Next generation applications

EXDCI WP3 – Applications Grand Challenges

S. Requena - PRACE / GENCI – EXDCI WP3 leader
and all the WP3 contributors

September 7, 2017

The EXDCI project has received funding from the European Unions Horizon 2020 research and innovation programmed under the grant agreement No 671558.

Outline

- Introduction on EXDCI WP3
- Few words about the overall context
- Overview of some grand challenge applications
- First recommendations

WP3 : Applications Roadmap Toward Exascale

- Objectives
  - Provide updated roadmaps of needs and expectations of scientific applications
  - Provide inputs to the update the **PRACE Scientific Case** in order to support PRACE in the deployment of its (Pre)Exascale pan European HPC research infrastructure

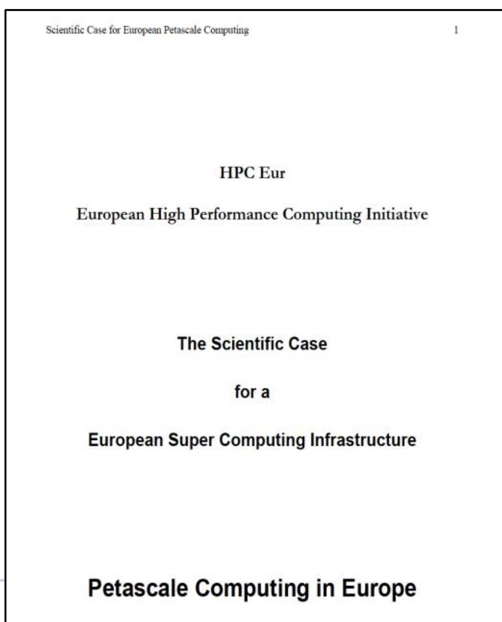
Scientific Case for European Petascale Computing 1

HPC Eur

European High Performance Computing Initiative

The Scientific Case

for a

European Super Computing Infrastructure

Petascale Computing in Europe

WP3 structure in a nutshell (1/2)

- Organisation
  - 4 working groups of 45 experts
    - ✓ T3.1 : Industrial and engineering applications
      (EDF : Yvan FOURNIER, University of Aachen : Heinz PITSCH)
    - ✓ T3.2 : Weather, Climatology and Solid Earth Sciences
      (Univ. Salento/CMCC : Giovanni ALOISIO, JCA Consultance: Jean-Claude ANDRE)
    - ✓ T3.3 : Fundamental Sciences
      (CEA : A. Sacha BRUN, JSC: Stefan KRIEG)
    - ✓ T3.4 : Life Science & Health
      (BioExcel CoE/ KTH : Rossen APOSTOLOV, CompBioMED / UCL : Peter COVENEY)
  - T3.5: Coordination, consolidation of application roadmaps and inputs to the update of the PRACE Scientific Case

H2020-FETHPC-2014

Coordination of the HPC strategy

EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D3.1
First set of recommendations and reports toward applications

Final

Version: 1.0
Author(s): S. Requena (GENCI) et al.
Date: 25.10.2016

- Deliverables
  - D3.1: First set of reports and recommendations toward applications (**M15**)
  - D3.2: 3rd version of the PRACE SSC Scientific Case – a full bottom-up new version of PRACE SSC Scientific Case, following the last one published in 2012 (**M27 = end 2017**)

WP3 structure in a nutshell (2/2)

