European Exascale Software Initiative

Deliverable D3.1

Progress Report of WP3

Applications working groups

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

Due date of deliverable: 30 July 2013

Actual submission date: 30 July 2013

Publication date:

Start date of project: 1st September 2012 Duration: 30 months

Name of lead contractor for this deliverable: GENCI

Authors: Stéphane Requena (GENCI), Norbert Kroll (DLR), Ange Caruso (EDF), Giovanni Aloisio (Univ. Salento), Jean Claude André (Jca Consultance), Godehard Sutmann (FZJ), M. Orozco (BSC), C. Laughton (Univ. Nottigham), U. Rude (Univ Erlangen).

Name of reviewers for this deliverable:

Abstract: This document describes reports the outcome of the first year of activity of EESI-2 WP3 dedicated to Applications

Revision 1.0

Project co-funded by the European Commission within the Seventh Framework Programme (FP7/2007-2013)		
Dissemination Level		
PU	Public	
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)	X

Glossary

Abbreviation / acronym	Description
API	Application Programming Interface
CAD	Computer Assisted Design
CERN	Centre Européen de Recherche Nucléaire
CFD	Computational Fluids Dynamics
CI	Combustion Instabilities
CMIP5	Coupled Model Intercomparison Project Phase 5
CTR	Center for Turbulence Research (Stanford)
DEEP	Dynamical Exascale Entry Platform
DNS	Direct Numerical Simulation
DSL	Domain Specific Langage
EC	European Commission
ENES	European Network for Earth System Modeling
EESI	European Exascale Software Initiative (Europe)
EPOS	European Plate Observation System
ERC	European Research Council grant
ESA	European Space Agency
ESM	Earth System Models
FAU	Friedrich-Alexander Universität Erlangen
FDA	(US) Federal Drugs Administration
FET	Future and Emerging Technologies, an EU pathfinder programme in information technologies
GPU	Graphical Processing Units
HPC	High Performance Computing
IDC	International Data Corporation
IT	Information Technology
MOOC	Massive Open Online Courses
NAFEMS	National Agency for Finite Element Methods and Standards, an independent, not-for-profit membership association dedicated to FEA and CFD

NASA	(US) National Aeronautics and Space Administration
ISV	Independent Software Vendors
LES	Large Eddy Simulation
PCM	Phase Change Memory
PETSc	The Portable Extensive Scientific toolkit, a numerical library
PRACE	Partnership for Advanced Computing in Europe, the European HPC research infrastructure
SPH	Smooth Particle Hydrodynamics
UQ	Uncertainty Quantification
WCES	EESI-2 WP3 Weather, Climatology and Solid Earth Sciences working group
WG	Working Group (in EESI)
WP3	EESI2 workpackage on Applications
WP4	EESI2 workpackage on Enabling technologies
WP5	EESI2 workpackage on Cross cutting issues

1. Introduction

The objective of EESI2 work package 3 (WP3) “Application working groups“ is to investigate on scientific and industrial application drivers for Peta and Exascale computing. It is organized into five representative working groups and a last one dedicated to WP3 management. Following the activity started during the previous EESI-1 project this work package aims to pave the path by:

- Investigating on key application breakthroughs, quantifying their societal, environmental and economical impacts and performing a gap analysis between current situation and Exascale targets;
- Evaluating the R&D activity performed by scientific and industrial communities, especially in applications redesigning and development of multiscale/multiphysics frameworks;
- Fostering the structuration of scientific communities at the European level;
- Integrating within WP tasks, cartography update through continuous use of network.

WP3 is organized around 5 applications working groups, the first four of them are in continuity with the existing EESI WP3:

- WG 3.1: Industrial and engineering applications
- WG 3.2: Weather, Climatology and Solid Earth Sciences
- WG 3.3: Fundamental Sciences
- WG 3.4: Life Sciences and Health

WP3 has been completed by 2 new additional working groups:

- WG 3.5: Disruptive Technologies
- WG 3.6: Coordination of WP3

Activities of WP3 are tightly linked with activities from WP4 “Enabling technologies Work Groups” and the newly created WP5 “Cross cutting issues Work Groups “ since WP3 and WP5 will provide inputs regarding scientific communities Exascale needs and WP4 will provide answers about what will be the technology and the tools available at this period.

During the first period, each of the five first working groups recruited around an average of 10 experts in the respective areas and started internal face to face meetings and regular telcos.

End of May 2013, all work package leaders and working group chairs, vive-chairs and topic experts met for a 2-day internal workshop to present the results of the individual working groups, to discuss update of their respective roadmaps, gap analysis and were involved into a first set of 12 project recommendations from EESI2.

This document summarizes the most important results, findings, conclusions and recommendations of these five WP3 working group reports during the first period.

2. Executive summary and main recommendations

During the first period, each of the five first working groups recruited around an average of 10 experts in the respective areas and started internal face to face meetings and regular telcos.
For the whole WP3 a total of 43 experts were enrolled during the first period by the 4 first WGs, these experts are coming from academia as well as industry with 7 companies represented. It was also important to ensure a wide geographical representation of the experts across 8 European countries as well as Russia and USA.

End of May 2013, all work package leaders and working group chairs, vive-chairs and topic experts met for a 2-day internal workshop to present the results of the individual working groups, to discuss update of their respective roadmaps, gap analysis and were involved into a first set of 12 project recommendations from EESI2.

This document summarizes the most important results, findings, conclusions and recommendations of these five WP3 working group reports during the first period.

In term of gap analysis and major breakthroughs since 2011 it's important to notice:

- For the industrial applications
  - These industries are more or less following the roadmap announced in EESI-1 with perhaps a little delay of 2 years. Companies like Airbus, EDF and TOTAL made recently major investments into large petascale HPC systems making them one of the biggest users of HPC in their respective domains.
  - The European industrial roadmaps have been challenged against the one from their US competitors and the results are strongly coherent.
  - In automotive there is a strong acceleration of the use of advance HPC for crash analysis with the enforcement of the crash regulation (EuroNCAP6 in 2015 or in 2020 with the EU recommendation imposing to halve the number of overall deaths) or the rise of composite in order to replace steel.
  - In combustion major results have been possible in US and Europe with the deployment of petascale systems by NSF/DoE and PRACE. The first jet noise simulation on one million core have been performed by a team of CTR (Stanford) on the Sequoia DoE system and German teams from Univ. Aachen have launched petascale simulations on PRACE systems.
    Some aerospace companies are now setting up ambitious roadmap for developing next generation of gas turbines for 2020. This will require to couple LES simulations et all the different stages of the turbine together with the need to develop both simulation codes as well as flexible coupler and optimised coupling techniques.
- For the Weather/Climate and Solid earth Sciences
  - The climate community is now preparing all the workflows and the simulations codes toward the next IPCC campaign planned to start in 2016/2017. As already stated into the EESI-1 roadmap this will need to work on data management issues which already became crucial as well as working in improving scalability of the applications, developing

dynamical cores, coupling multiple codes, assessing uncertainties of the models, …

- In Fundamental Sciences
    - In cosmology a lot of massive simulations especially in the field of dark matter or collisions between galaxies have been performed by US, Japanese or European research teams during the period. Such simulations are posing crucial issues for handling very large amount of data to be treated on the fly during the computation and made available during years to worldwide communities.
    - Two European fusion codes, GYSELA and PEPC have been scaled out to the full JUQUEEN machine at Juelich (on halt a billion core using 1,835,008 threads) showing massive scalability and such teams are engaged into one of the G8 funded project NuFuSe.
- In Life Sciences
    - Alongside the announcement of the Human brain Project FET by the EC, the US administration announced the support a similar imitative called BRAIN with a budget of 100M$ for 2014-2015.
    - On June 2013 a team leaded by University of Illinois physics professor Klaus Schulten and postdoctoral researcher Juan Perilla recently used molecular simulations on the NSF Blue Waters supercomputer to determine the chemical structure of the HIV capsid using NAMD on an unprecedented 64 million-atom simulation.
- The European Commission announced in early 2013 the launch of 2 major Flagships : Human Brain Project and Graphene with a funding of more than 1B€ during 10 years for each project. These two initiatives will lead to the creation in Europe of strong structured communities highly connected with existing research infrastructures like PRACE or GEANT.
- The full deployment of PRACE, the European HPC infrastructure with now 6 petascale systems in France, Germany, Italy and Spain offering a cumulated performance of more than 15 PFlops and high value services like training (through 6 dedicated PRACE Traning Centers), user support and code enabling and co development for Open Source applications.
  In January 2012, the access to PRACE has been extended to industrial users through for Open Research activities and until now 12 companies (large companies as well as SMEs) already benefited from access to PRACE systems and services.
- At the national level as well as the PRACE level some specific initiatives have been launched in order to help SMEs to discover, assess and use HPC in order to increase their competitiveness.

The experts also identified disruptive technologies in the following fields :

- Simulation and Optimization with uncertainties
- Simulation and Optimization in Complex Networks
- Data Driven Simulations
- Bridging scales by ultra-large simulations
- Advanced Adaptive Resolution
- Learning-on-the-fly

- Multiscale: hybrid and beyond
- In Silico Design of New Materials

The workpackage have been deeply involved into the EESI-2 recommendations especially regarding code couplers, training, Big Data, mini applications, ultra scalable solvers as well as UQ & Optimization methods and tools.
Additional recommendations from WP3 are the following:

- To reinforce the training program of PRACE across Europe and the support of HPC and numerical simulation trainings from education (under graduate programs) to permanent training. MOOC tools need to be developed in complement of physical trainings in order to train more people from new disciplines or coming from remote locations.
- To fund co-educational programs on the master level which provides a broad knowledge of hardware technology and system- and application software not only for computer scientists but also for students in scientific computing.
- Foster educational programs for training of postdoctoral and senior scientists, who have to follow the trends and developments of hardware architectures, software and optimization and tuning of codes.
- With the rise of Open Innovation between academia and industry, it become mandatory to industrialize the different codes developed by academia in order to make them usable by industry with higher TRL.
  For the European Commission as well as the others funding agencies there is a need to invest on structures/teams which are able to provide software engineering methodologies for developing new standard components (numerical libraries, i/o libraries and tools, data managements tools, pre/post processing packages, …) that could be used widely in a long tem period by various communities. Such example of shared components used daily by multiples communities is for example the PETSc numerical library funded by DoE in US with a strong team for developing, testing and supporting the tool. Developed at the European level in order to reach the needed critical mass and to gather a distributed expertise, such common components will avoid to European scientists and industries to re-invent the wheel every time and could allow Open Source software to cohabit with ISV commercial software by using the same underlying framework.
  As additional recommendations:
  - Define standards for sharing and (re-)using software including commercial parts.
  - Create code repositories, which offer numerical kernels, mini-apps, modules and numerical libraries, optimised for different architectures and provide long-term support and visibility.

    Again in this domain such an infrastructure like PRACE could be a good tested for hosting such service.

3. Scientific contest and tasks

3.1 Industrial and engineering applications

Industry is in the midst of a new, 21st century industrial revolution driven by the application of computer technology to industrial and business problems. HPC already plays a key role in designing and improving many industrial products — including automobiles, airplanes, pharmaceutical drugs, microprocessors, computers, implantable medical devices, golf clubs, and household appliances — as well as industrial-business processes (e.g., finding and extracting oil and gas, manufacturing consumer products, modelling complex financial scenarios and investment instruments, planning store inventories for large retail chains, creating animated films, and forecasting the weather).

HPC users typically pursue these activities with virtual prototyping and large-scale data modelling – that is using computers to create digital models of products or processes and then evaluating and improving the design of the products or processes by manipulating these computer models. Given their broad and expanding range of high-value economic activities, HPC users are increasingly crucial for industrial and business innovation, productivity, and competitiveness.

In a communication issued in February 2012 by the European Commission (EC) and called "High-Performance Computing: Europe's place in a Global Race » the EC emphazed the crucial role of HPC in allowing European industry to develop innovative products and services. The Council asked for a further development of the European High Performance Computing Infrastructure and a pooling of national investments in HPC in order to strengthen the position of European industry and academia in the use, development and manufacturing of advanced computing products, services and technologies. This is the high-level objective driving a renewed European HPC strategy.

In two studies conducted by IDC (International Data Corporation) and called 'A Strategic Agenda for European Leadership in Supercomputing: HPC 2020' and 'Financing a Software Infrastructure for Highly Parallelised Codes' almost 97% of the industrial companies that employ HPC consider it indispensable for their ability to innovate, compete, and survive.

Since 2012, PRACE the European HPC Infrastructure opened its HPC systems and services (training, code enabling, …) to industrial users through the Open R&D call for proposal. Companies are now eligible for performing open research alone or paired with academia. In less than 1 year this programme attracted 12 companies (larges groups as well as SMEs) for a cumulated allocation of more than 100 million cpu hours.

In parallel, some companies are also using their own HPC facilities and Europe with companies like TOTAL, EDF, Airbus, CGG-Veritas, Shell, to name a few are very well positioned in the level of HPC facilities owned by companies.

This is a very good signal but on the other side an important part of the European industrial competitiveness, the Small and Medium Sized companies (SMEs) are not using HPC and even advanced numerical simulation. In order to raise awareness and foster innovation and competitiveness of SMEs through HPC, some national initiatives in France, Germany, Italy, Netherlands and UK started few years ago and more recently PRACE launched a pilot of enabling services for SMEs (called SHAPE) and EC granted a project called Fortissmimo for setting up a pan European Cloud infrastructure offering commercial HPC services for "evangelised" SMEs.

3.2 Weather, Climatology and Solid Earth Sciences

Earth Sciences in general concern the sciences of the planet Earth. It encompasses a wide range of disciplines from the study of the atmosphere, the oceans, the biosphere to issues related to the solid part of the planet. Earth Sciences address many important societal issues from weather prediction to air quality, ocean prediction, climate change to natural hazards such as seismic and volcanic hazards, for which the development and the use of high-performance computing (HPC) play a crucial role. Among all these issues WCES concentrates on two main domains, on the one hand weather and climate, which share some similarities, and on the other hand solid earth. These do not include all the modelling aspects of Earth Sciences but include the two domains mainly concerned by exascale computing issues.

The two main objectives are then (i) to investigate and identify the main challenges connected with the Weather, Climatology and solid Earth Sciences (WCES) domains in the exascale timeframe, and (ii) to contribute to the establishment and up-dating of a coherent and integrated vision/roadmap towards Exascale focusing on these main challenges. More specifically the WCES working group is addressing such main issues as:

- Scientific data management: visualization, parallel files systems, parallelization strategies and frameworks for data analysis, programming models for big data, data compression, data mining, data quality, data provenance, metadata management, etc.
- Earth Sciences Modelling: finer resolution, development, evaluation of parameterization well adapted to those fine resolution, fast run for lower-resolution models, uncertainty quantification, quality control, new approaches (e.g., new grids and new dynamical cores), etc.
- Cross-cutting issues (multi-physics coupling software for 10^5-10^6 cores, resilience, and performance of individual modules and of complete coupled systems, parallel I/O) are also presented and discussed throughout the report.

The WCES working group will also provide a deeper understanding (wr.t. the EESI1 roadmap) of the WCES impacts on societal, environmental and economic aspects.

3.3 Fundamental Sciences

During the last years some effort and development has been seen in hardware- and software-development towards the exascale era. Fundamental sciences, comprising nuclear physics, laser-plasma physics, nuclear fusion, quantum chemistry, soft matter physics, materials science and astrophysics/cosmology, strongly rely on fast algorithms and efficient software implementations which are able to keep track of developments in hardware technology. However, there is no clear opinion, whether the transition to exascale can be achieved by a smooth transition from current implementations or whether a revolutionary step is necessary to allow for e.g. new programming models, new communication schemes, new memory management, new algorithms including fault tolerance and energy efficiency in order to make the transition from peta- to exascale.

Europe has a very strong position in the global scientific community and is leading in several fields of astrophysics/cosmology, nuclear/hadron physics, fusion research and materials sciences. European software development of community codes is very advanced in materials science, quantum chemistry, astrophysics, nuclear/hadron physics and plasma/fusion physics. The newer field of multiscale simulations will have a strong impact on future scientific leadership and competitiveness. With appropriate support, Europe would be well positioned to develop pioneering multiscale software frameworks.

3.4 Life Sciences and Health

This working group is focused on HPC applications for Life Sciences and Health. The Life Science community is very diverse and there is a large unbalance between the large size of experimentalists/biologists (that strongly depends on computational results) and the small size of computational biologists (that depend heavily on HPC resources). For this reason the work of computational biologists has a “multiplicative” effect on Life Sciences. The goal of computational biology and bioinformatics is to understand the mechanisms of living systems. With the recent advances in this area (e.g. next generation of DNA sequencing instruments) the generated data is becoming larger and more complex. In contrast to other communities there are no universal computer packages and software evolves very fast to adapt to new instruments. The problems faced by scientists working in molecular simulations and genomics are also very different, as are the computer algorithms used. The importance of having fast and flexible access to very large computational resources is crucial in the many fields of Life Sciences and the lack of suitable computers can block entire projects with important consequences for science and the society.

3.5 Disruptive Technologies

The notion of “disruptive technology” implies the meaning of “breaking traditions”. This might open the field for completely new approaches to modelling techniques and ways of thinking but, at the same time, implies a high risk of failure on the long run. In the context of application domains discussed in the context of exascale computing it is a challenge by itself to classify disruptiveness, since (i) exascale is a concept of the future for whose concepts of possible technologies are still under discussion and (ii) disruptive implies breaking with the past. However “applications” refer to actual scenarios, not just pure speculations. Also disruptive technologies must show a concrete potential to be successful in the coming Exascale era. Therefore, the selection of topics that are discussed in the present context, is based on fields where some development and progress has already been achieved, but which is not yet in a stage of maturity and which both bares a risk to be successful on the long run and contains a high potential for a break-through if methodological and algorithmic advances are further progressed.

Topics which are discussed within “disruptive technologies in applications” exhibit an inherently interdisciplinary and multi-disciplinary character. They often include a variety of length- and time-scales and a diversity of different methods, e.g. stochastics vs deterministics, lattice vs mesh-less methods, optimization, dynamics vs sampling methods including multiscale methodologies. Naturally, the selection of topics is limited and cannot claim to be complete.. Additionally it is hardly possible to define a natural scheme of priorities. Nevertheless, the topics selected for the present study seem to fulfill the criteria given above, i.e. baring risk and a high potential at the same time for a very large class of applications.

To give an example, multiscale techniques are widely discussed in literature and are sometimes considered as versatile method to characterise complex systems over various time- and length-scales. However, multiscale often implies multi-numerics, i.e. coupling different methods which have their own stability constraints and level of accuracy and a concurrent coupling of methods might not guarantee a possible error control. Promising schemes have been proposed and schemes like adaptive multi-resolution and learn-on-the-fly methods are very promising to advance efficient simulations on high performance architectures.

Exascale bares a chance to solve various types of problems related to the complexity of the problems. For example, very complex networks exist in nature and technological applications, which form a type of self-organized, synthesized or genetically determined structure. The function of such systems is difficult to recast in a straightforward and deterministic way. To understand the function of complex networks or to model and optimize functions within networks contains the potential to help solving socio-economic or environmental problems, like climate change, neurological diseases, efficient energy-production and -usage or market crashes.

