the process and the impacts. The panellists emphasized that the excascale capabilities must serve the industry purposes and its implementation must be pragmatic.

   To achieve these two objectives, it is important to use and improve the monitoring of the funded project and the measure of the HPC ROI. The risks are to disperse the investments and to create useless high-level computers.

   Effectiveness is the key to convince the funders and users and successfully gain the HPC revenues. If the EU fails in that part, the HPC added value will be developed abroad.

2. Anticipation. The timeframe of both policy makers and industry leaders needs anticipation to adapt. This is of most importance to remain competitive in a global market, as well as to design efficient policies. Anticipation has also its on role on collaboration, which needs trust, and trust requires time.

   To answer this, it is of most importance to have a common roadmap build on long-term perspective. If we fail to anticipate, we will lose crucial time on our implementation of efficient exascale, taking the risk to stand beyond the international competition.

3. Collaboration. Alone, nobody can reach the exascale challenges. Collaboration is multi-dimensional, involving: industry leaders, scientists and experts from different domains, dialog with other communities (Big Data, IT sector) and with the policy-makers as well as worldwide cooperation.

   Collaboration goes through the animation of an HPC ecosystem, sharing a general vision, in a proactive, open and cooperative way. Collaboration needs the design and use of common tools (ROI measure, SRA, Centres of Excellence, ...). Neutral bodies such as PRACE or ETP4HPC are also important.

4. Adaptation. The implementation of the exascale roadmap must be revised and adapted on a regular basis and the technological milestones monitored. The industry leaders need to adapt to remain competitive in a global market. Collaboration between multiple science domains, engineers and different communities needs the adaptation of our work process and, last but not least, going for breakthrough innovation requires that we put in place transition management.

6. Conclusion

The EESI2 project coordinator, Philippe Ricoux of TOTAL, summarised the final conference, of course by thanking all participants, but also summarizing the scientific, strategic and funding recommendations done by EESI2. Philippe was particularly satisfied with the broad support for the recommendations, not only by scientists but also by industry. The recommendations are concrete, many of them are disruptive. Next steps must be implementation together with industry to reach efficient exascale applications.

EESI Scientific Recommendations

- **Recommendations of the Tools & Programming Models pillar:** (Tools&PM)
  - High productivity programming models for Extreme Computing
  - Holistic approach for extreme heterogeneity management of Exascale supercomputers
  - Software Engineering Methods for High-Performance Computing
  - Holistic approach to resilience
- **Recommendations of the Ultra Scalable Algorithms pillar:** (Ultra Scalable Algorithms)
  - Algorithms for Communication and Data-Movement Avoidance
  - Parallel-in-Time: a fundamental step forward in Exascale Simulations (disruptive approach)
- **Recommendations of Data Centric Approaches pillar:** (Data Centric Approaches)
  Vision "Software for Data Centric Approaches to Extreme Computing"
  - Towards flexible & efficient Exascale software couplers (direct or not, exchange of massive data)
  - In Situ Extreme Data Processing and better science through I/O avoidance in High-Performance Computing systems
  - Uncertainties Quantifications tools evolution for a for better exploitation of Exascale capacities
  - Declarative processing frameworks for big data analytics, extreme data fusion e.g. identification of turbulent flow features from massively parallel Exaflops and Exabytes simulations

The amount of funding required to develop the entire software ecosystem towards exascale computing is large. For real competitiveness of European Industries and Research, critical mass could lead to:

- Specific research and Development projects (Recommendations): 20 projects, 3 M€ per year for each during 10 years, gives 600 M€ over 10 years;
- Centre of Excellence program: 5-7 CoE, 10 M€ per year for each, gives 500-700 M€ over 10 years;
- Exascale Infrastructure platforms: 500-700 M€ for 2 to 3 platforms, over 10 years;
- Technological transfer through a European Exascale Software Centre: 20 M€ per year, gives 200 M€ over 10 years.

This gives a cumulative estimated budget of around 2 billion €.

EESI Strategic & Funding Recommendations

Towards Exascale EU funding Budget must concern of course Infrastructure but also and largely the Software and Applications as pushed by EESI

It is a fact: Europe is in late vs USA (CoE, Examaths, In situ Data Processing, Resilience, …)

Exascale is Disruptive, So Europe must be disruptive in its strategy

Key R&D programs must be *specified to tackle Exascale issues* as described in the 3 EESI pillars, not mixed with classic HPC, and should be *funded at once* without waiting for 2017. Europe already lost 3 years!