- Contribution from 45 experts of 8 countries
- Academia & industry

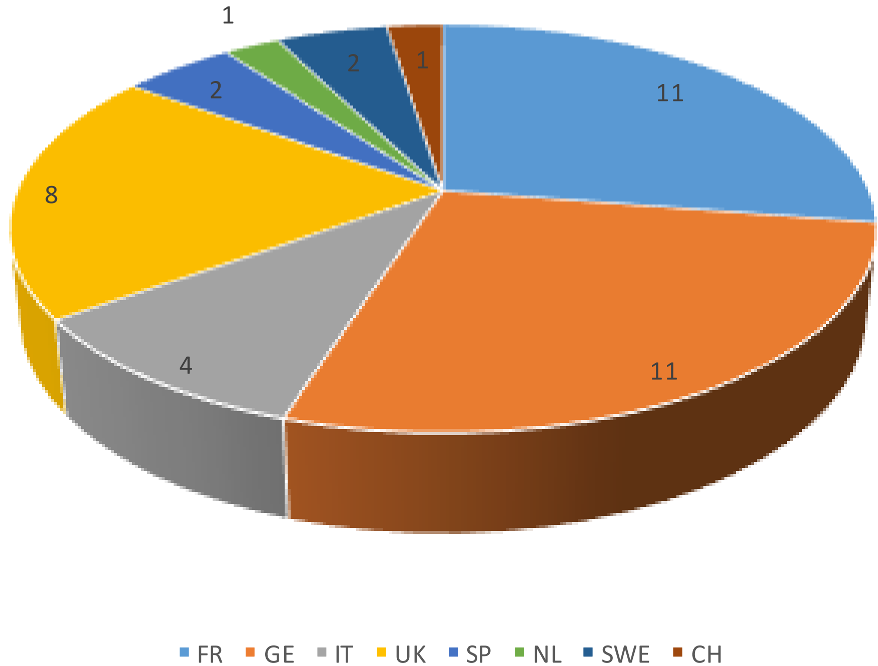

- Strong collaborations with the CoE :

esiwace CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE in WG3.2

E-CAM, E-CoE, NOMAD NOVEL MATERIALS DISCOVERY in WG3.3

bioexcel Center of Excellence for Computational Biomolecular Research, COMPBIOMED leading WG3.4

exdci European Extreme Data & Computing Initiative

The context : the road to Exascale

- Expected in
  - 2020 for China
  - 2021/22 for US, EU and Japan
- No more focus on peak performance
- Systems 50 to 100x faster than current ones on **real apps**