4. List of experts

During the first period all the five WP3 working groups working chairs and vice chairs worked heavily in engaging an average of 8 scientific experts into their respective WGs. For the 4 first WGs this list take as input the most actives experts previously participating into EESI-1 and additional experts from European and international projects and users communities:

- For the WG3.1 (Industrial and engineering applications) an initial set of 8 experts from EESI-1 WG was completed by two new experts following discussions with NAFEMS in order to increase the representation of ISV (Independent Software Vendors). This last point is very important as ISV are a key partners for industrial companies and by consequence a key domain for assessing needs and roadmap toward Exascale.
- For WG3.2 (Weather, Climatology and Solid Earth Sciences) the initial list has been extended with the ENES HPC Task force, an ad hoc body created by the ENES (European Network for Earth-System modelling) to specifically address climate modelling and data HPC issues and European initiatives/projects like IS-ENES, EPOS, ESA-CCI. The list of experts was progressively extended, by adding new experts any time a new topic was addressed requesting new inputs.
- For the WG3.5 until now no fixed working group on “Disruptive Technologies in Application Domains” has been formed. Due to the complex nature of the topic it was necessary to first develop a more concrete definition of the themes and topics. Discussions in the community were initiated on an informal basis and additionally; a survey among various communities was performed. It resulted in feedback with individual views and community demands. This has helped to clarify which research topics may need to be developed apart from traditional research lines in order to progress the field or to allow for possibilities enabling a more broad or detailed description, including the view with respect to exaflops computing.

For WG3.4 (Life Sciences and Health) the list of experts have been formed by using a reduce set of 8 experts from previous EESI-1 working group.

For the whole WP3 a total of 43 experts were enrolled during the first period by the 4 first WGs, these experts are coming from academia as well as industry with 7 companies represented. It was also important to ensure a wide geographical representation of the experts across 8 European countries as well as Russia and USA.

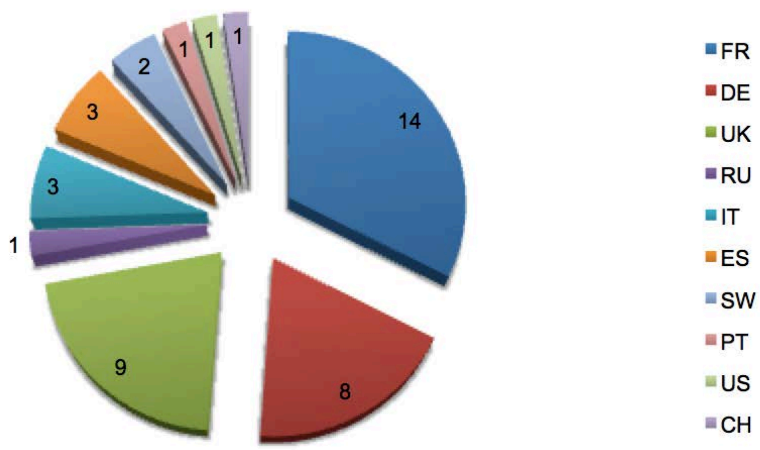

Figure 1 - Repartition per country of the 43 WP3 experts

More experts will be incorporated, especially for WG3.5 on disruptive applications and on-demand for specific issues for the others WGs.

After a first phase of enrolment of the experts, the first four first WGs started on assessing gap analysis between their previous EESI-1 recommendations and roadmaps, and focusing on new fields of research as well as identifying disruptive technologies in relation with WG3.5.
For the WG3.5 due to the complex nature of the topic it was necessary to first develop a more concrete definition of the themes and topics and perform an initial survey among scientific communities.
The working groups activity was performed though face-to-face meetings and regular teleconferences or WebEx sessions.
As notables facts of this activity the following elements can me noticed:

- For WG3.1: With the newly presence of ISV a specific activity started by drafting a questionnaire to be sent during summer 2013 to the members of NAFEMS in order to gather needs, challenges and recommendations of Independent Software Vendors. This activity will be monitored by Lee Margetts and Vladmir Belsky and outcome is expected for fall 2013
- For the WG3.2: ESM challenges were also discussed during two earlier workshops organized by the IS-ENES HPC Task Force, in Lecce, Italy, in Dec.2011, and in Toulouse, France, in Jan.2013.
  The WG benefits from strong links with ENES HPC Task Force, which is working through regular TelCo's (5 in 2012, 5 to date in 2013).
- For the WG3.4 an international conference on Exascale called "Towards in Silico Humans. A Challenge for Exascale Computing Area", to be held on September, 18th, 19th and 20th, 2013, in Barcelona, Spain will allow in-depth discussions with a wide panel of experts (including the ones from WG3.4) and produce an updated roadmap and gap analysis. The agenda of this 3-day meeting is the following :

SESSION	DEBATE	DAY
The Future of Supercomputing for Life Sciences		18-set-13
9:15-9:30	Wellcome	
9:30-11:00	European resources and Future supercomputing	
11:00-11:30	Coffee break	
12:00-13:00	Investing in Disruptive Technology	
13:00-15:00	Lunch Break	
15:00-17:30	European Exascale Software Initiative *(incl coffee break)*	
Multi-scale simulation of living organisms		19-set-13
9:00-11:00	Molecular Simulation	
11:00-11:30	Coffe Break	
11:30-13:30	Life systems simulation	
13:30-15:00	Lunch Break	
15:00-18:30	Open Session *(incl coffee break)*	
Future Computing for Pharmaceutical Industry		20-set-13
9:30-10:30	Computing for Genomics	
10:30-11:00	Coffe Break	
11:00-12:00	Computing for Systems Biology	
12:00-13:00	Computing for Drug Discovery	

All the WP3 chair and vice chairs also participated on the first internal EESI2 conference on May 28-29 in Tremblay, this was an opportunity to present some first results of their activity, share their views about cross cutting issues and participate on the collective brainstorming on the 12 EESI2 first recommendations.

For the whole WP3 regular teleconferences with chairs and vice chairs has been organised and all the materials have been uploaded on the EESI2 internal wiki.

5. Update on key challenges

5.1 Challenges in industrial and engineering applications

5.1.1 Challenges on Aeronautics and Aerospace

The impact of computer simulation in aircraft design has been significant and it continues to grow with the expectations of regulators (Acare and Flightpath) and the need to provide optimized designs and reduced development risks and costs.
The latest ACARE (Advisory Council for Aeronautics Research in Europe) (and the Flightpath 2050 Europe's Vision for Aviation) reports are expecting for 2050 a CO_2 reduction by 75%, a NOx reduction by 90%, by 65% the reduction in perceived aircraft noise (regarding a 2000 typical aircraft), by 80% the accident rate, while air companies are expecting reliable and less kerosene hungry planes with a 3x increase of the traffic and 99% of flights within 15' of schedule!

To meet the challenges of future aircraft transportation (Greening the Aircraft), it is indispensible to be able to flight-test a virtual aircraft with all its multi-disciplinary interactions in a computer environment and to compile all of the data required for development and certification with guaranteed accuracy in a reduced time frame.

In order to maintain European global leadership and serve such societal needs the aircraft industry will need to increase product performance, perform step changes in aircraft technology and foster new design principles.
As already stated into previous EESI-1 report, this digital aircraft vision includes the following major simulation and optimisation challenges:

- Accurate physical modelling through the complete flight envelope: Improved physical modelling for highly separated flows (turbulence / transition) by using in production massive LES simulations on industrial cases (using complex geometries with high Reynolds numbers) as well as accurate integrated prediction tools. Due to the requirements of high resolution models in time and space, LES simulation for a full aircraft will most likely not be feasible with an Exascale computer, so Exascale will be a mandatory stepping stone towards more computing capacity.
- In flight aircraft real time simulation: the full real time simulation of a manoeuvring aircraft will require the efficient coupling of all relevant disciplines using high-fidelity methods (aerodynamics, structural mechanics, aero-elastics, flight mechanics, aero-acoustics, engines, ...). This will require the availability of efficient couplers as well as scalable individual codes. Exascale systems used in capacity mode will allow simulating several manoeuvres simultaneously (flying the aircraft by equations).
- Aerodynamic and aeroelastic data production: This requires methods to calculate in the shortest possible time frame aerodynamic loads and derivatives of the aircraft for every imaginable flight situation. The data are used for the structural layout of the aircraft, the flight simulator database and the development of the flight control system. The required input is currently obtained mainly from experiments or from simple empirical methods with strong issues between surrogates and reduced order models Exascale systems used in farming mode will allow to support the vision of flying the aircraft through the database.
- Prediction of noise sources and impact: This requires the full development of noise source mechanisms, acoustic radiation and noise impact simulation tools which compute acoustic disturbances on top of aircraft flow
- Multidisciplinary aircraft design based on high-fidelity methods: This requires a fully coupled simulation of the flow around a parameterized aircraft configuration and surface shapes covering a reactive structural model within a sophisticated optimization process. In addition uncertainty quantification needs to be included to allow for a robust aircraft design.

The coupled large-scale simulations will run multiple times, leading to farming applications on Exascale systems.

Such Digital Aircraft vision will have an impact on the use of HPC by all the aeronautics value chain, requiring the availability of leading edge HPC resources in both capacity (or farming mode) as well as capability. This will also require strong software development efforts for:

- Increasing the scalability of individual CFD, structure, acoustics, … simulation codes by working on new scalable numerical solvers
- Introduction of non HPC technologies (CAD) in multi-disciplinary optimisation
- Developing automatic grid generation tools for handling complex geometries
- Next generation of couplers for handling multi scale and multi physics heterogeneous applications
- Next generation of uncertainties/optimisation framework
- Handling and visualisation of big data

The requirements of HPC for Airbus provided into previous EESI-1 report are still valid and consistent with the roadmap proposed by similar industries in US (see chapter 6.1.2 for the NASA CFD 2030 roadmap) with perhaps few years of delay regarding the final 2030 objective.

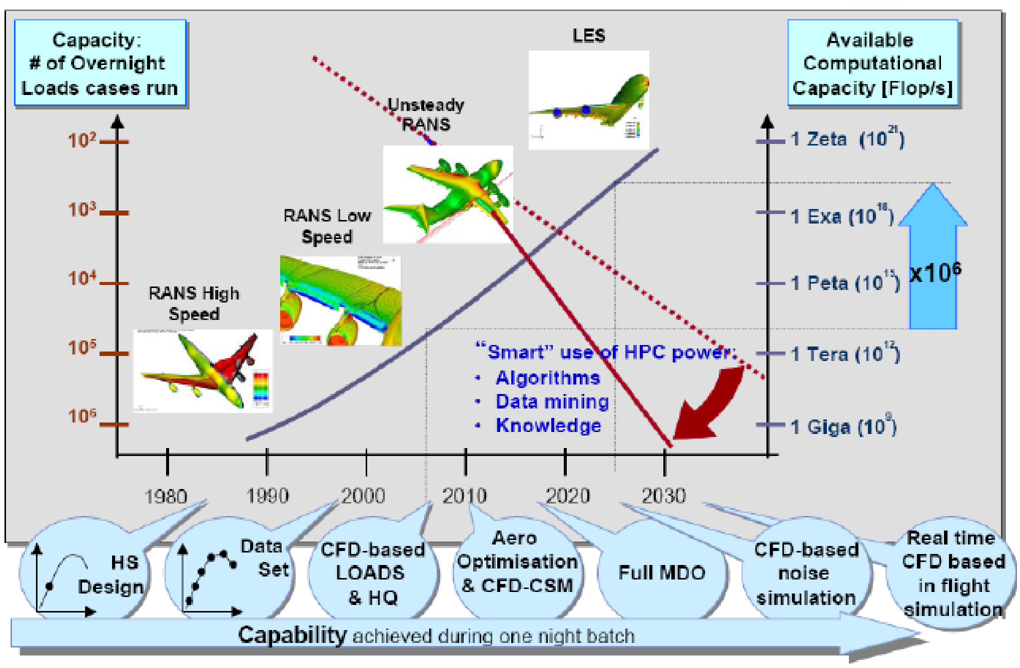

Figure 2 - HPC requirements of Airbus for a overnight run

5.1.2 In Oil and Gas

In the oil & Gas industry which is already one of the first (or the first) user of HPC resources the need of Exascale (and beyond) is mainly driven by the requirements of oil exploration for imaging more and more complex fields and the need to improve the oil recovery efficiency by reservoir modelling techniques.

In oil & gas exploration, the objective is to produce from a seismic campaign the best estimation of the underground topography in order to optimize reservoir delineation and production by solving the Full Inverse Wave Equation.

This application is largely embarrassingly parallel, and the higher performing the HPC system is, the better the approximation of the underground topography.
A roadmap of the steps of this kind of approach, showing the different necessary and more complex methods of approximations of the physical reality (e.g. elastic, visco-elastic, etc), is given below with courtesy of Total.
Such advanced algorithmic advances for oil & gas exploration and recovery could even unlock new types of reserves (called unconventional), including coalbed methane, tight sandstones and methane hydrates, potentially ushering in another "energy revolution".

The following roadmap provided by TOTAL during the previous EESI-1 roadmap and adopted by almost all the oil & gas majors is still valid. Almost all the majors and contractors (CGG Veritas, Schlumberger, …) have now massive computing capacities of more than 1PFlops each making Oil & Gas as one of the main user of HPC worldwide.

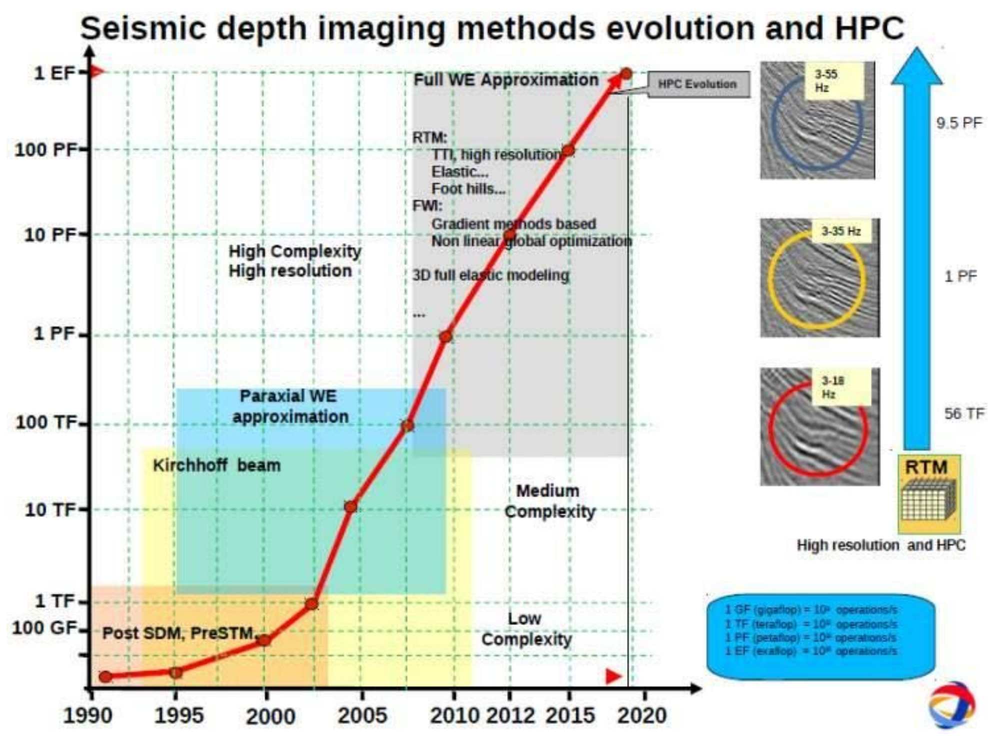

Figure 3 - TOTAL HPC roadmap for seismic modelling

When moving downstream oil & gas companies will also need to rely on HPC for running multiphase fluids dynamics simulation in reservoir modelling (for increasing the oil recovery ratio from a reservoir in production), pipeline modelling where the diameter is measured in fraction-of-meter, whilst the length of the pipeline normally is measured in kilometres or Global optimization of plants including multi-disciplinary interactions, e.g. simulating full plant lifetime and control systems, emulation of operations and risk analysis.
For these applications there is a need to rely on HPC infrastructures able to execute applications in a farming model.

Such roadmap for the oil & gas industry will need to work on the scalability of numerical solvers, the availability of multi-scale and multi-physics efficient couplers, the management of massive amount of data as well as the need to develop tools for handling uncertainties and verification and validation.

5.1.3 In Automotive

Beyond the requirements expressed by the automotive industry in the first EESI-1 report the experts of the working group presented new fields of research which raised with the enforcement of the crash regulation (EuroNCAP in Europe) and some disruptive technologies in the automotive industry like the use of composite instead of steel for lowering the weight of the car and by consequence the consumption and the exhausts.

The computational requirements to support the digital design process in the automotive industry (and in particular in RENAULT) evolve in different directions, all leading to an increase in the compute cost:

- Numerical models will continue to grow faster than the evolution of the equipment, as well as the volume of data generated. The evolution of the models is a consequence of the requirement of accuracy, but also the arrival of new materials, the digital representation is more complex (composite). Crash models, for example, will increase by an order of magnitude as the number of finite elements, which will increase the CPU consumption of two orders of magnitude. New calculation methods, eg implicit mesh-less and crash analysis, provide a gain in accuracy but high cost calculation;
- The number of different task profiles used for validation of a technical solution increases as the project decision is increasingly multi-objective, using a profile of extreme mission validation will be replaced by the envelope a set of profiles (car-to-car crash, dummies population modelling rather than "average" driver ...)
- Demand in robustness studies to validate an entire production rather than just nominal vehicle increases for all services and simulated motor vehicle. Existing techniques for propagating uncertainties involve an increase in the cost calculation;
- The weak coupling of models is mandatory to improve the accuracy of the simulations, the model becomes a work-flow. For example all the stampings of a vehicle body must be treated by a stamping code (which modifies the materials properties) before the crash calculation. Similarly, the arrival of composites involves a systematic coupling product / process;
- Responsiveness must also increase to better integrate simulation tools for displaying results, so as to multiply rapidly examining a large number of scenarios ("what if" method);
- The main benefits of the vehicle (comfort and performance, aerodynamics, safety) become active, they should be included in procedures for point-to-control laws (aerodynamics component, for example);
- The design process includes optimization studies; the number of parameters is increasing dramatically with the routine optimizations form. This also implies an increase in the number of simulations.

Automatic reduction of large models and the size of the design space will become checkpoints for the control and optimization, optimization e.g to produce a validated numerical models with heavy solution but can use scale models to find this solution.
The activity of the WG regarding automotive has been quite focused for the moment on crash and structure mechanics, in the future this work will also include advanced CFD dealing with aerodynamics, aerothermal exchanges and combustion.

5.1.4 In Power generation and nuclear plants

In this industrial domain, the objectives are multiple: first Improvement of safety and efficiency of the facilities (especially nuclear plants), and second optimization of maintenance operation and life span. This is one field in which physical experimentation, for example with nuclear plants, can be both impractical and unsafe. Computer simulation, in both the design and operational stages, is therefore indispensable.