Research critical mass must be assigned to each key program and each CoE during 10 years

Europe must develop a THINK TANK for Exascale (EESI … or equivalent)

For real competitiveness of European Industries and Research , Critical mass could lead to:

- **Specific Research and Development projects (Recommendations)**: 20 projects; 3M€/year each during 10 years → **600 M€ over 10 years**
- **Center of Excellence program**: 5-7 CoE , 10M€/year each → **500 – 700 M€ over 10 Years**
- **Exascale Infrastructure platforms: 500 – 700 M€ for 2 to 3 platforms, over10 years**
- **Technological transfer** through a **European Exascale Software Center**: 20 M€/year → **200 M€ over 10 years**

Cumulative Global estimated budget : ≈ 2 Billions € over the next 10 years (Confirming EESI1)

EESI2-Final-Conf-Dublin-28May2015-PhR 12 **Much Higher than EC plans!**

7. Annex – EESI2 Conference Participants

1	Andre	Jean-Claude	Jca Consultance & Analyse
2	Asch	Mark	Anr
3	Badia	Rosa M	Barcelona Supercomputing Center
4	Bala	Piotr	ICM, University of Warsaw
5	Bauer	Peter	
6	Bergeaud	Vincent	Cea
7	Berthou	Jean-Yves	Anr
8	Bidot	Thierry	Neovia Innovation
9	Bode	Mathis	RWTH Aachen University
10	Bodin	François	
11	Bogaerts	Serge	Cenaero
12	Bolten	Matthias	University Of Wuppertal
13	Botten	Lindsay	National Computational Infrastructure
14	Brown	Nick	Epcc
15	Browne	Michael	ICHEC - Irish Centre For High-End Computing
16	Chami	Fatima	Durham University
17	Demetrescu	Camil	Sapienza University Of Rome
18	Demichel	Patrick	Hp
19	Desplat	Jc	Ichec
20	Dewar	Mike	NAG Ltd
21	Duceau	Eric	Airbus Group Innovations
22	Duchaine	Florent	Cerfacs
23	Duff	Iain	STFC - Rutherford Appleton Laboratory
24	Epicoco	Italo	University Of Salento
25	Erbacci	Giovanni	Cineca
26	Falter	Hugo	ParTec Cluster Competence Center GmbH
27	Filippone	Salvatore	University Rome Tor Vergata
28	Foujols	Marie-Alice	Ipsl
29	Furlong	Michelle	Institute For Computational Cosmology
30	Giraud	Luc	Inria
31	Girona	Sergi	Prace Aisbl
32	Gonnord	Jean	Cea
33	Grigori	Laura	Inria
34	Große-Wöhrmann	Bärbel	High Performance Computing Center (HLRS)
35	Grothey	Andreas	University Of Edinburgh
36	Hanke	Michael	KTH Royal Institute Of Technology
37	Ho	Yoon	Rolls-Royce PLC
38	Kågström	Bo	Umeå University
39	Kamps	Fieke	SURFsara
40	Kennedy	Alison	EPCC, University Of Edinburgh And PRACE Aisbl

41	Koehler	Axel	Nvidia
42	Lambert	Catherine	Cerfacs
43	Lavignon	Jean-Francois	Bull
44	Lefevre	Corentin	Neovia Innovation
45	Ltaief	Hatem	Kaust
46	Ludwig	Thomas	German Climate Computing Center
47	Markidis	Stefano	Kth
48	Meerbergen	Karl	KU Leuven
49	Miah	Wadud	
50	Michielse	Peter	SURFsara
51	Mohr	Bernd	Jülich Supercomputing Centre
52	Nielsen	Keld Lund	ENI SpA
53	Nomine	Jean-Philippe	Cea
54	Oorsprong	Marjolein	PRACE Aisbl
55	Ostasz	Marcin	Barcelona Supercomputing Centre
56	Ott	Michael	Leibniz Supercomputing Centre
57	Ottaviani	Maurizio	Cea Irfm
58	Papka	Michael	Argonne National Laboratory
59	Parsons	Mark	EPCC, The University Of Edinburgh
60	Peng	Ivy	
61	Perseil	Isabelle	Inserm
62	Petiton	Serge	U. Lille 1, And Maison De La Simulation/CNRS
63	Poghosyan	Gevorg	Karlsruhe Institute Of Technology
64	Requena	Stephane	GENCI / PRACE
65	Ribes	Alejandro	Edf - R&D Sinetics
66	Richardson	Harvey	Cray UK Ltd
67	Ricoux	Philippe	Total Sa
68	Saint Georges	Marc	Adaptive Computive
69	Sawley	Marie-Christine	Intel IPAG
70	Segers	Philippe	Genci
71	Stewart	Dean	Rogue Wave Software Uk Ltd
72	Tsarchopoulos	Panagiotis	European Commission
73	Valero	Mateo	Barcelona Supercomputing Center
74	Vanroose	Wim	University Of Antwerpen
75	Verhoeven	Dionne	SURFsara
76	Walsh	John	Rogue Wave Software
77	Wolf	Laura	Argonne National Laboratory
78	Zyulkyarov	Ferad	Barcelona Supercomputing Center