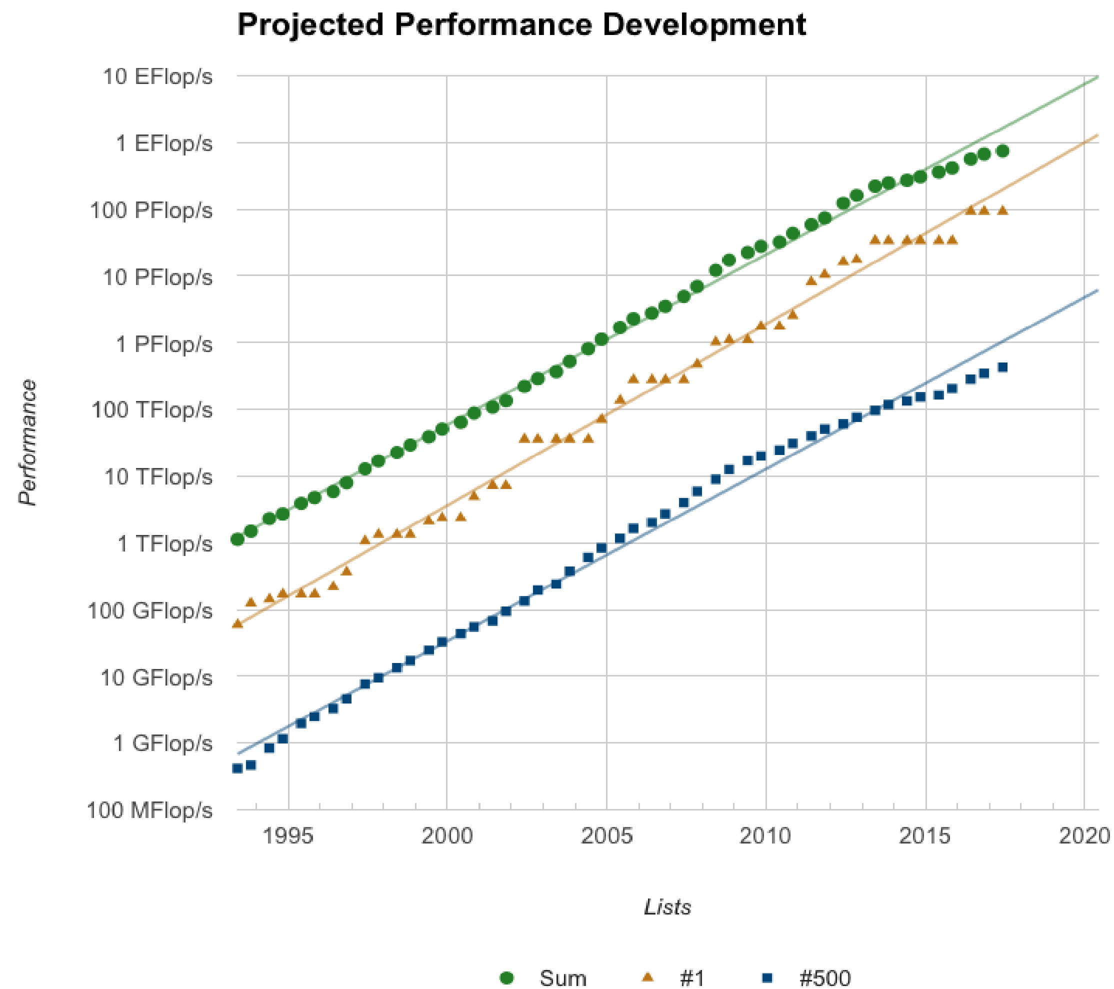

Source top500.org

The context : Convergence between HPC, HPDA and AI

Adapted from illustration in *Abduction and Induction: Essays on their relation and integration*