In the thermal hydraulics CFD the improvement of efficiency may typically involve mainly steady CFD calculations on complex geometries, while improvement and verification of safety may involve long transient calculations on slightly less complex geometries.
This will require HPC for the study of flow-induced loads (to minimize vibration and wear through fretting in fuel assemblies), flow-induced deformation in PWR cores and the use of detailed simulations designed to verify and increase safety.

In order to validate such models, it twill be mandatory to run quasi-DNS type calculations on subsets of the calculation domain may be necessary.
Such studies will require multi-Exaflop capacities during few weeks.

In the neutronics domain the needs will target the development of transport calculation based on kinetic Monte Carlo methods using coupled multiphysics codes as described:

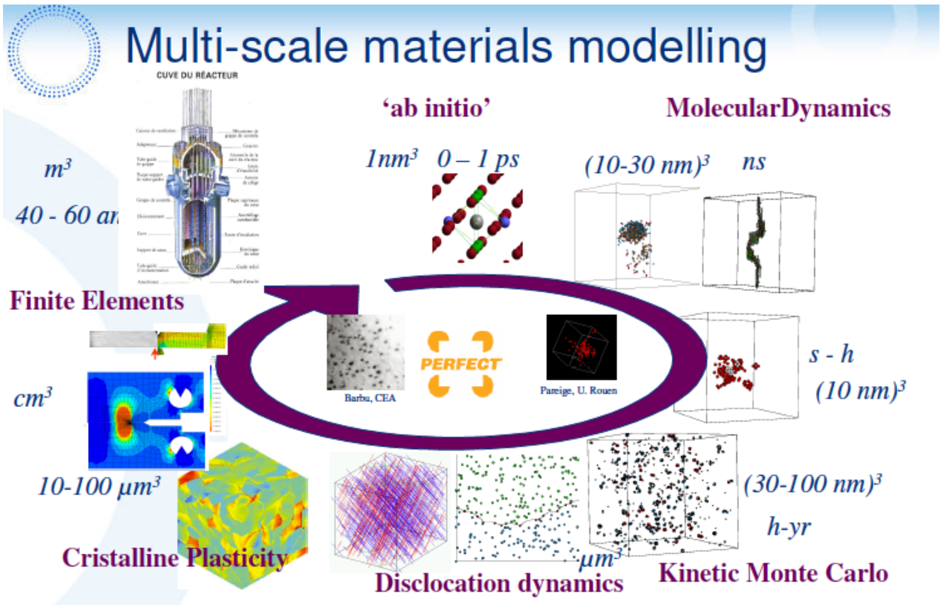

Figure 4 - Multi scale framework used for Neutronics - Courtesy of EDF R&D

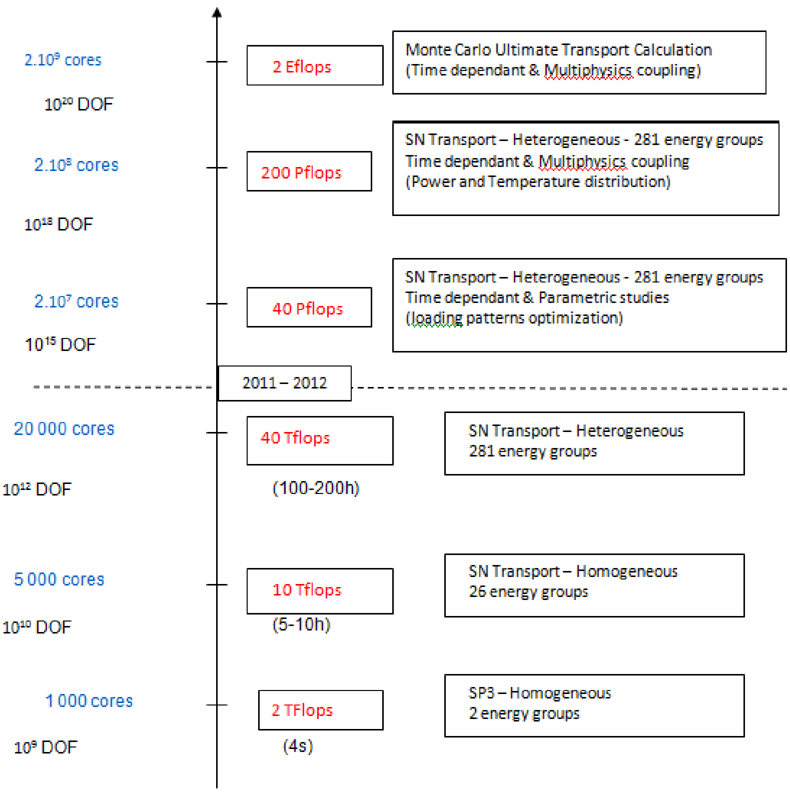

Figure 5 - EDF R&D roadmap for Neutronics

In this field HPC is mandatory to reach the goals of efficiency and safety of the plants. Next challenges will rely on the development of improved CFD codes as well as the availability of next generation of couplers and uncertainties quantification platforms.

5.1.5 In Combustion

As already sated, combustion is used to produce around 90 % of the earth energy and is essential for ground and air transportation, electricity production, industry applications or safety.
The central position of combustion in our world will not decrease before a very long time and science, especially simulations must allow combustion to continue with the lowest impact on climate and the highest efficiency.

In the previous EESI-1 roadmap the need to develop leading edge CFD methods like LES or DNS with a strong expertise in Europe (CERFACS, CORIA or Ecole Centrale Paris in France, University of Aachen in Germany, Lund University in Sweden, University of Manchester in UK, …) has been highlighted by the experts.

During last year such teams have been able to use massive HPC configurations in Europe (PRACE) as well as in US (through the Incite DoE calls) to demonstrate the efficiency of their LES and DNS codes for improving combustion process and reduce instabilities.

In June 2013 a new project leaded by T. Poinsot (CERFACS) thanks to an ERC grant and called INTECOCIS just started. It aims at introducing recent progress in the field of High Performance Computing (HPC) for combustion simulation into studies of Combustion Instabilities (CI).
The target is to build simulation tools that can predict the occurrence of CI in future combustors before their construction. In order to achieve this goal, the simulation tools used today for CIs must be revolutionized to integrate recent HPC capacities and have the capabilities and brute power required to compute, understand and control CI phenomena. A second objective is to add Uncertainty Quantification (UQ) methods to CI tools.
The project will integrate experimental validations and industrial applications. The tools will be made available to European laboratories working in the field of combustion instabilities. INTECOCIS will be based on two teams: the PSC (Particle Sprays and Combustion) team at IMFT (CNRS) that will provide the CI expertise and the CFD team at CERFACS that masters HPC aspects required for combustion simulations.

The next challenges for these teams are now to integrate such technologies into the industrial world by not only solving flow field but also being able to perform multiphysics coupled simulations on complex industrial geometries.
This will require to use not only the CFD code into the combustion chamber but also to integrate all the physics into the whole reactor, powerplant, piston engine, … and to make it used daily by industry in collaboration with academia.
A lot of aeronautics or gas turbines companies developing engines or power plants are now planning for 2020 to be able to simulate the whole reactor by coupling multiples instances of high fidelity CFD codes with the acoustics, thermic…

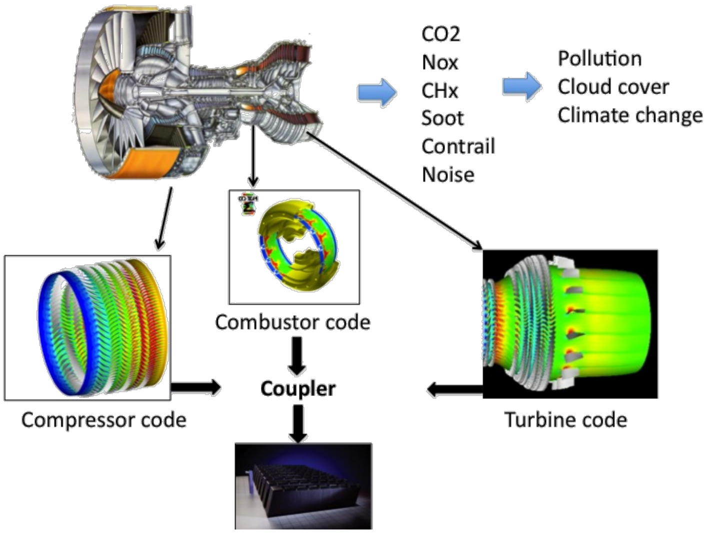

This will require to work on the improvement of the scalability of each single code but also to develop a new generation of code couplers for handling efficiently high frequency exchange of large amount of data between codes, management of different time and space resolutions, heterogeneity of data structures, heterogeneity of HPC resources used, asynchronous communications, management of I/O, standard API, …

This was the purpose of the proposal elaborated by CERFACS and all and submitted during the EESI2 face-to-face meeting in Tremblay end of May 2013 which has been selected as one of the 12 global EESI2 recommendations to EC.

5.1.6 Common key issues for reaching Exascale

The energy and transportation industries, and more generally engineering industries, are looking for multi-petaflopic and exaflopic computing resources: not only the next supercomputers themselves but also the ways to know how to program on 10^5-10^6 cores machines.
Industry is in fact expecting the best coupling between Architecture / Algorithm / Application, in order to address and solve on exaflop systems crucial issues in energy and transportation, along with other economically and socially important engineering challenges.

For all the experts engaged into EESI2 WG3.1 on industrial and engineering applications, a broader use of Exascale systems by European various industries will be possible if:

- On the infrastructure side it will be important to support both capability simulations (the experts state that some "hero" application will be able to scale out to a full system) and capacity simulations (for supporting loosely independent executions of multiphysics and multiscale coupled simulations as well as uncertainties/ensemble and optimisation studies.).
- On the software side 5 major items have been identified by the experts:
  1. The need to develop a next generation of couplers in order to support efficiently multiphysics and multiscale simulations. As described in the previous chapter for combustion a lot of issues need to be addressed by these tools and a concrete proposal has been elaborated by the working group. This includes the development of advanced parallelization job-scheduling systems
  2. The development of an automatic and highly parallel grid generation tool for Exascale
  3. The improvement of the scalability of numerical algorithms and again the working group is part of the proposal elaborated by EESI2 on that side (so called ultra scalable algorithms).
  4. A standardized and efficient data management from end to end with challenges in pre and post processing, visualisation, computational steering as well as data analytics, curation/archive and dissemination to various stakeholders.

5. The need to support both legacy and new software platforms by using unified simulation frameworks and associated services.

5.2 Challenges in Weather, Climatology and solid Earth Sciences

During the first period the WG 3.2 focused its activity on assessing more deeply and refreshing their roadmap regarding 2 majors challenges: Scientific Data Management and Earth Sciences modelling.

5.2.1 Challenges in Scientific Data Management

Exascale challenges in the domains covered by this WG are equally a big data and HPC challenge. The preliminary thinking emphasized some challenges connected with data production, data analysis, data storage, and data used by a larger community and mainly related to I/O issues and parallel files systems, scientific data discovery and visualization, parallelization strategies and new frameworks for parallel data analysis, novel programming models for big data, new data structures and storage models, data provenance, and metadata management.

Architecture

To address big data management issues, several changes at the architectural level need to be taken into account. New storage systems are needed to seamlessly integrate heterogeneous storage technology like tape, disk, flash, PCM and for some workloads memory targets. Automatic partitioning and distribution of data could also enable parallelism whereas indexing, replication, and data management hierarchies could improve execution efficiency and throughput. Such a new storage systems would offer intelligent auto-tiering, support novel user-interfaces and allow consistent monitoring. In this regard, one advantage of the novel user-interfaces would be their capability to integrate and take into account several aspects (like semantic ones) of data as well as its entire life-cycle to improve fault tolerance, availability and efficiency. For example, depending on the intended data usage the appropriate storage targets should be automatically identified. The same would apply to automatic background backup strategies and data migrations.
From an architectural point of view, new storage systems should provide new storage models, strongly related to the intrinsic n-dimensionality of the WCES data model to better address the WCES scientific (big) data management challenges. Such novel architectures would be strongly tailored to specific needs like sub-setting in space and time (as many end-users often wants to extract only a very small part of the data), time series analysis and data reduction.
Finally, implementation of data streaming (in memory) for online parallel post-processing will be key to reduce the “time to insight”.
From an operational point of view, since the novel storage system described above would support legacy schemes such as POSIX or MPI-IO this transition could happen smoothly. Anyway, a transition to a new interface is mandatory to gain advantage of the complex future storage landscape. A different support in terms of file-system would be also needed to address WCES challenges like data provenance. In this regard, scientific data would need a stronger support at the file system level to describe the context (metadata information) into more complex “objects” rather than “files”.
Projecting the current state different systems will co-exist, e.g., big data archives, extreme computing machines, databases, storages, in-situ visualization and parallel file systems. Presumably accelerating technology such as burst-buffers for file systems will be used. This will increase the burden on users to manage data, especially since each technology offers other APIs and will lead to decreased efficiency and research productivity. As a consequence, revolutionary alternatives such as the sketched novel integrated heterogeneous architecture needs more research efforts.
The amount of data needed/created by WCES simulations forces WCES scientists to do a part of the pre/post-processing on tens/hundred of cores directly on the storage system seen by the machine where the simulations are performed. As part of the co-design aspect, active storage management research would allow moving the analysis data kernels at the storage layer, drastically reducing I/O transfers via in-storage computations. This research will need a strong involvement of both hardware

and software competencies linked with interdisciplinary groups.

Workflows

The current workflow, from simulation to analysis includes several steps like: (i) configuration of model run, ensemble members, code, scripts and data (even including observational data), (ii) compiling, linking and defining IT sources, (iii) pre-processing initial and boundary data from distributed data nodes and diverse storage systems, (iv) reading input and boundary conditions fields from distributed data sources, (v) running WCES model experiments, which includes both writing data, metadata during processing, online visualization, etc. and comparing results to reference data or previous results, (vi) I/O during model run which may include post-processing and using new data directly for next run as well as producing provenance and meta data and storing data, (vii) I/O of final results from parallel environment to other distributed sources for post-processing, adding provenance and metadata and storing the results and finally (viii) archiving results after quality checking.
Looking to the future, as storing all data produced at high-resolution would not feasible, a typical workflow could be the following: (i) agree, with other scientists and impact people, on a number of spatio-temporal domains and number of variables of particular interest, (ii) run a first high-resolution simulation, store only a reduced number of parameters at global lower resolution (for global analysis of simulation) and store the variables of particular interest on the domains of particular interest at high resolution and (iii) if a further particular study needs additional high-resolution data, rerun the simulation to produce this additional data. Such a workflow could also include on-line statistical diagnostics (e.g., mean, min, max, variance, skew-ness, kurtosis, percentile) of ensemble experiments in order to avoid storing each member.
On the other hand, since the WCES community publishes and shares the output of WCES simulations into national and international networks like IS-ENES, Earth System Grid Federation, SeaDataNet or MyOcean, the workflow commonly used for scientific data discovery is generally based on the following steps: (i) search, (ii) locate, (iii) authN/authZ, (iv) download and (v) analyze steps. This workflow will not be feasible at peta-exascale. The model will fail for several reasons including: ever-larger scientific datasets, time- and resource- consuming data downloads, and increased problem size and complexity requiring bigger computing facilities. Peta-exascale data will require a different workflow based on data-intensive facilities close to data storage and server-side analysis capabilities. Only the final results of an analysis (e.g., images, maps, reports and summaries typically megabytes or even kilobytes) will need to be downloaded—and even those data may be server-side managed (it could be even stored in a cloud-based environment). Such an approach will reduce the downloaded data, the make-span for the analysis task, and the complexity related to the analysis software to be installed on client machines.
In all cases the data post-processing, analysis and mining plays a key role for the WCES workflows to get some insights from the WCES data.
Currently most of the available software is based on sequential codes. Compression methods such as netcdf4 or jpeg2000 must also be investigated to reduce the amount of data but these methods are often weakly compatible with (strong) parallelism.
Scalable, parallel and robust software frameworks (runtime environment) to support the aforementioned workflows are needed especially in view of running large ensembles of model experiments. Additionally, the WCES community would strongly need:

- on-line statistics on ensemble simulations, and data reduction,
- efficient compression,
- parallel tools for analysis and visualization,
- efficient portable I/O library or server, including on-line reduction of data,
- user friendly and robust (error handling) workflow environment to run and re-run simulations easily changing the configuration,
- quality control of data,
- alternative data formats (e.g. HDF5) to reduce the number of files to be moved and/or generated on the HPC systems

A plan for achieving interoperability relates to hardware (in terms of comparable / sufficient per node memory on different systems, fast network and sufficient temporary storage space for pre- and post-processing and re-arranging data), software (in terms of APIs for doing data processing, analysis, mining, visualization, etc.) and conventions for data, models and metadata.

Taxonomy

The use of data involves various levels of data-centric computing facilities such as:

- direct data sub-selection and data download: data reduction at data centre,
- specific data analysis at data centre and delivery of results: widely required derived data products,
- generic data analysis at data centre: supporting data centre federation wide distributed data analysis workflows.

To address scientific needs it is key to carry out (i) maintenance and publication of powerful consistent metadata repository users can search to answer simple data metadata related questions (content and context of data) – ("first level data portals"), climate impact gateways ("second level portals") for climate impact community members with specific research interests accessing first level data centres and data portals and finally (iii) climate service centres with specialists for specific climate data science questions.

A data driven mini-application demonstrating the new usage model in the climate domain is related to the analysis of climate indicators from distributed (*e.g.*, CMIP5) petabyte data archive (existing worldwide climate data node federation).

It relies on the following steps:

- interaction with global metadata index to perform a data discovery task,
- data centre near deployed compute services generate climate indicators from local data,
- local results are collected and used to compute final experiment results,
- all data, metadata, services (including triggered analysis code) etc. are persistently identified using persistent identifiers,
- generated data is linked to involved data, metadata as well as processing services to enable data provenance,
- the generated data is published and shared in a collaborative workspace,
- user group can reproduce, discuss, evaluate results which include their provenance,
- final result is communicated to decision makers (*e.g.*, insurance agency) – data results are reproducible and fully described according to their provenance.

From the geospatial community, the super site exploitation platform highlighting the use of interferometry on radar data to prepare for the sentinels satellites can be considered another notable example. For seismology, the VERCE project provides some demos for CPU-intensive and data intensive seismic use cases.

Finally, the most relevant cross-cutting concerns for big data are data integrity, data reproducibility, data quality and data integration. About data quality, the exponential growth of climate simulation data requires new methods of data quality assurance and flexible data annotation mechanisms. In terms of data integration, integrating structured data (*e.g.*, existing petabyte scale climate archives) and semi- as well as un-structured data (derived results, quality related information, user annotations, data discussions and data citations etc.) is not trivial due to the heterogeneity of the involved data sources. Data reproducibility is strongly connected with data provenance which helps reproduce analyses and products.