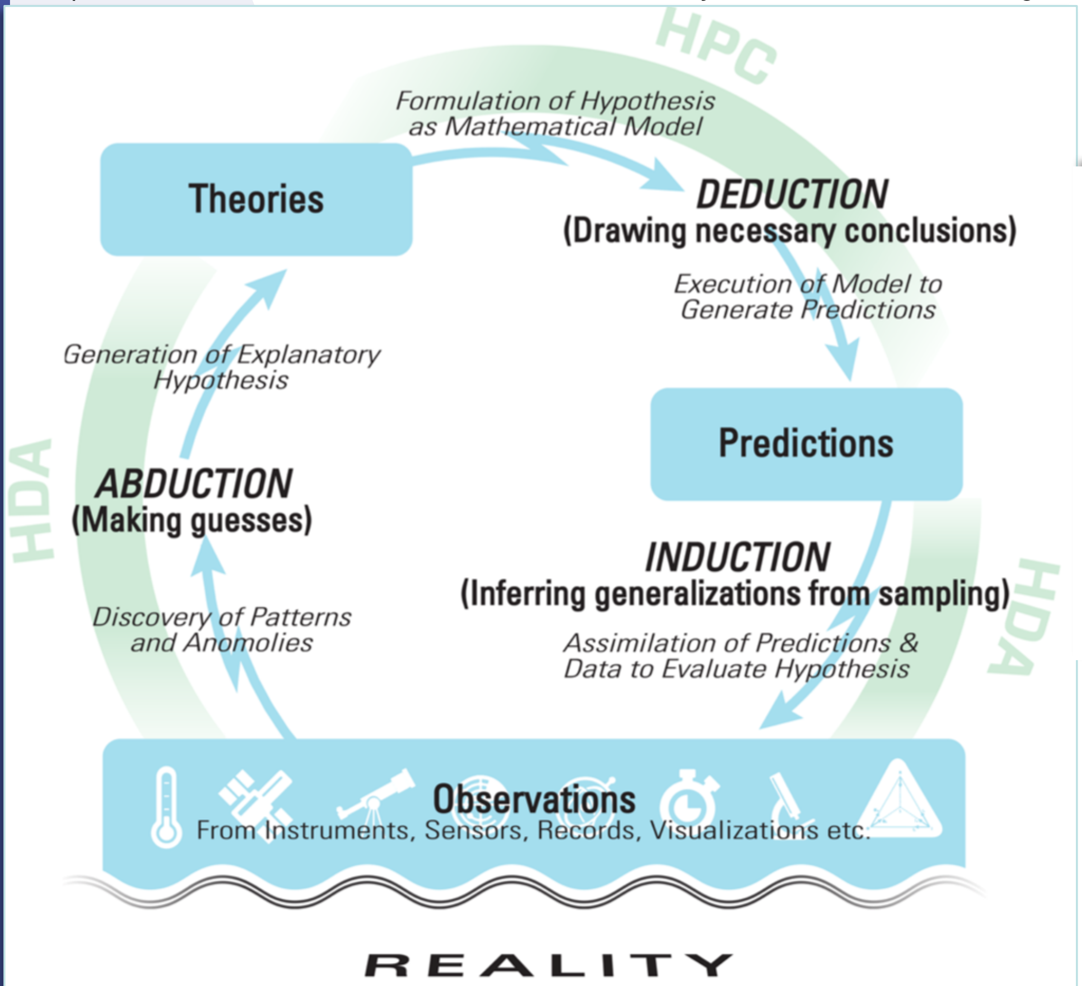

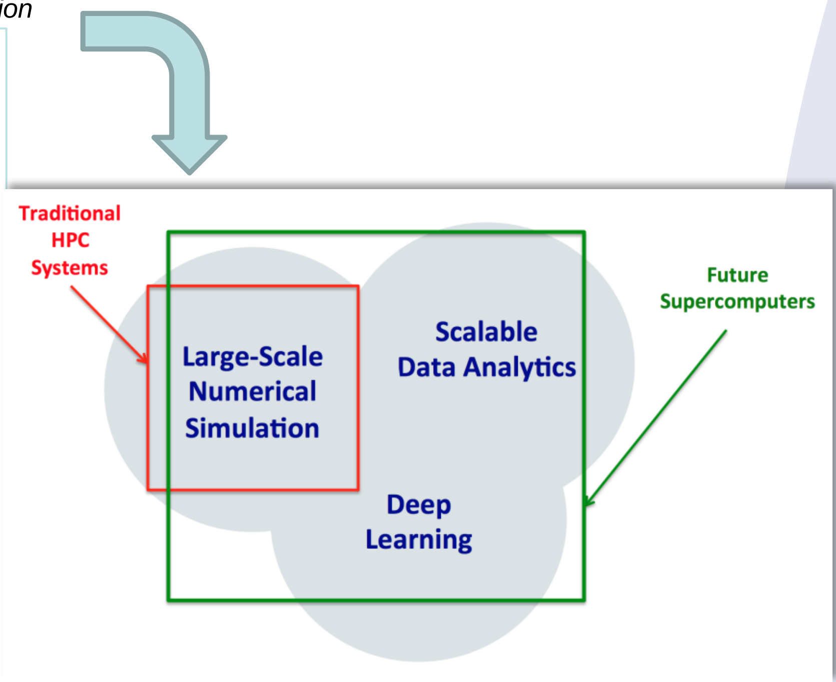

- Challenges : cohabitation of SW stacks, containers, security, smart resource managers, end to end workflows, edge computing (at the source), …

exdci European Extreme Data & Computing Initiative

The context : applications are THE meaning of HPC and Big Data

Europe owns a lot of applications
(>50% in chemistry/material sciences)

Europe is the biggest generator of data
(C. Moedas, Eu Open Science Cloud conf., 19 April)