Software

The WCES community exploits a wide set of software to manage and explore the data. In particular to manage the data the most notable software used at the supercomputing and data centres are:

- Oracle DBMS, Postgresql, MySQL relational databases,
- Solr/Lucene
- iRods
- HPSS (High Performance Storage System)
- Parallel file system (GPFS)
- ftp, rsync for data transfers

On the other hand to explore data the most relevant software/middleware are:

- grid-middleware like Globus or gLite,
- web services exposing specific data reduction, analysis and visualization capabilities,
- distributed P2P based data federation (petabyte scale) for data sharing and publication
- climate data analysis tool box (cdo) and netCDF Operator (nco)
- IDL, Matlab, R, webgis like arcgis,
- in-house statistical software (based on fortran, python libraries)
- cmor for data format and verification
- NCL for analysis and visualization

Additional work is needed to improve the data analysis, parallel I/O, parallel post-processing, automated metadata generation. For the majority of them a server-side approach is the proper solution to analyze the data where they are produced, taking advantage of the parallel machines available at the data centres.

However, along with data issues at the software level, additional work is also required at the metadata level to address a scalable metadata and persistent identifier management (for instance involving new distributed cloud key/value technology).

Looking to the future, the main gaps that need to be filled are related to the (i) integration of security solutions for distributed – multi-organizational – data management, (ii) integration / interlinking of user community added value data with data centre managed data, (iii) data quality assurance, (iv) complete provenance tracking.

Interoperability Challenges

In terms of interoperability challenges, there are several examples connected with metadata handling for provenance, data quality, descriptive information, etc. which help in better describing both the content and the context of the data, even including its history, quality and links with other entities, thus enabling data interoperability.

In this regard, some relevant examples are (i) the integration of annotation info in file level metadata (*e.g.*, in netcdf, hdf data headers), (ii) the automated data centre side analysis code to generate data quality related annotations, (iii) the users / user side data analysis.
From a service-level point of view, in such a context a basic persistent identifier service relates externally managed annotations (*e.g.*, data from user surveys, questionnaires and cloud storage based commenting tools) with persistently identified stored data in data centre. higher level services exploit this basic correlation layer to enable more flexible and complex services (*e.g.*, search). In this regard a key prerequisite is a scalable persistent identifier infrastructure (as discussed in the research data alliance RDA) as well as tools exploiting this infrastructure.
Information systems for providing semantic capacity in terms of effective translation between data and conceptual models used by different communities are ontology-like system.
The systems currently used for providing information about the actual use of both observational data and simulated data are usually web server, data server analysis systems, user support systems, data digital object identifiers (DOIs) for referencing data in papers and a wide (and often weakly linked) set of web portals.

Social aspects

Social challenges are becoming more important as scientific efforts in various disciplines move towards large-scale experiments/environments (*e.g.*, the Large Hadron Collider at CERN for High Energy Physics and the Coupled Model Intercomparison Project, phase 5 (CMIP5) for the Climate Science). In this regard, new collaborative environments devoted to ultra-scale data analytics and connected with social networks could change the way scientists interact each other both inside (for research purposes) and outside (for dissemination purposes) the scientific community. Finally, training and academic courses on exascale computing and data intensive science, both from a software and a hardware perspective, will be needed to prepare the next generation of scientists and fill the skills gap.

Revolutionary vs. evolutionary approaches for scientific data management

We suggest that in the 2020 timeframe the scientific community should take the chance to exploit revolutionary strategies for large-scale scientific data management. There are strong domain-based I/O and real-time analysis requirements, a strong need for more efficient and scalable data structures, parallel primitives, programming models and parallel frameworks to manage, compare and visualize big datasets. At the same time, simulations, observations, and re-analyses are producing new, richer and more complex data. Evolutionary approaches may require less short-term effort but are unlikely to succeed in all cases and at large scales. On the contrary, revolutionary approaches would require a from-scratch design and long-term implementation strategy. They may need larger investments, but would definitely target and focus on large-scale data intensive issues and requirements. They will permit hardware-software co-design strategies to be considered from the beginning, as part of the overall design.

5.2.2 Challenges on Earth Science Modeling

Earth Science Modeling challenges were already addressed in the PRACE report “The Scientific Case for High Performance Computing in Europe” by the PRACE thematic panel on WCES (chaired by ENES). Moreover, ESM challenges were also discussed during two earlier workshops organized by the IS-ENES HPC Task Force, in Lecce, Italy, in Dec. 2011, and in Toulouse, France, in Jan. 2013. A summary of these challenges and the main conclusions from the IS-ENES workshops are reported in the following.

Running very high-resolution models to better understand, quantify and predict extreme events and to better assess the impact of climate change on society and economy on the regional scale

Modelling the climate system is a challenge because it requires the simulation of an extremely large number of interacting and complex processes as well as their analysis at different time and spatial scales. Climate system modelling requires sophisticated numerical models, due to the inherently non-linear governing equations. Huge computational resources are needed to solve billions of individual variables describing the physical processes at different scales. Indeed model simulations are required to represent both modification of the larger-scale, global, state (inside which extreme events are developing) and the fine-scale temporal and spatial structure of such events (storms, cyclones, intense precipitation, etc.).

Currently global climate models have typical grid spacing of 100-200 km and are limited in their capacity to represent processes such as clouds, orography effects, small-scale hydrology, etc. The latest generation of models, under development or just starting to be used, have grid spacing in the 20-50 km range and there is evidence that a number of important climate processes are better represented at this resolution (*e.g.*, ENSO – El Nino Southern Oscillation- , blocking, tropical storm numbers etc.). A priority is to continue the development of coupled models at such high resolution and use them in multi-member multi-model inter-comparisons focused on key climate processes.

In weather forecasting applications, much higher convective-resolving limited domain models are now being used operationally. However, these models cannot be run globally for climate because of the prohibitive cost of associated computing resources and limits in model scalability. The climate community first “grand challenge” on the longer term is then to develop global climate models that resolve convective scale motions (nominally around 1km horizontal resolution).

These very-high resolution models will directly resolve convective systems, allow a better representation of orographic effects, atmosphere and ocean energy and matter transport, and provide greater regional details. They will allow determining whether or not convective scale resolution is necessary for credible predictions of some important aspects of regional climate change. Developing such very high resolutions will require developing scalable and more efficient dynamical cores and improving physical parameterisations.

Moving from current climate models towards full Earth System Models

Today it is clear that models must also include more sophisticated representations of non-physical processes and subsystems, which are of major importance for long-term climate development, like the carbon cycle. Scientists are strongly interested to know the sensitivity of predictions not only to unresolved physical processes (as, *e.g.*, the cloud feedbacks mentioned above), but also to non-physical ones, like those related to biology and chemistry (including, for example, those involving the land surfaces, and -greenhouse gases reactions).

It should be noted that including the representation of biogeochemical cycles using different biochemical tracers and aerosols typically increase computing time by a factor of 5 to 20 (depending on the complexity of the parameterizations and the number of tracers). An increase of computing power by a factor of 5 to 20 is then required to better account for the complexity of the system.

Quantifying uncertainty

Future projections of climate change are uncertain for a number of reasons. The future forcing by greenhouse gases and aerosols are uncertain, and climate variations have both a natural and anthropogenic component and both need to be represented in climate models. The models are also inherently imperfect owing to physical processes that are either not completely understood or yet to be adequately represented because of limited computer power.

To better understand and predict global and regional climate change and climate variability using numerical models, there is a wide range of underlying scientific issues that need to be solved by the international community, as reported in the WCRP strategy COPES "Coordinated Observation and Prediction of the Earth System" 2005-2015 (http://wcrp.ipsl.jussieu.fr/).

The major issues, taking into account the interests and strengths of the European climate science community and the aim to answer societal needs, are related to (i) the predictability – and its limits - of climate on a range of timescales, (ii) the range of uncertainty that can be fully represented using the models currently available and (iii) the sensitivity of climate and how much we can reduce the current uncertainty in the major feedbacks, including those due to clouds, atmospheric chemistry and the carbon cycle.

The consensus approach to solving these problems is to assume that the uncertainty can be estimated by combining multi-model multi-member experiments. Running in a coordinated European way multi-model experiments allows to investigate the sensitivity of results to model parameters. Moreover, running multi-member ensembles of climate integrations allows accounting for the chaotic nature of climate and thereby enables systematic assessment of the relative roles of natural climate variability and man-made climate change.

Use of different scenarios of emissions of greenhouse gases is also mandatory to warrant the production of experiments allowing probing the future course of climate. And quantifying uncertainty of future climate change will also require to investigate the sensitivity of results to specification of the initial state; the latter initialization issue is particularly important for decadal timescale predictions where both natural climate variability and man-made climate change need to be predicted within the model. It should be noted that computing requirements scale directly with the number of ensemble members, which are required to better represent uncertainties associated to both internal variability and model parameterizations, but the number of members required to keep the same signal-to-noise ratio in climate forecasts also increases as the spatial resolution increases! Ensemble experiments are then computationally expensive – a factor 10 to 100 for each experiment – but will bring enormous economic benefit as they will improve reliability of models and our understanding of uncertainties in forecasts.

Investigating the possibility of climate surprises with longer simulations

In a complex nonlinear system such as the Earth system, minute actions could cause long-term, large-scale changes. These changes could be abrupt, surprising, and unmanageable. Paleoclimatic data indicate the occurrence of such rapid changes in the past. For example, it is crucial to determine if there are thresholds in the greenhouse gas concentrations over which climate change could become irreversible. The Atlantic thermohaline circulation (THC) might undergo abrupt changes as inferred from paleo-records as well as from some long simulations of future climate. The possible climatic consequences of such a slowdown in Atlantic THC are still under debate. Surprises may also arise from ice-sheet collapse and large amount of fresh water in the ocean.

Some key questions arise on the possibility (i) to model and understand glacial – interglacial cycles, including changes in carbon cycle and major ice sheets, (ii) to use observational evidence from past climates to calibrate the sensitivity of complex climate models and respective adjustable an model parameters and (iii) to attribute signals in the period of the instrumental record to understand Earth system processes (from weather scales to those typical of anthropogenic climate change). The need for longer historical runs, both current-era hindcasts and palaeoclimates, and for longer runs to investigate possible future nonlinear changes is evident, and the computing needs scale accordingly.

Investigating climate surprises requires longer simulations on future and past periods with medium to high resolution and various degrees of complexity. A factor 10 to 1000 is required which, however, cannot be achieved just by an increased number of cores but also by an increase of power for each individual cores.

Main conclusions from the IS-ENES workshops

As already mentioned various issues related to modelling challenges in the field of climate and Earth-systems have been addressed during the two IS-ENES workshops. The conclusions of these workshops will be revisited shortly with a more general WCES point of view. It is nevertheless of interest to mention some of the preliminary points of general agreement (see André et al. [ISENESWS2013]):

- agree on metrics (*e.g.*, for scaling, definition of the variables to be compared –simulated years/day, model configuration, horizontal resolution, etc.) in order to facilitate inter-comparisons of modelling approaches and to disseminate best practices;
- develop interdisciplinary teams where computational and application scientists can work together to address specific scientific issues
- new dynamical cores (*e.g.*, based on new grids) have been the subject of intense development and runs are now possible up to 105 cores. The need is to keep reviewing their advancement and their use in various parts of climate models
- using GPU’s and other types of hybrid cores for climate modelling would require solving the issues of low memory per core and the available bandwidth to access it, the number of threads per task per core, the silent errors, …
- need for revisiting the model code structure with respect to questions such as
    - How to ensure more modularity in the codes (component approach) and better isolate the “science” from the underlying technical software layers (code infrastructure, *i.e.* utilities for parallelization, I/O, etc., and code superstructure, *i.e.* the shell assembling and interconnecting the components)?
    - How to separate the scientific software from underlying implementation using underlying software kernels which might utilize unfamiliar programming models?
    - How to access to more efficient algorithms working with much higher parallelism? This is seen as a major disruptive technology with high positive influence on climate modelling techniques;
    - More generally and on a longer time scale, whether or not we should try to converge on common code infrastructure and superstructure and how to increase their adaptability and robustness.
- address issues related with the production, storage and use of big data (see section 4.2.1).
- how to increase the scalability to achieve higher resolution and the shorter time-steps required to run at higher resolution and to include more complexity in the models. As the number of calculations per grid-point rises, more scalability is needed. The scalability also needs to be increased for addressing (very) long simulations with small or medium spatial

resolution, as it is the case for paleoclimatic studies and long-duration climate scenarios. The issues here are quite different from those related to high-resolution

- how information can be exchanged between the increasing number of components in the system linked to the increase in complexity (by implementing more efficient coupling)
- and how data can be moved around the system at every level more efficiently: CPU to cache to memory to disk to archive to post processing platform and back to memory, as well as between cores and nodes.

5.3 Challenges in Fundamental Sciences

The first part of the European Exascale Software Initiative was finished in Sept. 2011. Since then the development of hardware has seen a further development towards exaflop computing by an increase from about 1 Pflop/s (K computer) [T500-611] to 33+ Pflop/s (Tianhe-2) [T500-613] in compute performance, while the power consumption increased only from about 10 MW [T500-611] to 17.8 MW [T500-613]. Although this trend is showing in the correct direction (increasing performance while relatively decreasing power) exaflop computers will need different a further strong improvement, since the announced limit of 20 MW power consumption is nearly reached. A similar trend is observed in software development and the solution of grand challenge problems. There are developments towards higher performance and better scalability, but to make a substantial progress to solve the grand challenges, which were outlined in the first report of EESI, the time is too short. Therefore, in this section some progress is reported for several software developments application domains from fundamental sciences, demonstrating progress in keeping track of developments in HPC architectures. Indicative for performance considerations the Gordon Bell price can be considered, which is awarded once a year for best performances during the Supercomputing Conference in US. Also the portfolio of large scale projects running in a European context within PRACE projects can provide some insight about progress in application areas.

5.3.1 Astrophysics and cosmology

In 2011 and 2012 codes from materials science and cosmology, running on the K-computer and Tsubame in Japan as well as at LLNL in US attracted high attention. In 2012 cosmology could be identified as a driver for high performance computing, as it was awarded in two categories for a large simulation as well as demonstration of large scale data-analysis, where efficient evaluation of two-point correlation functions was demonstrated for systems, composed of multi-billion particles and which is used to characterize the distribution of matter/energy in the Universe. The latter shows that not only production of data is now important for demonstrating efficiency and performance, but also data analysis. This directly relates to the topic of Big Data, which currently heavily discussed in the context of HPC, since faster performance on larger physical systems not only implies more insight and details but first of all more data and more elaborated analysis tools. Cosmological simulations focused

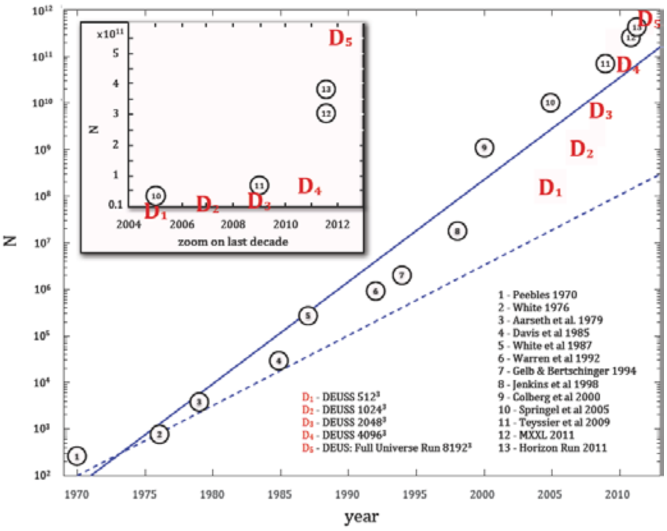

on dark matter distributions. Number of particles which can be treated by now has increased to more than a trillion. The performance of that simulation, carried out by a Japanese group on the K computer was 5.67 Pflop/s, corresponding to 98% of peak performance of the machine. A similar simulation, carried out by a US group at LLNL on Sequoia, achieved even 14 Pflop/s.
From the European side also a paper from cosmology was highly ranked demonstrating for the first time a simulation of the complete observable universe using 3 different initial dark matter distributions. This simulation was carried out on Curie and the simulation, performed for more than 500 billion particles combined with an adaptive grid with more than a trillion grid points (O(10000^3) grid points) and it was shown to scale to the entire machine. This simulation on more than 76 000 cores generated at a sustained performance of 50 GB/s more than 150 PB of rough data which has been reduced on the fly by an ad-hoc post processing toolchain to a final amount of 1.5 PB of refined data.
These results will be used as inputs maps for calibrating one of the next space telecscope called EUCLID which will be launched by ESA in 2019 for studying and making a first cartography of the dark matter. This is one of the first time that computational models are going to feed back large scale instruments.

5.3.2 Soft Matter

Complex soft matter systems have attracted attention in the field of coupled fluid/particle simulations in applications for blood flow through coronary arteries by European/Japanese and US/European-groups. Simulations have demonstrated that it is possible to perform simulations, in which blood cells are modelled as molecular dynamics particles, which are embedded into a mesoscopic fluid model, based on Lattice-Boltzmann description. The code showed very good scaling properties on a GPU

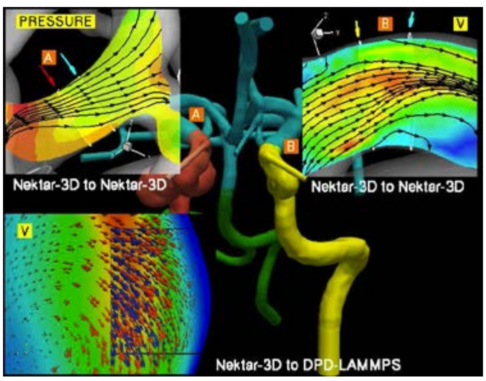

based system up to 4000 GPUs. A very similar simulation was carried out CRAY XT5 and BG/P architectures and included a description based on molecular dynamics, dissipative particle dynamics and Navier-Stokes solver.

5.3.3 Material Science

In the field of materials science, progress is to be reported for order N transport simulations, combining tight binding methods with wavepacket propagation, based on linear response theory. With order N methods it is possible now to simulate disordered systems with up to 100 million atoms.

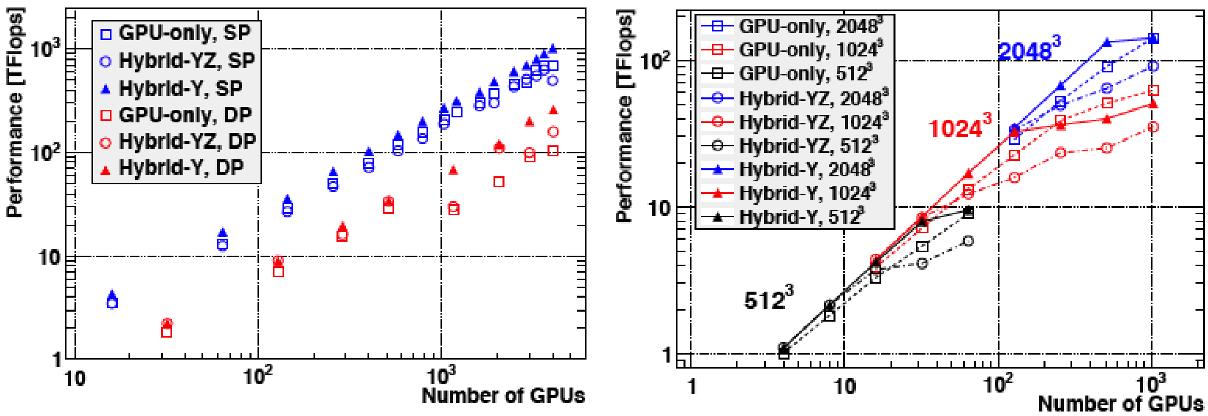

Within the Graphene European flagship, which was launched after the time of the last report, this type of simulation will be applied to evaluate properties of Graphene in the context of transport and chemical. Further progress was made towards atomistic nanoelectronics device engineering simulations, where sustained performances of more than 1.4 Pflop/s could be achieved [Luis11].
Not only in first principles simulations but also in mesoscopic approaches, like the phase-field method, progress was made towards highest performance.
Simulations for dendritic solidification was carried out on the hybrid CPU-GPU based system Tsubame-2, where a performance of 2 Pflop/s could be achieved [Shi11].

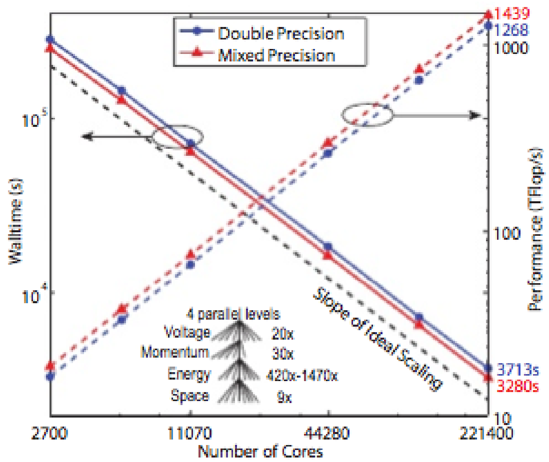

5.3.4 Plasma / Fusion Physics

Progress is made towards high performance implementations on leader class systems. Supercomputing Centre Jülich just recently launched the High-Q Club [JuQ13], where all code owners and developers are accepted whose code scales up to the whole JUQUEEN machine (BG/Q, No.7 top-500). Two of these codes originate from Plasma-Fusion community.

The GYSELA code is a non-linear 5D global gyrokinetic full-f code which performs flux-driven simulations of ion temperature gradient driven turbulence (ITG) in the electrostatic limit with adiabatic electrons. It solves the standard gyrokinetic equation for the full-f distribution function, i.e. no assumption on scale separation between equilibrium and perturbations is done. This 5D equation is self-consistently coupled to a 3D quasineutrality equation. The code also includes ion-ion collisions in global simplified magnetic geometry (concentric toroidal magnetic surfaces with circular cross-sections). The system is driven out of thermodynamical equilibrium by versatile sources which are capable of injecting separately parallel momentum and heat. Alternate sources have also been studied, in order to excite supra-thermal particles. Finally, the code has the originality to be based on a semi-Lagrangian scheme and it is parallelized into an hybrid OpenMP/MPI paradigm. Recent simulations have investigated: (i) the impact on Energetic-Particle-Driven Geodesic Acoustic Modes on Turbulence, (ii) ion transport barriers triggered by plasma polarization and (iii) turbulent momentum transport in core tokamak plasmas.

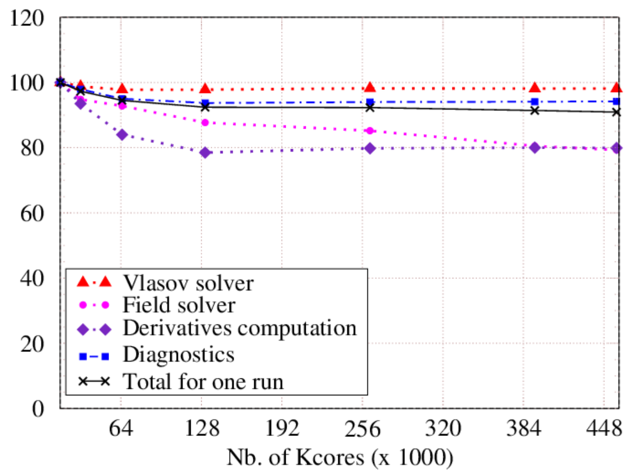

Figure 6 - Scalability of the Gysela code on IBM BG/Q systems

The code demonstrated scalability on IBM BlueGene/Q (JUQUEEN) on 458,752 cores using 1,835,008 parallel threads as well as on up to 65,536 cores on x86 systems (Curie and Helios).

The other code, PEPC is a tree code for solving the N-body problem. It is not restricted to Coulomb systems but also handles gravitation and hydrodynamics using the vortex method as well as smooth particle hydrodynamics (SPH). PEPC is a non-recursive version of the Barnes-Hut algorithm with a level-by-level approach to both tree construction and traversals. The parallel version is a hybrid MPI/PThreads implementation of the Warren-Salmon 'Hashed Oct-Tree' scheme. The long-range interactions are computed using multipole groupings of distant particles to reduce computational effort. The Barnes-Hut algorithm is well suited to dynamic, nonlinear problems and can be combined with multiple-timestep integrators.

This code is used in the G8-project NuFuSE to study kinetically consistent schemes for refinement and transport of plasma in strong magnetic fields. It was demonstrated that the code scaled up to 458,752 cores, invoking 1,668,196 parallel threads on BlueGene/Q as well as on 294,912 cores on a BlueGene/P.

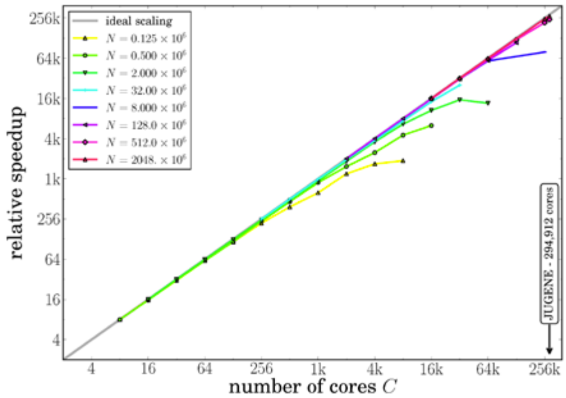

Figure 7 - Scalability of the PEPC code on IBM BG/Q systems

5.4 Challenges in Life Sciences and Health

5.4.1 New trends in the Pharma industry

The number of new NME per year is showing a positive and optimistic trend (Nature Reviews Drug Discovery 12, 87-90). The FDA approved 39 new drugs in 2012, a significant increase of the average over the last 20 years.
There are recent success cases of computer-aided drug design in Pfizer, Eli Lilly, GlaxoSmitKline, Sanofi (Drug Discovery Gets an Upgrade, WSJ 2013).

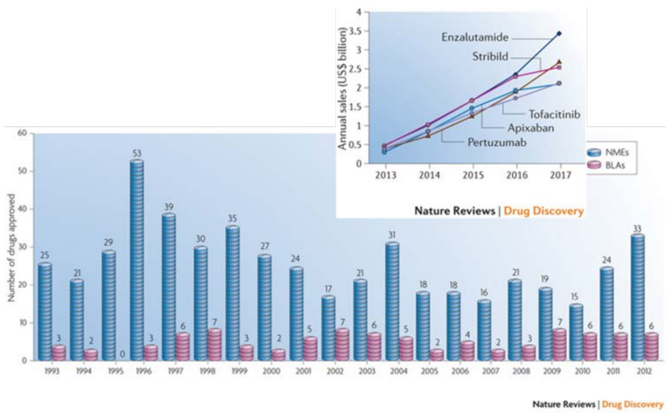

5.4.2 New challenges in Genomics

ENCODE. Initial results of the project were released in a coordinated set of 30 papers (September 2012) published in the journals Nature (6 publications), Genome Biology (18 papers), and Genome Research (6 papers). Identify all regions of transcription, transcription factor association, chromatin structure and histone modification in the human genome sequence. Thanks to the identification of these functional elements, 80% of the components of the human genome now have at least one biochemical function associated with them. modENCODE will continue the project targeting Drosophila melanogaster and Caenorhabditis elegans

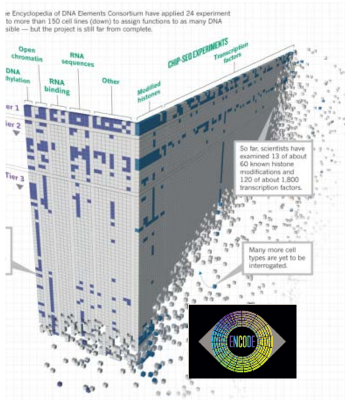

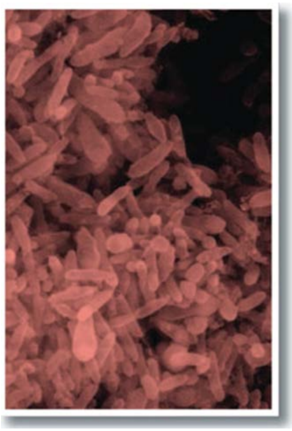

Metagenomics. FP7-Metahit - Science runner-up breakthrough 2011. Portfolios of resident microorganisms vary from individual to individual—even twin to twin—and from body part to body part. Our internal microbial communities fell roughly into three enterotypes, with a dominant microbe in each. The classifications weren't correlated with people's age, weight, sex, or nationality. Each enterotype differed in how it processed energy and in which vitamins it produced, factors that could influence the health of the human host. Provided more clues about the microbiome's role in disease, development, and immune function

5.4.3 New challenges in System Biology

Whole-Cell simulation. Whole-cell computational model of the life cycle of the human pathogen Mycoplasma genitalium from Genotype (Cell 2012). Entire organism (525 genes) modeled in terms of its molecular components. Integration of multi-format data and very fragmented data. Experimental analysis directed by model predictions identified previously undetected kinetic parameters and biological functions.

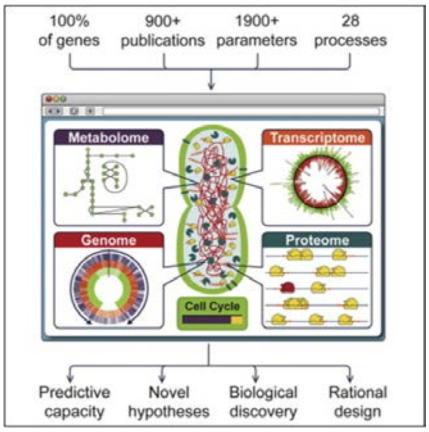

5.4.4 Challenges in molecular simulation

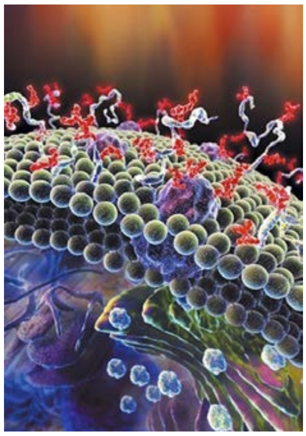

Automatic tools for high throughput MD simulations (such as MdWeb and NaFlex) and for multi-ensemble MD simulations coupled to Markov models to reproduce rare long-time transitions using thousands of shorter replicas are becoming more widely used. This mode of working maps well onto Exascale architectures.

The field is now facing significant challenges related to the rapidly growing mass of data being generated and the difficulty of storing and accessing this data, both within the community and for interactions with other academic and industrial partners. Actions are needed to create standards for simulation data and sharing protocols

Individual Anton-based long trajectories have provided detailed information on key elements of the regulation of cell growth. Schulten's group has shown (Nature 2013) the power of MD on ultra-large systems (the entire HIV1 capsid; 64 million atoms).

5.4.5 Challenges in Organ simulation

One of the main breakthrough and gap analysis wrt previous EESI-1 roadmap has been the decision of the EC to support and fund the Human Brain Project [HBP] in 2012. It aims to integrate the biological actions of neurons to create theoretical maps of different subsystems, and eventually, through the magic of computer simulation, a working model of the entire brain. It was one of the two FET 1B€ Flagship (ITFoM dropped) with Graphene. Ramp-up phase (2013-mid 2016), the project should focus on setting up the initial versions of the ICT platforms and on seeding them with strategically selected data. At the end of this phase, the platforms should be ready for use by researchers inside and outside the project.

The HBP flagship will rely on the use of massive HPC resources and the following chart shows the roadmap expected to be followed:

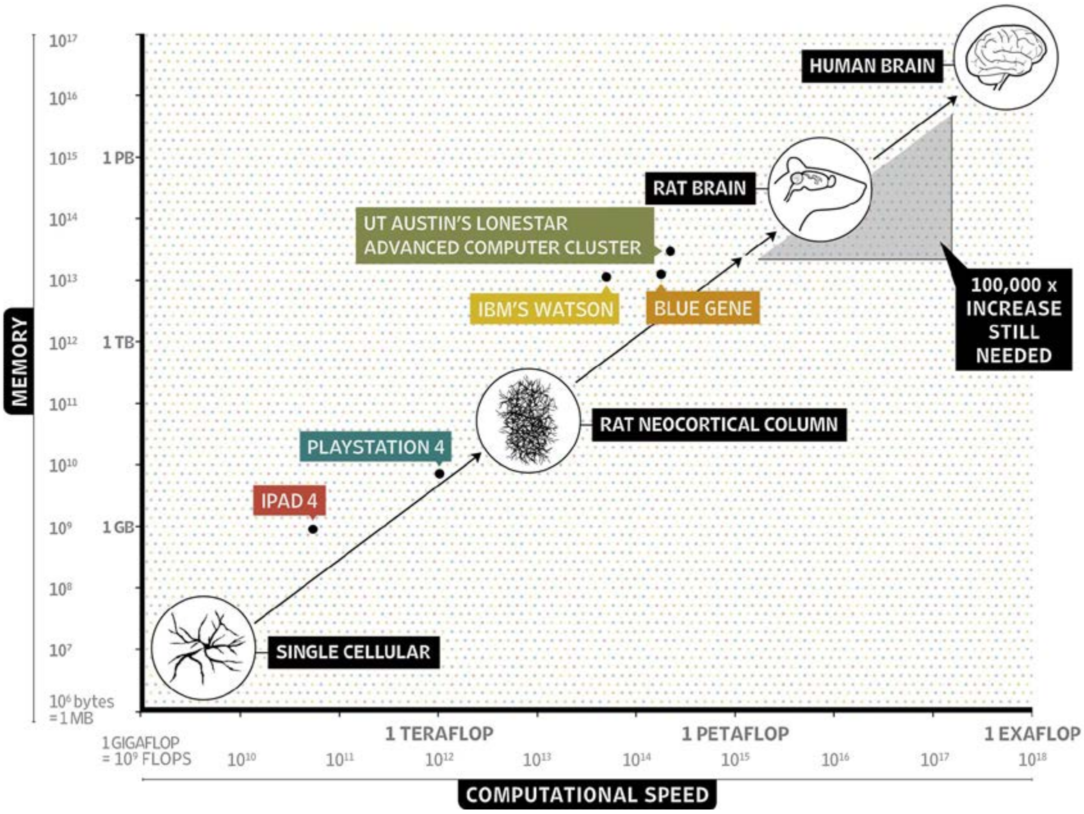

Figure 8 - Computational requirements of the HBP Flagship

On their side, the USA announced in April 2013 a competitor project ($100 million in 2014) called the Brain project.
It is different from the Human Brain Project, however, in that it has, as yet, no clearly defined goals or endpoint. Coming up with those goals will be up to the scientists involved and may take more than year.
The effort will require the development of new tools not yet available to neuroscientists and, eventually, perhaps lead to progress in treating diseases like Alzheimer’s and epilepsy and traumatic brain injury. It will involve both government agencies and private institutions.
The initiative, which scientists involved in promoting the idea have been calling the Brain Activity Map project, will officially be known as Brain Research Through Advancing Innovative Neurotechnologies, or Brain for short; it has been designated a grand challenge of the 21st century by the Obama administration.
Three government agencies will be involved: the National Institutes of Health, the Defense Advanced Research Projects Agency and the National Science Foundation. A working group at the N.I.H., described by the officials as a “dream team,” and led by Cori Bargmann of Rockefeller University and William Newsome of Stanford University, will be charged with coming up with a plan, a time frame, specific goals and cost estimates for future budgets.

VPH. Technological framework to enable collaborative investigation of the human body as a single complex system. Collections of anatomical, physiological, and pathological data stored in digital format, by predictive simulations developed from these collections, and by services intended to support researchers in the creation and maintenance of these models.

6. Report on Disruptive applications

6.1 Industrial and Engineering applications

The WG3.1 started to work in identifying potential disruptive technologies at the level of the application (and not at the level of the algorithms or the hardware/software which are already covered by other working groups in WP4 and WP5 of EESI2). This action is performed in strong relation with EESI2 WG3.5 which aims to gather all these inputs into a single recommendation for Applications.

In the industrial and engineering domain, one of the disruptive technologies identified by the automotive domain is the rise of the use of composite materials instead of steel. This will allow car makers to develop even more robust and light cars and this will reduce the carbon footprint of such cars. On the other side, as already mentioned into previous section this will have for consequence to rethink totally the way to perform crash simulation (a car will be composed by many less pieces than the current ones with steel) by developing new models and integrating from the beginning the precise properties of a composite into the design process at the stamping stage.
After that the last question for the automotive market is the right moment to launch this technology as a massive market. For the moment some companies are already quite advanced in the use of composite and even some of them have already designed a full car body in composites. For being price competitive such technologies must be spread on a massive market and the reduction of weight and by consequence oil consumption must be accelerated by a high price of oil.

6.2 Weather, Climate and Solid Earth Systems

A consolidated view of the climate experts on disruptive technologies was discussed under the IS-ENES Task Force and then sent to WG 3.5. It was later agreed by WCES, with only limited additions and it resumed in the following.

Technical challenges

To tackle the scientific grand challenges reported in Sections 4.2.1 and 4.2.2, the following technical aspects need to be carefully considered:

- Emerging HPC technologies are based on many levels of parallelism (different processor types in on system, many core per processor, many threads per core ...)
- Technologies relating to data are not keeping pace with peak performance characteristic of systems: optimization is needed of the slow data flow through the numerous layers, between application layer to hardware layer, which are influencing each other non-linearly in many ways, disadvantageous for performance.
- Adequate volumes for storage are not available: both fast storage for model products while simulations are running, and for later analysis (whether fast or not).
- Energy is an increasing cost burden.
- Programming approaches to overcome these issues is increasingly complex

which poses difficulties in an area where large models, with their own scientific and functional complexity, need to be easily and rapidly understandable and adaptable to scientists. We need therefore many more competent and well-educated scientists in these areas!

- Major portions of application codes had to be rewritten (software refactoring) and algorithms had to be re-engineered for applications to run efficiently and productively on exascale architectures. The effort involved in software refactoring can be substantial and often surpasses the abilities of individual research groups. Then, deep synergies between climate and computer scientists need to be achieved, planning community projects to reach a breakthrough in their collaboration.