EXASCALE on critical path

But BIG DATA ANALYTICS even more

Some Exascale applications challenges

- In the Oil & Gas : toward FWI in seismic exploration in 2020

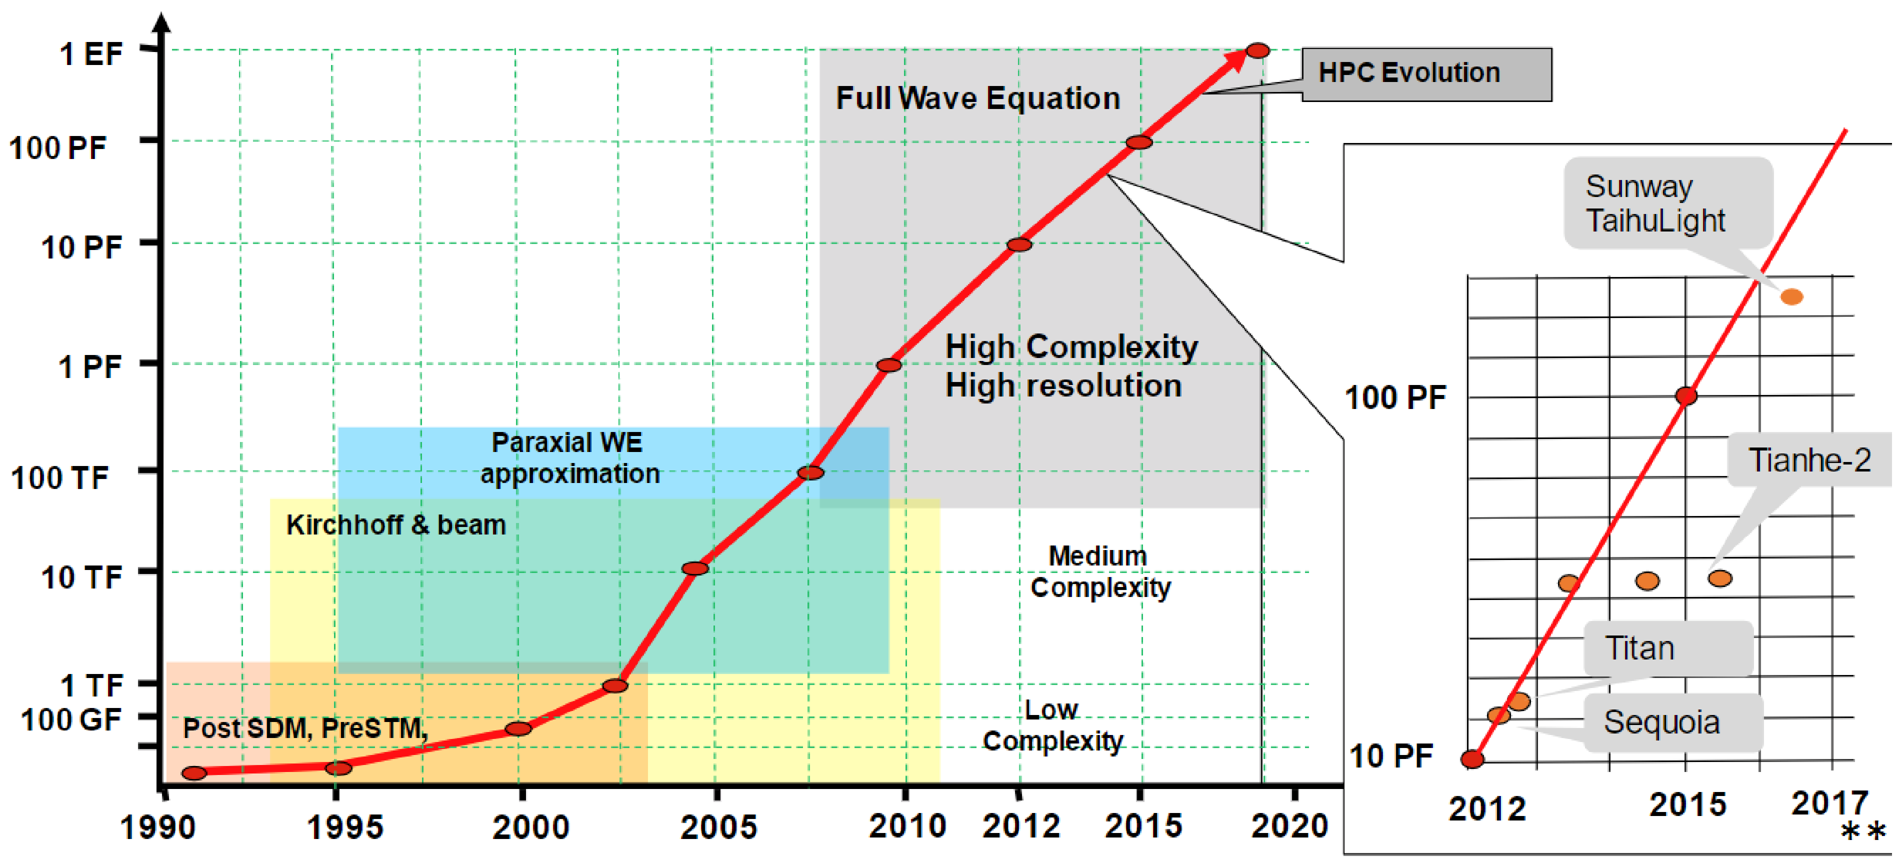

* Rice Oil & Gas HPC Workshop, 2011 (Henri Calandra, John Etgen & Scott Morton)
** Rice Oil & Gas HPC Conference, 2017 (Henri Calandra, Scott Morton, & Sverre Brandsberg-Dahl)

Some Exascale applications challenges

- In aeronautics : toward a full multiphysics virtual bird in 2030 !