Possible positive disruptions

The dominant climate models will be those that can run at high resolution and increasing complexity if this can be achieved. The characteristics of these models would be:

- New more scalable dynamical cores,
- Asynchronous IO servers, asynchronous communications, asynchronous algorithms,
- Highly optimised kernels,
- Higher number of kernels running in parallel (i.e. functional parallelism and instruction-level parallelism -SIMD-isation or vectorisation- to overcome exhausted data parallelism),
- Increased overlap between communication and parallelism.

Otherwise, there will be an increasing reliance on ensembles models that should be considered as evolutionary rather than disruptive.
Parallelism in time is theoretically possible and would be disruptive for very long climate simulations, even if it is not likely to be applicable in climate codes for the next decade or so.
The advent of active storage systems may be disruptive in a positive sense if such storage systems can be optimised for climate simulation data. It may include heterogeneity (in memory: local vs. remote, local vs. remote accelerators; in processors: accelerators, heterogeneous cores, etc.).
In any case, better techniques for parallel data analysis, an optimised hardware will be key to exploiting future data volumes.
The issues on “highly optimized kernels” and on “higher number of kernels running in parallel” are included in recommendation on “Development of ultra scalable algorithms with quantifiable performance for realistic apps “ (see below). The one on “dynamical ores” is clearly a challenge concerning weather and climate modelling, and should be addressed by international programs (involving joint development and/or inter comparisons). The issue on “parallelization in time” is seen as theoretically possible, even though not likely to be applicable for the next decade or so.
The issues on “asynchronous IO servers, asynchronous communications, asynchronous algorithms“, on „active data storage systems” and on „better techniques for parallel data analysis” are included in recommendation on “End-to-end techniques for efficient I/O and data analysis”. The issues on “increased overlap between communication and parallelism” and on “functional parallelism” are of more general nature.

Some considerations about possible negative disruptions

Because the climate problems is fundamentally interconnected and data intensive, the trends to massively parallel systems with relatively poor data bandwidth is a disruptive technology in the HPC

market that will negatively impact the climate community. This would lead to increasingly complex coding which would in turn slow down the scientific development of increasing complex models.
New forms of scientific compromises will need to be made due to the characteristics of future architectures. For example, more scalable grids can have undesirable numerical characteristic and the need for parallelism may imply the need to use less than optimal ordering of calculations to make more algorithms independent when they are not; this may break rules of the physics we are trying to model.

6.3 Disruptive technologies

Based on an extensive survey performed by WG3.5 during the first month of activity following first sets of 8 disruptive applications have been identified.

6.3.1 Simulation and Optimization with uncertainties

Many applications rely on uncertain data. Though this is an almost universal statement, it is particularly true for modern applications of simulation technology in fields such as bio-medicine and the earth and life sciences. Other complex systems, such as energy of traffic networks and user behavior do require stochastic descriptions, Thus modeling uncertainties and stochastic behaviour in the models will have large impact on simulation based science in the future. Quantifying uncertainties typically generates increased computational load and many of the important applications will only come into reach with exascale capabilities.

6.3.2 Simulation and Optimization in Complex Networks

Our society and modern economics is increasingly influenced by and relies on complex networks. This includes physical networks such as for traffic and logistics, energy networks, and information networks. Modeling, simulating, and optimizing such networks is rapidly gaining importance. Other applications require the analysis of such networks with the goal of extracting meta-information, The page-rank problem for web-search and classifying web pages is just one example. Other networks, such as the power grid and its development into a “smart” grid require highly complex models in which continuous information (such as voltage and power, or also price) are combined with discrete information (such as the operation of a consumer), and stochastic information (such as consumer behaviour or weather conditions). For such applications, new modelling and simulation paradigms must be developed. For many relevant networks, the size of the data volume can only be handled with HPC Technology. The impact and success of such networks requires novel design and development tools that must be based on a yet to develop exascale computational analysis methodology.

6.3.3 Data Driven Simulations

Many real world scenarios hare characterized by a situation where models and experimental measurements are simultaneously available, but where both the model, and the data are incomplete. Typical examples are in biomedicine, where e.g. models for biological tissue lack the correct material parameters, while medical imaging delivers partial information about the deformation in a specific load case. Thus neither the classical forward simulation based on a classical deterministic model, nor just the analysis of the image data is sufficient. Today, in such a situation, the unknown material parameters may be tuned manually in an ad-hoc fashion so that simulation results agree with image data. However, in the future, we expect that systematic approaches will be developed to combine the techniques of data assimilation on one side, forward simulation on the other side, and techniques for solving inverse problems as a systematic means for bridging both model knowledge and experimental data together. We expect that such hybrid approaches, where both conventional simulation models and data analysis techniques are combined, may have a disruptive impact in many fields where currently simulation technology is not yet fully adopted.

6.3.4 Bridging scales by ultra-large simulations

In simulation technology we traditionally distinguish between the atomistic scale, where atoms and molecules are described and a macroscopic scale, in which we can describe materials or fluids on a continuum level. In between is a so-called mesoscopic scale that is characterized e.g. by the size of a biological cell, a pore in the soil of an aquifer, or the grain the structure of a material. This classification is over simplistic and must be extended by even smaller scales (the quantum level) and larger scales (of geo-science and astrophysics).

From the perspective of simulation technology, it is impossible to simulate most manmade artifacts (i.e. processes and devices that belong to macroscopic level) using micro-scale atomistic simulation. This is simply so, since the number of atoms or molecules needed on the macro-scale typically exceeds anything that can be handled by the computing power that will be at our disposal in any foreseeable future. However, it will increasingly become possible to bridge between the mesoscopic scale and the continuum scale. We note that the number of blood cells in the human body is estimated as 2.5x1013 red blood cells containe din 5l of blood.

ExaScale computing will deliver memory and compute capacity for such large ensembles. In fact, with 1018 operations per unknown, an exascale system will deliver 4x104 Flops per red blood cell per second. This may open the door for completely new approaches to simulation of complex systems, where e.g. macroscopic systems can be modeled based on a mesoscopic resolution. If this leads to yet unforeseen new applications and scientific results, then ultra-large scale applications will a disruptive effect.

6.3.5 Advanced Adaptive Resolution

Adaptive schemes are well known in grid-based simulations, where adaptive meshes are constructed to control the accuracy of solutions on a grid, according to a given criteria for a local error measure. This is well established for unstructured grids in the context of finite element methods and multigrid techniques but also applies to finite volume or finite difference methods, most often applied on structured grids. For complex simulation scenarios, including particles, adaptive resolution schemes are a rather recent development and were proposed in various contexts to form a link to multiscale simulations. Especially for expensive descriptions of particle interactions, adaptive resolution schemes offer a large potential for efficient scaling in space and time. Often, the interesting regions in a simulation are rather localized and a majority of particles serves as environment which acts as a heat bath or elastic medium in cases e.g. of chemical reactions or propagation of shock waves through a material. Adaptive resolution schemes therefore require a change in the description of the nature of particles or an adaptive coarse graining including a smooth transition from one description to the other. For dynamical simulation like molecular dynamics, this requires careful consideration of stability constraints and continuous transitions of thermodynamical, mechanical and transport properties. Progress has recently be obtained by constructing adaptive resolution schemes based on a Hamiltonian description, ensuring consistency with mechanical properties including a smooth transition of energies and forces between regions. Potential advantages of an adaptive resolution description of a system can be recast to a reduced memory consumption and smaller time-to-solution. However, until now efficiency is limited due to the overhead in the administration of the different description in spatial regions, which have to be matched by transition regions, in which a weighting between fine- and coarse-description is done. From the point of view of modeling this methodology is challenging due to the stability constraints concerning dynamical and structural properties. Although promising, this approach was not yet adopted in massively parallel simulations which might partly recast to an involved load-balancing requirement. In the context of exascale this approach is particularly challenging due to the dynamical nature of the problem. Since regions of interest are often time-dependent and in general may vary in size and number, a runtime system is required that controls both the local workload and memory consumption on a compute node and which redistributes load dynamically over the nodes. Since memory is a critical issue on future architectures, load-balancing has not only to consider workload and communication efforts but also memory usage.

Adaptive resolution techniques have a large potential for a wide range of applications but call for requirements in accuracy, robustness and efficiency. Further novel developments lie in coupling particle based methods to mesoscopic methods (such as the Lattice Boltzmann method), or to macroscospic methods based on models from continuum mechanics. Such coupled multi-scale simulations are currently an academically interesting research topic whose systematic development in concrete application scenarios may have disruptive impact in a number of fields, such as mircrofluidics, and micro-reactors, bio-medicine, nano-techology, lab-on-a-chip technology, fracture mechanics, aging and degradation of materials, etc.

6.3.6 Learning-on-the-fly

Simulations of complex systems strongly depend on the level of accuracy of description of individual components and their interactions. Often a proper account of all details of a process has to change the level of description as it is the case e.g. for chemical reactions or activated processes, where individual events cover only a small part of the total simulation time. After a transient period, a sub-system, where an event occurred, often can be described again with a simpler model or a coarse grain description. However, it is not evident that such a description is available a priori and learn-on-the-fly methods aim at providing sufficient information to continue simulations with a description of system components which is evaluated iteratively during the simulation or by a parameter optimization on-the-fly matching the description between a coarse and a fine level. The idea of learn-on-the-fly is thereby to continuously run the simulation and to ensure a proper and stable description of the system. Learn-on-the-fly methods have proven to be a valid method in the context of molecular dynamics simulation for the description of e.g. crack propagation in solids, where classical force-fields are adjusted and improved during the simulation to allow for bond-breaking and a proper description of non-equilibrium dynamics. Similar to adaptive resolution schemes, the description of system components is changed during the simulation. Components are characterized by a multi-parameter description, which allows for an efficient adjustment from an expensive to a more cheap description of components. Learn-on-the-fly methods can be used to adaptively change the level of resolution in e.g. particle simulations or to change a mesh resolution when upscaling in e.g. reservoir simulations. Therefore, learn-on-the-fly techniques offer the great potential to extend the range of applicability of cheap descriptions of system components by intermediate fine-level calculations and adjustments of parameters in a predictor-corrector way. Similar to adaptive resolution methods, learn-on-the-fly calls for a highly dynamic adjustment of the workload on the compute nodes.

6.3.7 Multiscale: hybrid and beyond

A large class of phenomena exhibit processes on various length- and time-scales, which naturally introduces problems to methods which are bound by stability constraints to relatively small time-steps in a simulation. Multiscale simulation approaches can be classified into concurrent- and sequential-methods. In the concurrent approach, various levels of resolution are active in the same simulation applied to different regions requiring a fine-level or allowing a coarse-level description. This type of multiscale approach is often applied to (sub-) systems, embedded into a material, which can be characterized by effective parameters, acting as an environment which provides natural boundary conditions for a system and minimizes finite size artifacts. In solid state physics this is often modeled by an atomistic description inside a given domain, which is matched to a finite element description in an outside region. Depending on the complexity of processes inside the atomistic region, different levels of accuracy might be introduced in a hierarchical fashion, similar to the adaptive resolution scheme. Sequential multiscale on the other side relies on the simulation of properties evaluated in a unique description within a given level of accuracy and combines simulations of different levels of accuracy to feed in parameters from one level to the other. Usually an upscaling scheme is applied where from fine-level simulations with high accuracies, system parameters are evaluated which serve as materials properties on a coarser level. In solid-state physics, for example, elastic constants can be calculated from relatively small systems, which can then be introduced into coarse-level descriptions like dislocation-dynamics simulations. Sequential multiscale simulations therefore stay on the same level of approximation model and therefore have a unique time- and length-scale, which changes during the sequence of sequential simulations on different levels. Concurrent multiscale challenges the simulation by introducing different levels of description at the same time, therefore covering different characteristic length and time-scales. This more complicated way of simulation offers the advantage to cover non-equilibrium scenarios over different scales and to include a direct feedback between fine-to-coarse and coarse-to-fine levels.

Multiscale methodology offers a large potential for future simulation scenarios, especially in the perspective of exaflop computing. The total problem is usually split into parts, which are calculated on a different level of resolution, applying different methods and algorithms and eventually by coupling different programs dynamically. This requires either a module- or library-based concept from the point of view of multiscale program environment, including e.g. ab initio-, classical force-field-, grid-based mesoscopic or macroscopic methods. Working on their own partitions in a parallel environment and coupling software components by interfaces, allows to execute individual methods on smaller parts of a large parallel machine. Since highly accurate calculations, e.g. ab initio, on the finest level, is often related to a large memory consumption and high computational work per degree of freedom, multiscale simulations have the potential to find a balance between workload, memory consumption and complexity of the problem, so that not every software component of a code needs to fulfill hyper scaling properties, which might allow to apply even parts of existing programs, which could not be mapped as monolithic components to massively parallel environment.

6.3.8 In Silico Design of New Materials

The advancement of computationally efficient and accurate atomistic descriptions of matter, offers the possibility to systematically screen compositions of materials and composites. Defining a target function for a desired property the task is related to modify the compositions of an input set of particles, to calculate physical and mechanical properties of the system and to compare with a target function. This can be done in a systematic way by creating data bases of properties for materials and to use chem-informatics and materials-informatics methods to calculate properties in multi-dimensional phase diagrams for mixtures, as it is done e.g. in the CALPHAD[CAL] projects. Large effort in this direction is reported from the US based project Materials Genome Initiative[MGI1] which aims e.g. to find substitutes for critical minerals or designing light weight materials for reducing oil dependence for transportation, thereby facing important societal problems related to the economical use of finite resources. Another closely related US project is the Materials Project [MP1] providing open access data and software to the scientific community and offering Web based applications for e.g. searching materials information, chemical reaction enthalpies, calculating phase diagrams or predicting structures of compounds. The realization of in-silico design of new materials related to exascale compute resources has the large potential to come up with optimized materials for given applications as well as to dramatically reduce time-to-market, which is a critical issue not only for economical competitiveness but also for societal needs, i.e. to provide timely optimal solutions to problems and to keep track with societal, environmental and economical needs. Although this approach is a large step forward to find and optimise new material compounds, in many cases it is still limited. The basic assumption is often that an accurate calculation based on small samples can be considered as representative for bulk materials properties. Often, however, bulk properties are strongly influenced or dominated by the microstructure of a material, which cannot be predicted from very small sample calculations. This applies e.g. to metallurgy and metallic alloys like steel exhibit very rich phase diagrams having their origins in different compositions of iron plus alloying elements, of which the most prominent is carbon, having a weight content of 0.002% to 2.1%. Taking into account the less abundant alloying elements Mg, P, S, Si, O, Ni and Al plus even smaller concentrations of other elements show that current compute resources will not be capable to perform in-silico design of alloys and complex compounds. For these cases a combination of extreme scale compute resources and a further development of multi-resolution algorithms is necessary to open the field for materials design which can take into account the important properties of the microstructure, where some advances were already achieved [MPIE]. For the case of materials design, a strong demand to high throughput calculations is given. This fulfills the need for relatively small sized calculations for a large number of different compositions of materials. However, considering more complex scenarios, taking into account atomistic information plus the microstructure, a combination of capacity and capability computing is required, which calls for exaflop compute resources.

7. Gap analysis

7.1 Industrial and Engineering applications

7.1.1 Recent breakthroughs in Aeronautics/Aerospace

In US

In the aero-acoustics domain, one of the major breakthrough has been performed in early 2013 by a team of the Center for Turbulence Research (Stanford) who has set a new record in computational science by successfully using a supercomputer with more than one million computing cores to solve a complex fluid dynamics problem—the prediction of noise generated by a supersonic jet engine.

Joseph Nichols, a research associate in the center, worked on the newly installed Sequoia IBM BlueGene/Q system (16.3 PFlops peak) at Lawrence Livermore National Laboratories (LLNL) funded by the Advanced Simulation and Computing (ASC) Program of the National Nuclear Security Administration (NNSA).

The exhausts of high-performance aircraft at takeoff and landing are among the most powerful human-made sources of noise. For ground crews, even for those wearing the most advanced hearing protection available, this creates an acoustically hazardous environment. To the communities surrounding airports, such noise is a major annoyance and a drag on property values.

Understandably, engineers are keen to design new and better aircraft engines that are quieter than their predecessors. New nozzle shapes, for instance, can reduce jet noise at its source, resulting in quieter aircraft.

Predictive simulations—advanced computer models—aid in such designs. These complex simulations allow scientists to peer inside and measure processes occurring within the harsh exhaust environment that is otherwise inaccessible to experimental equipment. The data gleaned from these simulations are driving computation-based scientific discovery as researchers uncover the physics of noise.

For this massive simulation, the team from CTR used an in-house LES code code called CharLES on the whole 1 572 864 cores configuration of Sequoia.

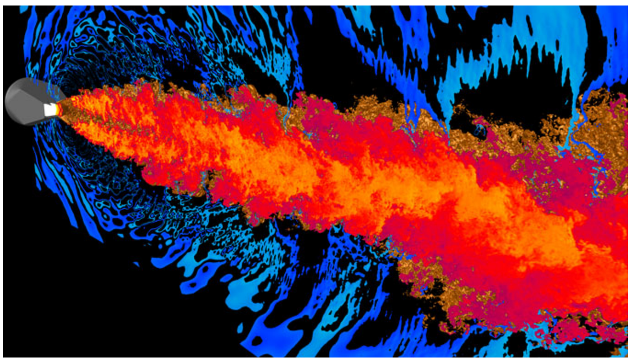

An image from the jet noise simulation. A new design for an engine nozzle is shown in gray at left. *Exhaust tempertures are in red/orange. The sound field is blue/cyan. Chevrons along the nozzle rim enhance turbulent mixing to reduce noise. (Illustration: Courtesy of the Center for Turbulence Research, Stanford University)*

Beyond this very impressive result performed by this team from CTR, there is also the need in Europe to deploy and sustain such amount of computing performance into PRACE in order to keep European scientific teams at the same level of innovation.

In Europe

With the raise of PRACE petaflops machines available since mid 2010 some European teams performed also some major breakthroughs in CFD especially in the simulation of turbulent flow using DNS methods.

At University of Aachen the team of H. Pitsch worked on a massive simulation of highly resolved isotropic forced turbulent flows applied on a grid of 4096^3 (or 68.7 billion) points on the JUQUEEN machine at Jülich Research Center (Germany) in January 2013. This simulation performed using the pDNS code developed by Pitsch 's team for pseudo-spectral DNS (3D FFT for spatial discretization) ran on up to 524,288 threads on this IBM BG/Q system (the largest HPC system in Europe at date).

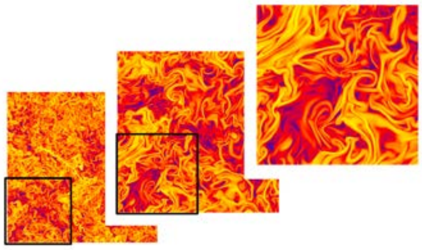

In the field of energy, gas turbines are a crucial contributor to the world's power capacity. In the future, the role of gas turbines in electricity generation in Germany will become even more important, as gas turbine power plants have been identified as a potential replacement for nuclear power plants. The vision "Energiekonzept 2050" by the German government emphasizes the importance of stationary gas turbines as it aims to increase the share of renewable energy. Therefore, it is crucial to reduce pollutant emissions from future gas turbines, which are operated under lean premixed conditions.