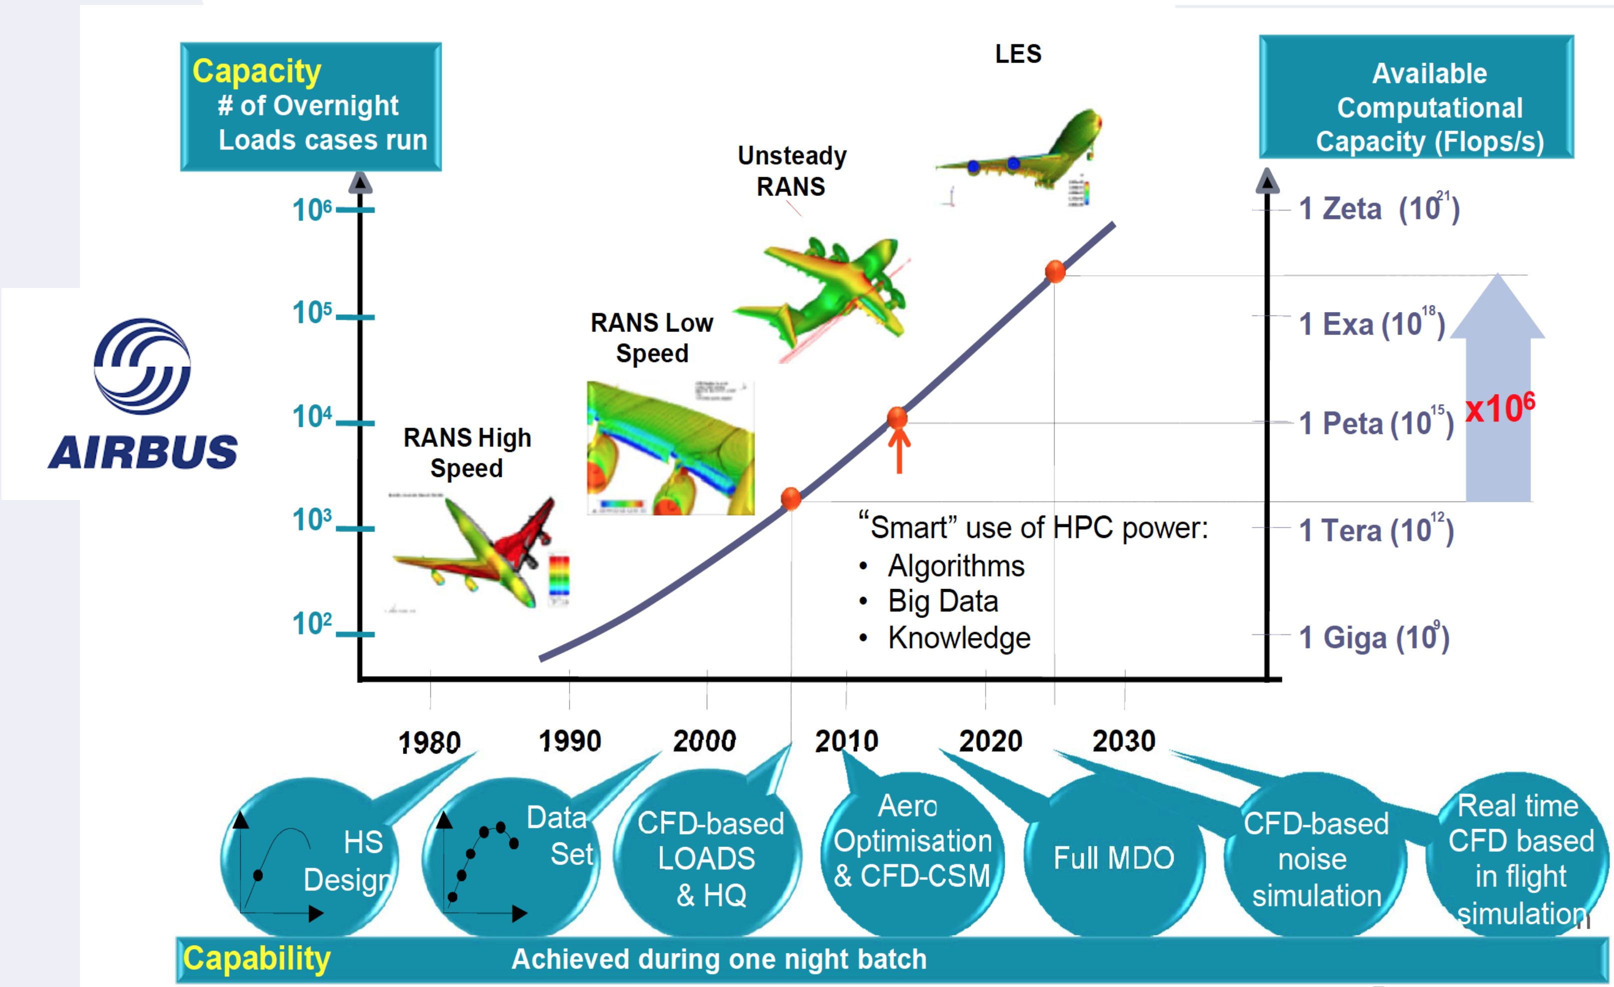

Some Exascale applications challenges

- Toward a full LES simulation of a 3.5-stage gas turbine compressor
- Project « Cœur Numerique 2020 » led by Safran