Same team conducted also very innovative simulations on temporally evolving premixed jet flame with finite chemistry (DNS) in the field of Large-Scale Simulations and Modeling of Pollutant Emissions in Turbulent Premixed Flames. Simulations on up to 70 000 cores of the SuperMUC system at LRZ (Germany) were performed on massive 10 billion grid cells meshes and showed the strong weight of the chemistry into the total simulation time. This could lead to a hybrid version of the code where chemistry could be executed on GPUs or manycore accelerators while the rest of the code will remain on scalar system.

7.1.2 New US roadmap elements in the Aeronautics/Aerospace

In August 2012, NASA organised in Hampton, Virginia a full seminar dedicated to the presentation and exchanges between US aeronautics and aerospace experts of the NASA 2030 Vision.

As all the proceeding are online, this give us the opportunity to have a look on some HPC recommendations of the panellists:

- The need to work with more accurate models for hybrid turbulence
- The assessment of uncertainties will also play an important role in the future and this will require the support of massive parametric studies
- Important to develop efficient and robust mesh adaptation for complex configurations
- The experts are considering that the use of accelerators will be a key for achieving performance for capability codes as well as less scalable capacity codes.

And some important factors of success regarding 2030 CFD:

- Throughput: End-to-end cycle time must be faster than process based on wind tunnel testing ☐ Consistent quick turnaround essential in a schedule-driven development program ☐ Engineering development consists of development of large databases for loads, stability and control, and flight simulation

- Cost: Database creation using CFD must be cheaper than using wind tunnels
- Accuracy: Must be superior to wind-tunnel
- Error / Uncertainty Quantifications: Must be able to quantify the error and uncertainty associated with results
- Expertise: Must be automated and operational using project engineers, not just CFD experts

Its interesting to notice that most of the recommendations given by the panellist of this 2030 CFD workshop are already integrated and present into the ones given by the aeronautics/aerospace European experts as well.

7.1.3 Breakthroughs in Automotive

Since mid 2012 the PRACE calls for proposals for accessing to the 6 European petascale systems deployed are now open to industrial users through the Open R&D offer. Companies can apply twice a year to PRACE calls and can be granted for a one year free of charge access on the basis of scientific excellence of their proposal and if they commit on publishing results at the end of the grant period.

At this date almost 12 companies benefited from the PRACE resources and more than half are SMEs.

In the automotive domain its interesting to illustrate as breakthroughs two of the first PRACE Open R&D grants in the field of advanced CFD for HydrOcean and optimization of crash processing for Renault.

The first case is coming from HydrOcean, a spinoff of Ecole Centrale de Nantes, specialized into hydrodynamics simulation using an in-house SPH (smooth particle hydrodynamics) code called SPHFlow. This SME has been of the first companies granted by PRACE and they are now performing massive SPH simulations on more than 32 000 cores on the Hermit system at HRLS (Stuttgart, Germany) for assessing the aquaplaning effect on automotive.

The second example come from Renault (with ESI an ISV in physics of materials for manufacturing industries and Ecole des Mines de St Etienne), one of the European leading car company which has been granted by PRACE with 42 million CPU hours on CURIE at TGCC/GENCI (Bruyères-Le-Chatel, France) for performing massive crash optimization. This unprecedented case in the automotive industry is applied to an optimization study with more than 200 different parameters (more than 4x the state of the art) applied on meshes of more than 20M finite elements (more than 3x the state of the art).

As previous example from HydrOcean was performed in 32 000 cores showing the need to support capability for industrial applications, the research project from Renault will spawn hundreds of individual executions running on 512 and 1024 cores in capacity mode.

Such two examples illustrate the need for research infrastructures like PRACE to support both capability and capacity for supporting industrial needs.

Such level of complexity was not reachable for current Renault HPC facilities and this project will allow Renault to increase the safety of their car and to anticipate the future strongest EuroNCAP regulations coming in 2015.

These two cases illustrate the crucial need for European industries (large companies as well as SMEs) in order to increase their competitiveness to use advanced simulation and HPC and to rely on leading edge sustainable research infrastructure for performing Open Innovation projects with academia.

Finally during the PRACE fifth industrial seminar held in Stuttgart in April 2013, a representative from Porsche made a strong testimony on the use of HPC for the development of their products. HPC represents 0.26% of the development budget and 1.14% of the IT budget but it has a major impact for the company. As an example the use of HPC allowed Porsche to reduce by 16% the time to market of the Panamera and by 28% for the 991 C2.

Massive simulations allowed also to reduce weight of the car by 26 kg on the 991 C2 and Porsche performed more than 15 000 multi disciplinary optimization simulations in 2012 on the Automotive Simulation Center in Stuttgart.

HPC is s strong enabler for crash analysis and all the in-car security systems in order to fit with an EU recommendation imposing to halve the number of overall deaths until 2020.

7.1.4 Breakthroughs in Oil & Gas

One of the most important breakthroughs in the oil & gas industry was the acquisition and the installation by TOTAL of the largest industrial HPC system in their R&D center in Pau (France) on April 2013.

This 2.3 PFlops peak performance SGI Altix system called Pangea has been ranked in #11 position of the latest top500 on June 2013. It will allow TOTAL to develop and use the latest generation of RTM (Reverse Time Migration) and beyond advanced numerical methods. The performance of this machine will be doubled by TOTAL in 2015.

7.1.5 Breakthroughs in Nuclear Energy

As TOTAL for Oil & Gas domain, EDF one of the first European leading player in electricity and energy increased its HPC internal facilities by installing in their R&D center a new 1 PFlops peak performance x86 cluster provided by IBM. This machine will complement the current 4 racks IBM BG/Q (836 TFlops peak) already in production since more than 1 year.

With close to 2 PFlops available, EDF is the leader in its market in term of HPC facilities installed and one of the first industrial players worldwide.

7.2 Weather, Climate and Solid Earth Systems

In this section a preliminary gap analysis on the two main topics of this report is presented. The following table mainly highlights the key climate challenges, their current status and the expectations in 2020. It should later be completed in order to account for other WCES domains than climate.

Challenges	Where we are now	What is expected in 2020

Stand alone model: atmosphere, ocean,	• Run effectively well on o(1000) cores • New dynamical core: work in progress, international cooperation (CMIP exercise every 4 years) • Recognition of the need to separate natural science layer and computer science layer (example: UK Gung ho project)	• Improved physics • Use new dynamical core operationally
Earth system model	Used in CMIP5 like experiments	More complexity, larger ensembles, longer simulations, (see PRACE scientific cases)
Data Analysis and Visualization	Client-side data analysis and visualization, lack of parallel tools for data analysis and visualization, visualization and analysis tools weakly coupled	Parallel data analysis at the data center. Novel visualization tools with parallel engines and better coupling/integration with parallel I/O libraries
Data quality	Current procedures not sufficiently standardized, automated and comprehensive	Programmatic and automated data quality enabled on storage clusters, with extensive range of functions and checks.

7.3 Fundamental Sciences

7.3.1 Report on national and international initiatives

FET Flagship Projects

Flagship projects from the Future Emerging Technologies (FET) are ambitious large-scale, science-driven, research initiatives that aim to provide a vision for future developments. The scientific advance is aimed to provide a broad basis for future technological innovation and economic exploitation in a variety of areas, as well as novel benefits for society.

One of the two flagship projects which has been chosen to be funded with a 10 years budget of 1 billion Euro is Graphene. The consortium consists of 9 partners from 7 countries. From the start in

2013 the Graphene Flagship will coordinate 126 academic and industrial research groups in 17 European countries. Through calls, activities will be broadened towards fundamental and engineering aspects of graphene research. The mission of Graphene is to take graphene and related layered materials from academic laboratories to society, revolutionize multiple industries and create economic growth and new jobs in Europe.

Currently there is no direct link to high performance computing or exascale initiatives in the foreground of the project. However, there are a number of theory groups related to the project, which are involved in computational- and high-performance computing activities, so that computing will be advanced at least as a side effect. First demonstrations of highest scalability on the European PRACE Tier-0 architecture Curie are reported for Graphene related simulations.

FP7 Exascale projects

The three European exascale projects, Mont Blanc, DEEP and CRESTA, have started in October 2011:

- Mont-Blanc: This project is mainly concerned with hardware development for energy efficient computing, as the energy consumption will be one of the most important concerns for reaching exascale performance. The goals of Mont-Blanc are to assess and design a new type of computer architecture capable of setting future global High Performance Computing (HPC) standards that will deliver Exascale performance while using 15 to 30 times less energy than current HPC platforms.
  Although the project by itself is mainly related to set up a new hardware environment, several workpackages are devoted to software as well. WP3 is devoted to kernel selections and WP4 to porting and optimization of larger program packages. Related to the fundamental sciences, represented by the present EESI working group, are the programs BQCD (QCD), BigDFT (Material Science, Computational Quantum Chemistry), EUTEURPE (plasma physics), MP2C (soft matter physics), PEPC (plasma physics), QuantumESPRESSO (materials science, computational quantum chemistry). All programs have proven to scale up to a large number of processors on conventional HPC archtictures. In the Mont-Blanc project kernels and programs were first ported to run on ARM-processors, as this was the original target platform. The purpose of this first porting was not to provide decidable benchmarks for the potential of this architecture. Such features will come with the next generation of Mont-Blanc platforms like PedraForca in which each compute node will be boosted by a mobile GPU and will have a much more performing inter connect. An important extension of the codes, used in the project, is the extension to programming models based on OMPSs, a task based model, which supports to run independent tasks in parallel and to facilitate load-balancing.

- The DEEP (Dynamical Exascale Entry Platform) consortium will develop a novel, Exascale-enabling supercomputing architecture with a matching SW stack and a set of optimized grand-challenge simulation applications. DEEP takes the concept of compute acceleration to a new level: instead of adding accelerator cards to Cluster nodes, an accelerator Cluster, called Booster, will complement a conventional HPC system and increase its compute performance.

Directly related to activities of the current WG is one application, using the TurboRVB program package, performing Quantum Monte Carlo simulations. The aim is to study realistic materials in order to determine whether high-temperature superconducting properties can be quantitatively understood and therefore predicted within an ab-initio approach. The TurboRVB package is developed at SISSA (Scuola Internazionale Superiore di Studi Avanzati di Trieste) and runs on massively parallel architectures. The TurboRVB is well suited for the DEEP Architecture since the Cluster part can be exploited for the optimisation of the many-body wave function, letting the Booster perform the evolution of the Monte Carlo walkers.

- CRESTA brings together four of Europe's leading supercomputing centres, with one of the world's major equipment vendors, two of Europe's leading programming tools providers and six application and problem owners to explore how the exaflop challenge can be met.
  The project has two integrated strands: one focused on enabling a key set of co-design applications for exascale, the other focused on building and exploring appropriate systemware for exascale platforms.
  There are six application program packages foreseen to be prepared for exascale architectures, originating from engineering (CFD, Nek5000, OpenFOAM), life sciences (molecular dynamics, GROMACS), nuclear fusion (gyrokinetics, ELMFIRE), weather/climate (weather prediction, IFS (Integrated Forecast System)), physiology (virtual human, HemeLB), i.e. there is one package related to the present WG 3.3, i.e. related to nuclear fusion energy-production. Compared with other programs it seems, however, that ELMFIRE still needs some improvements, e.g. extension to domain decomposition, to reach hyperscaling. Progress in this direction will be tackled within CRESTA.

G8 projects

International initiatives including participation of EU, US and Japan are seen in the G8 projects (see e.g. [G8web]), where projects from weather/climate, earth sciences, biophysics and physics are funded to prepare codes, algorithms and workflows for architectures of highest performance. One of the six projects, NuFuSE, is addressing Nuclear Fusion, demonstrating porting, optimization and highest scalability for a number of European and international codes.

Progress was obtained by improving data locality in PIC codes using e.g. data sorting of particles for improved memory access on a grid. Current activities include highly multi-threaded algorithms to address the memory latency as well as development of GPU and hybrid versions of codes, invoking MPI-OpenMP models in domain decomposition. A materials physics part of that project works on simulation techniques for rare events, including temperature accelerated dynamics, parallel replica dynamics and hyper-dynamics. These types of simulation techniques work on a number of replica systems and may be considered to scale very well on very big machines, as individual simulations might run efficiently on small partitions which, from time to time, exchange data between each other to progress the system over larger time scales.

High Performance Computing Projects in Germany

BMBF

Since 2009 there were, under the research program "IKT 2020 - Research for Innovations", overall three subsequent calls for “HPC-software for scalable parallel computers” [HPC-1,HPC-2]. Within each of the calls about 12 german-wide network projects were supported in diverse areas of research of numerics, system- or application-software. Although some projects contained a link to real world applications, the idea was not to promote specific applications and to solve grand challenge problems, but to prepare a software stack for applications which could also be used by industrial partners. Another goal of this initiative was to form a German HPC community, which was fostered by common meetings and conferences.

SPPEXA

SPPEXA is another German initiative to foster high performance computing and to form a community for HPC in Germany. Although the name suggests a close relation to exascale computing, this initiative works in a different way than e.g. FP7 exascale projects. SPPEXA is mainly driven by computer scientists, physicists and mathematicians, and focuses more on the algorithmic and software engineering aspects of implementations. Highest efficiency is aimed by optimizing software for current available HPC architectures and first impressive results were reported for e.g. the Earth mantle dynamics based on a multiscale space-time approximation, built on modern finite element technology and communication-avoiding ultra-scalable multigrid. Porting to the Blue Gene/Q in Jülich (JUQUEEN) was achieved onto 393,216 cores, running 983,040 parallel threads [TERRA].

Thematically close to the present WG 3.3 are the projects EXASTEEL, focussing on scale-bridging methods for multiphase steels [EXAST], EXAHD, working on two-level sparse grid methods for plasma physics applications [EXAHD] and EXAMAG, preparing simulations for highest performance in the field of evolution of the universe including magnetic fields.

Community activity through CECAM

The Centre Europeen de Calcul Atomique et Moleculaire is traditionally organizing workshops and tutorials in the field of atomistic simulations for e.g. materials science, soft matter, biophysics and fosters methodological and algorithmic research in scale bridging techniques. A strong interest of CECAM is the formation of a scientific community in Europe and to develop an infrastructure for information and knowledge exchange in the field of complex computer simulations. Ongoing initiatives foster the contact to industry to bring academic research closer to industrial needs on the one side but also to demonstrate to industry the potential benefit from cooperation with academia. Since recently not only workshops and tutorials are supported by CECAM but also discussion groups and working groups are organized focussing on specific topics. Noteworthy in the present context is the Scientific Software Development Group which consists of international members, discussing the actual needs and fostering research in direction of e.g. code interfacing, data basis for molecular simulations, verification of methods and codes.

7.4 Life Sciences and Health

Some Life Sciences projects are facing problems involving massive management of data in processes requiring little inter-processor communication. These problems might require hundreds of thousands of processors, but without requirement for fast interconnection between processors. The vast amount of data to be managed hampers however the use of cloud or GRID initiatives as general solution to the problem, since data movement among internet can become rate-limiting step. Suitable and flexible access to computer resources is crucial in this area.

Many Life Sciences areas are on the stage of collecting data. In a short period of time Life Sciences will converge to biomedicine simulations and molecular simulations, defining then a complex multi-scale scenario that might require Exascale capabilities and a specific simulation framework.

ExaFlops capabilities will allow simulating larger bio-molecular systems, but probably not for long enough time scales, since no infinite scaling is possible in systems of finite size. We envisage that Exascale capability will power biased-sampling techniques, which would require parallel computing. We believe also that Exascale capability will make possible today impossible “in silico” experiments, for example in proteome-scale screening of chemical libraries to find new drugs or in the study of entire organelles and even cells at the molecular level.

Simulation of a cell, the human brain or the whole human will require Exaflop computing, but this capability will not be sufficient, since integration of experimental information, ability human-interaction with calculations and refinement of underlying physical models is also instrumental for success.

Breakthroughs in molecular simulation

On June 2013, University of Illinois physics professor Klaus Schulten and postdoctoral researcher Juan Perilla recently used molecular simulations on the NSF Blue Waters supercomputer to determine the chemical structure of the HIV capsid. The capsid protects the virus’s genetic material and plays a key role in debilitating the immune system. However, although the capsid is an attractive target for antiretroviral drugs, researchers have struggled to understand its structure.

Researchers have used various lab techniques to look at individual parts of the capsid. (In fact, the Blue Waters simulations integrate data from experiments performed at the University of Pittsburgh, Pennsylvania, and Vanderbilt University, Tennessee.) Until now, they were unable to simulate the entire capsid, which consists of more than 1,300 identical proteins at the atomic level.

Simulations on the structure of the HIV capsid use highly scalable NAnoscale Molecular Dynamics (NAMD) software, developed at the University of Illinois by Schulten and computer science professor Laxmikant ’Sanjay’ Kale. NAMD's parallel scalability is based on the prioritized message-driven execution capabilities of the Charm++/Converse parallel runtime system. These efficiencies enable far better simulations: in 2006, the first million-atom simulation fully characterized the structure of the tiny satellite tobacco mosaic virus. Running on a petascale computer, NAMD now allows biomolecular modelling like Schulten’s and Perilla’s 64 million-atom simulation of the full HIV capsid.

Such capabilities are key to developing potent new antiretroviral drugs that suppress the HIV virus and stop the progression of AIDS. Armed with the detailed capsid structure, researchers can investigate how to disrupt its functioning – effectively interfering with the timing of its opening to prevent HIV infection.

8. Views concerning the 12 general recommendations

We give below some preliminary comments related to the 12 general recommendations prepared during the EESI-2 Technical meeting (Le Tremblay, France, May 28-29, 2013) and indicate how they may apply to the domains covered by the industrial and engineering applications working group.

8.1 Training and education from HPC to Exascale

This subject is one of the first issue raised by the working groups inside its first EESI1 report, training and retaining people will be crucial for European research laboratories and industries and this deserve a special and ambitious European programme from education to permanent training.

PRACE just opened its training programme to industry in 2012 through the Open R&D access, European companies can now be trained into the 6 PRACE Advanced Training Centers (in France, Germany, UK, Italy, Finland and Spain) on programming languages, parallelistion and optimization techniques, use of petascale systems, pre and post processing as well as specific open source software and pertinent methodologies for industry like an upcoming training on uncertainties.

As already done in France, it would be useful to label Ph. D. thesis during which the student has developed a clear competence and skills on HPC and Tier-0 machines. Such a label would possibly help the new doctors to find a position within academics or industry at a time when these skills are more and more needed.

Some other recommendations from WP3 will be to:

- fund co-educational programs on the master level which provides a broad knowledge of hardware technology and system- and application software not only for computer scientists but also for students in scientific computing.
- Foster educational programs for training of postdoctoral and senior scientists, who have to follow the trends and developments of hardware architectures, software and optimization and tuning of codes.

These actions need to be amplified by other physical training sessions and completed by the use of (massive) online tools for training remotely, this is becoming very popular in the US with the rise of specialized services provided by academia as well as new startups performing MOOC (Massive Open Online Courses).

8.2 Big Data

All the WP3 Applications working groups are very sensitive to the issue of handling huge amount of data generated by large instruments (telescopes, sequencers, particle accelerators, network of seismic sensors in oil & gas, inflight data from an airplane, …) or massive 3D simulations in both capacity and capability models. This amount of data must be managed in massively parallel systems, requiring not only computation but also I/O parallelization and smart intermediate storage management and methodologies.

The increase in amount of data and the efficient analysis and interpretation of data is an important task in scientific computing. When increasing the computational performance by three orders of magnitude when entering the exascale era with respect to the current state, the storage and processing of data becomes a major concern. Especially in high energy physics, cosmology, life sciences and climatology where large experiments or massive/ensemble simulations will increase, which produce a vast amount of data to be processed, high throughput of data has to be achieved combined with intelligent tools to filter interesting events.

The speed of data generation in simulations is continuously increasing compared to data writing onto or reading back from disk due to an increasing gap between computational performance and I/O-bandwidth. This calls for in-situ methodology for data-reduction, on-the-fly data analysis and data-filtering to automatically detect hot spots. This has to be considered together with asynchronous work flows taking into account a balance between data generation, data-reduction, data analysis and reduced data I/O, scalable on massively parallel architectures and to be demonstrated in complex simulation environments.

The WP3 experts have been involved with WP4 into the overall EESI2 recommendation regarding Big Data.

8.3 Towards next generation couplers and associated methodologies

As stated into the combustion section (4.1.5), the WP3 working groups considers as crucial the need to develop next generation of couplers and associated methodologies for managing multi scale and multi physics simulations requiring exascale facilities.

Different programming languages, data structures, parallelization strategies or scalability of software components makes it often hard to achieve a horizontal integration of codes in standardized way. With respect to concurrent multiscale simulations, which require different levels of detail in the description of a physical system, the requirement for integrating existing software, which is tested and verified, becomes a strong need. The same is true for new codes which rely on functional units and could be included as modules or libraries.

The WP3 working groups have been active into the promotion of such an EESI2 global recommendation.

8.4 UQ & Optimization methods and tools, improving code VV UQ (Verification, Validation and Uncertainly Quantification)

This domain is also very important for the industrial and engineering working group and UQ has been already stated as one of the important field to invest in order to cover all the potential usage of numerical simulation and HPC. Uncertainties will come from input data, from models, from numerical simulation as well as output data and need to be qualified in order to ensure robustness and assessing associated risks. Beyond the traditional Monte-Carlo sampling, a number of uncertainty propagation and probabilistic simulation algorithms have been and have to be developed, such as accelerated sampling (importance sampling, particulate methods etc.), reliability techniques, stochastic developments (e.g. chaos polynomials) and response surface techniques, ... The coupling between stochastic and deterministic methods is also a key for solving stochastic optimization problems.

Earth-system and climate modeling are faced with uncertainty quantification (the fluid mechanics part of these models is subject of deterministic chaos). Up-to-now this issue has been addressed through Monte-Carlo simulations, but the number of control parameters (initial conditions, model parameters …) is definitely much too large to adopt a simple method of varying a few of them. The numerical weather prediction (NWP) models are using methods in which the individual simulations correspond to different initial conditions, each of them corresponding to a perturbation associated with a fast-developing linear instability. Methods to account for uncertainty qualification in models with a very large number of degrees of freedom would definitely benefit to WCES domains.

These 2 working groups have been active into the promotion of such a EESI2 global recommendation.

Same concerns apply to optimization frameworks.

8.5 Exascale software simulators and mini apps

This recommendation has not been discussed among the working group but there is a possibility for companies who are developing their own software to work on extracting mini apps from a representative usage. In the field of advanced seismic methods, a company like TOTAL already worked on a mini app coming from their RTM software in order to work with the ECR (Exascale Computing Research) Intel laboratory for co designing such application on manycore devices.

In the WG3.2, It should however be stressed that the community is very much interested in developing benchmarks based on coupled model (e.g., atmosphere-ocean), as up-to –now benchmarking in WCES domains rely only on either atmosphere or ocean models without any coupling between.

A possible data driven mini-application demonstrating the new usage model in the climate domain is related to the analysis of climate indicators from distributed (e.g. CMIP5) petabyte data archive (existing worldwide climate data node federation)

8.6 Development of ultra-scalable algorithms with quantifiable performance for realistic apps

This recommendation is one of the top issues raised by the most of the WP3 working groups. This includes, e.g., numerical methods for solving more complex wave equation formulations, multiphase flows, adaptive methods for heterogeneous platforms, hybrid solvers, advanced numerical acceleration techniques, wavelets, multi-grid, better and simpler pre-conditioners are key issues that have to be addressed.

8.7 End-to-end techniques for efficient I/O and data analysis

This concerns is the same as the one regarding Big Data.

8.8 High productivity programming models, Dynamic Data Structure

The experts from the WP3 working groups stressed the fact that its important to develop novel programming models which will be able to manage and hide the underlying complexity of the Exascale platforms but such development must lead to the emergence of standards. No industrial will commit on porting real production codes on novel languages if such languages are not perene and supported in time.

This lead to the following recommendations:

- Provision of programming environments, which abstract the logic of applications from actual resources. Embedded Domain Specific Languages (DSL) can communicate between the programmer and the specific hardware, also enabling rapid prototyping.
- Enable exploitation of the heterogeneity of parallel platforms from runtimes that are able to exploit aspects such as data locality and load balancing

8.9 Software engineering methodologies

There is an important gap between codes developed by academia and used in lowest TRL by industry for research and production codes with high levels of TRL used by industry. The challenge for academia and industry is how to industrialise such developments done in an Open Innovation contest and make them usable, supported in time by industries.

For the European Commission as well as the others funding agencies there is a need to invest on structures/teams which are able to provide software engineering methodologies for developing new standard components (numerical libraries, i/o libraries and tools, data managements tools, pre/post processing packages, …) that could be used widely in a long tem period by various communities. Such example of shared components used daily by multiples communities is for example the PETSc numerical library funded by DoE in US with a strong team for developing, testing and supporting the tool.

Developed at the European level in order to reach the needed critical mass and to gather a distributed expertise, such common components will avoid to European scientists and industries to re-invent the wheel every time and could allow Open Source software to cohabit with ISV commercial software by using the same underlying framework.

Along the lines with the need to couple different software components is the need to provide basic routines, functions, modules, libraries, which possess a level of (nearly) optimal implementation for a given hardware platform. Therefore, various implementations of the same functionality might exist for different hardware, which would be made available for the scientific community. The Open Source model, often applied in academia might be one way to offer functional components to the community. Another way is via Open Innovation, which also includes sharing of commercial products and which is practiced between a lot of different companies in order to share ideas and products to increase its own productivity and to stimulate inventions.

As additional recommandations:

- Define standards for sharing and (re-)using software including commercial parts.
- Create code repositories which offer numerical kernels, mini-apps, modules and numerical libraries, optimised for different architectures

Applications and models have typical lifetime of 10 year or more. They are developed either in the academic context, where a number of developments are done by students and scientists who are not permanent members of the models teams, or in the operational meteorology context under very strict developing rules and quality control. General methodology for software engineering would be helpful.

8.10 System software that adapt to resource variations

The WP3 working groups have a general interest for such type of software and more generally to smart runtime systems able to perform dynamic load balancing but has nothing particular to contribute at this stage.

8.11 New consistent (across software layers) resilience approaches

Resilience has been identified are a key challenge for the upcoming Exascale architectures with potential millions of hardware and software components. A multi level software approach for handling hardware and software errors could be a smart approach in order to tackle this problem, from the level of the OS, to the runtime and then the final application itself.

WG3.2 has a general interest for such type of approaches and first exploratory studies have started (see Maisonnave [NEMO2012]) even if nothing is published and no particular methods are proposed at this stage.

8.12 Simulator investigating and deciding which direction to take

Same remark as the previous one for Exascale software simulators and mini apps

9. Assessment of Centre of Excellence

Following the outcome of EESI1 report regarding industrial applications, some early discussions started among the working group toward the possibility to gather strong European expertise in CFD and especially in combustion or turbulence. Similar US initiative is the CTR (Center for Turbulent Research) established since the 80s by the Stanford University and NASA, a research consortium for fundamental study of turbulent flows.

The principal objective of the CTR is to stimulate significant advances in the physical understanding of turbulence
and related non-linear multi-scale phenomena. These advances are directed to improving capabilities for control of turbulence and to modelling turbulence for engineering analysis. Particular emphasis is placed on probing turbulent flow fields developed by direct numerical simulations and/or laboratory experiments, on using new diagnostic techniques and mathematical methods, on concepts for turbulence control and modelling, and on complex effects on turbulence. These effects include, complex geometry, chemical reactions, complex fluids, multi fluid phases, magneto-hydrodynamics and hypersonics.

The main elements of the Center are an extensive visiting Fellows program, a biennial summer program, seminars and workshops, and a core of PhD students and postdoctoral researchers.

Beyond similar activities toward training and fellowship programs such a center of excellence could foster European research and collaboration toward the development in Open Innovation with industry of innovative products.

In the WG3.2, the Earth-system community has organized itself over the past 15 years or more, with successful initiatives in the field of modelling (e.g., CMIP[1] initiatives, IS-ENES and other international modelling workshops), of data management (e.g., IS-ENES, GO-ESSP, ESG/ESGF) and of inter-comparisons of model results (e.g., IPCC[2] assessment reports). The question is now to appreciate what would be the most appropriate organization for this community, either a distributed European centre of excellence or a permanent infrastructure to support its various and numerous far-reaching activities.

The WG3.4 also expressed into their previous EESI-1 report the wish to structure the European life science community into a strong Centre of Excellence targeting personalised medicine with the use of omics or tissue modelling.

Since last year and the announcement from the European Commission of the funding of the Human Brain Project, this will pave the path to the creation of a structure largely beyond the scope of a Centre of Excellence, a flagship which will foster the structuration of an European task force addressing the crucial issue of the modelisation and the understanding of the Human Brain and all the associated outcomes in term of research that will be required in a wide variety of areas including high-performance computing, neuro-morphic computing (emulating brain circuitry to perform complex calculations), brain-machine interfaces and robotics.

---

[1] Coupled Model Intercomparison Project. The CMIP-5 has been completed and the CMIP-6 is in its preparatory phase

[2] International Panel on Climate Change. The 5th assessment report is about to be published

10. References

[PRACE] PRACE Report: “The Scientific Case for High Performance Computing in Europe”, http://www.prace-project.eu/IMG/pdf/prace_-_the_scientific_case_-_full_text_-.pdf

[ISENESWS2013] André, J.C., G. Aloisio, J. Biercamp, R. Buedich, S. Joussaume, B. Lawrence and S. Valcke (2013): “The 2nd IS-ENES Workshop on “High-Performance Computing for Climate Models” (Toulouse, Jan.30th-Feb.1st, 2013)”. Submitted to Bull. Amer. Meteor. Soc.

[NEMO2012] Maisonnave, E., A resiliency strategy for NEMO ocean model, *Technical Report*, TR/CMGC/12/13, CERFACS, Toulouse, France

[G8web] http://www.dfg.de/en/research_funding/international_cooperation/research_collaboration/g8-initiative/faq_g8/index.html

[HPC-1] http://www.gauss-allianz.de/en/projects/1-call

[HPC-2] http://www.gauss-allianz.de/en/projects/2-call

[SPPEXA] http://www.sppexa.de/

[EXAST] http://www.mi.uni-koeln.de/numerik/de/left/research/exasteel/

[TERRA] http://www.fz-juelich.de/ias/jsc/EN/Expertise/High-Q-Club/Terra-Neo/_node.html

[EXAHD] http://ipvs.informatik.uni-stuttgart.de/SGS/EXAHD/

[NviOpt] http://www.theinquirer.net/inquirer/news/2227038/nvidia-says-large-gpgpu-speed-up-claims-were-due-to-bad-original-code

[Luis11] M.Luisier, T.B.Boykin, G.Klimeck and W.Fichtner; ***Atomistic nanoelectronic device engineering with sustained performances up to 1.44 Pflop/s,*** **SC '11 Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, Article No.2.**

[Shi11] T.Shimokawabe et al.; *Peta-scale Phase-Field Simulation for Dendritic Solidification on the Tsubame 2.0 Supercomuter*, SC '11 Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, Article No.3.

[JuQ13] http://www.fz-juelich.de/ias/jsc/EN/Expertise/High-Q-Club/_node.html

[HBP] http://www.humanbrainproject.eu

11. Annex

11.1 List of experts of WG3.1

EESI2 - WG 3.1	**Industrial and Engineering applications**			
Name	Organization	Email	Country	Area of Expertise
Stéphane Requena (Chair)	PRACE/GENCI	stephane.requena@genci.fr	FR	HPC
Ange Caruso (Vice Chair)	EDF	Ange.caruso@edf.fr	FR	Nuclear Energy
Norbert Kroll (Vice Chair)	DLR	Norbert.kroll@dlr.de	DE	CFD, aeronautics
Philippe Ricoux	TOTAL	Philippe.ricoux@total.fr	FR	Oil & Gas
Thierry Poinsot	CERFACS	Thierry.poinsot@cerfacs.fr	FR	Combustion, CFD
Eric Chaput	Airbus	Eric.chaput@airbus.com	FR	Flight Physics, Aeronautics
Heinz Pitsch	University of Aachen	h.pitsch@itv.rwt-aachen.de	DE	Automotive, CFD, combustion
Lee Margetts	NAFEMS / University of Manchester	Lee.margretts@manchester.ac.uk	UK	ISV
Yves Tourbier	Renault	Yves.tourbier@renault.com	FR	Automotive
Vladimir Belsky	Dassault Systèmes	Vladimir.belsky@3ds.com	USA	ISV

11.2 List of experts of WG3.2

EESI2 - WG 3.2	**WCES (Weather, Climatology and solid Earth Sciences)**			
Name	Organization	Email	Country	Area of Expertise
Giovanni Aloisio (Chair)	UNIV. SALENTO & CMCC	giovanni.aloisio@unisalento.it	IT	Exascale Computing
Massimo Cocco	INGV	massimo.cocco@ingv.it	IT	Solid Earth, Seismology
Sandro Fiore	UNIV. SALENTO & CMCC	sandro.fiore@unisalento.it	IT	Scientific Data Management
Pierre-Philippe Mathieu	ESA/ESRIN	pierre.philippe.mathieu@esa.int	IT	Earth Observation Science and Appl.
Jean-Claude André (Co-Chair)	Jca Consultance & Analyse	jc_andre@sfr.fr	FR	Earth-system Modelling
Sylvie Joussaume	IPSL	sylvie.joussaume@lsce.ipsl.fr	FR	Earth-system Modelling
Sophie Valcke	CERFACS	valcke@cerfacs.fr	FR	Coupled Climate Models
Jean-Pierre Vilotte	IPGP	vilotte@ipgp.jussieu.fr	FR	Solid Earth, Seismology
Marie-Alice Foujols	IPSL	marie-alice.foujols@ipsl.jussieu.fr	FR	Earth-system Modelling
Sebastien Masson	LOCEAN - IPSL/UPMC	smasson@locean-ipsl.upmc.fr	FR	Oceanography
Bryan Lawrence	NCAS, STFC, UNIV.READING	bryan.lawrence@stfc.ac.uk	UK	Earth-system Modelling
Christopher Maynard	METOFFICE	christopher.maynard@metoffice.gov.uk	UK	Exascale Computing
Joachim Biercamp	DKRZ	biercamp@dkrz.de	DE	Exascale Computing
Heiner Igel	LMU	heiner.igel@geophysik.uni-muenchen.de	DE	Solid Earth, Geophysics
Colin Jones	SMHI	colin.jones@smhi.se	SE	Earth-system Modelling

11.3 List of experts of WG3.3

In June 2013, the Vice Chair of the working group 3.3 has changed from Romain Teyssier (CEA and ETHZ) to Alan Sacha Brun (CEA).

EESI2 - WG 3.3	**Fundamental Physics**			
Name	Organization	Email	Country	Area of Expertise
Godehard Sutmann (Chair)	JSC	g.sutmann@fz-juelich.de	DE	Quantum Chemistry
Alan Sacha Brun (Vice Chair)	CEA	sacha.brun@cea.fr	FR	Astrophysics
Thierry Deutsch	CEA	thierry.deutsch@cea.fr	FR	Quantum Chemistry
Nicola Marzari	Univ Oxford	nicola.marzari@materials.ox.ac.uk	UK	Material Sciences
Maurizio Ottaviani	CEA	Maurizio.Ottaviani@cea.fr	FR	Fusion
Volker Springel	Garching, MPI Astrophysik	volker@mpa-garching.mpg.de	DE	Astrophysics
Mike Payne	University of Cambridge	mcp1@cam.ac.uk	UK	Quantum Chemistry
Louis Silva	Universidade Tecnica de Lisboa	luis.silva@ist.utl.pt	PT	Laser Plasma Interaction
Ulf Meißner	University of Bonn /FZ Jülich	meissner@hiskp.uni-bonn.de	DE	Hadron/Nuclear Physics

11.4 List of experts of WG3.4

EESI2 - WG 3.4	**Life Sciences and Health**			
Name	Organization	Email	Country	Area of Expertise
Modesto Orozco (Chair)	IRB Barcelona	modesto@mmb.pcb.ub.es	SP	Modeling and Bioinformatics
Charles Laughton (Vice Chair)	U Nottingham	Charles.Laughton@nottingham.ac.uk	UK	Biochemistry & Drug discovery
Paolo Carloni	German Research School	p.carloni@grs-sim.de	DE	Computational Biophysics
Peter Coveney	University College London	P.V.Coveney@ucl.ac.uk	UK	Integrative medicine
Alfonso Valencia	CNIO	avalencia@cnio.es	SP	Biological Sciences
Richard Lavery	University of Lyon	richard.lavery@ibcp.fr	FR	Simulation
Erik Lindhal	Royal Institute of Technology (KTH)	lindahl@cbr.su.se	SE	Molecular Dynamics
Henry Markram	École Polytechnique Fédérale de Lausanne	henry.markram@epfl.ch	CH	Organ simulation
Darren Green	GSK	darren.vs.green@gsk.com	UK	Drug Discovery and HPC
Supporting Experts				
Shantenu Jha	Rutgers	shantenu.jha@rutgers.edu	US	
Martin Guest	Cardiff	GuestMF@cardiff.ac.uk	UK	

11.5 List of experts of WG3.5

EESI2 - WG 3.5	Disruptive Applications			
Name	Organization	Email	Country	Area of Expertise
Godehard Sutmann (Chair)	JSC	g.sutmann@fz-juelich.de	DE	Quantum Chemistry
Ulrich Rude	U. Erlangen/FAU	uli.ruede@fau.de	DE	Computer Science