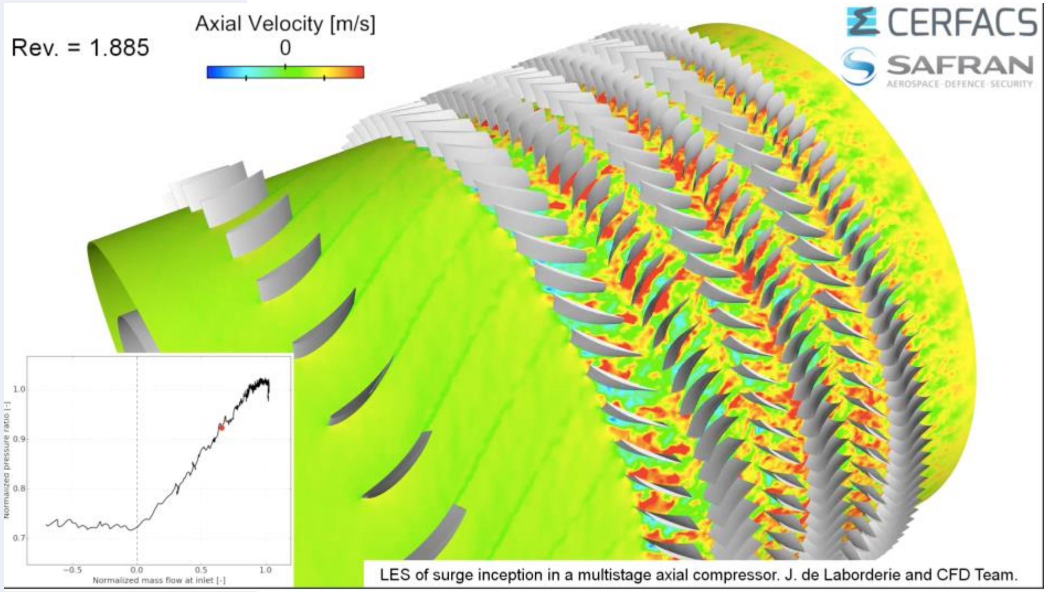

LES of surge inception in a multistage axial compressor. J. de Laborderie and CFD Team.

- importance of smart and scalable code couplers
- importance of MPMD on exascale (performance/intelligent load balancing)

Some Exascale applications challenges

- EDF grand challenge in fuel assemblies

  - 157 - 241 fuel assemblies in a PWR
  - Each fuel assembly has 17*17 fuel rods or guide tubes and 8 to 10 grids
    - complex, non symmetrical geometry
    - Different vendors and models
  - Many constraints / stakes
    - If head loss/lift too high, stronger springs needed to keep FA down, leading to possible bowing and deformation
    - Good heat exchange: need mixing grids to generate turbulence
    - Low vibration: loss of cladding integrity may result from vibration induced fretting
  - Complex geometry
    - Need for better automated meshing tools capable of generating very large and high quality meshes

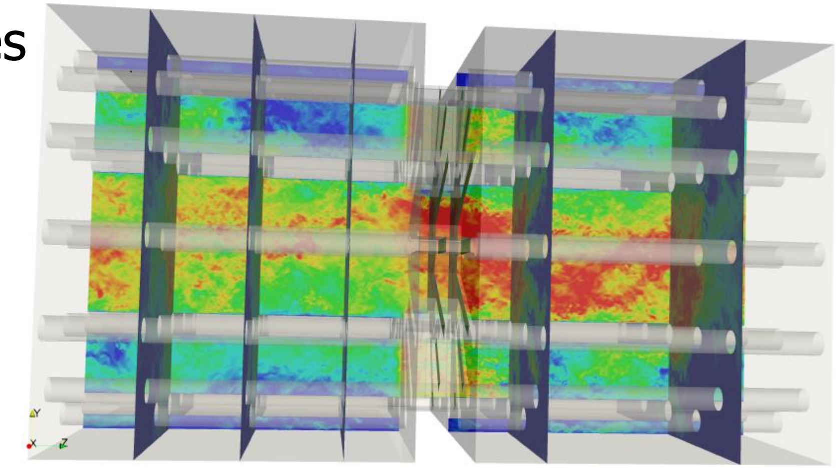

Billion-cell LES to obtain pressure loads on fuel rods; This is only a small subset of the full geometry

deformation

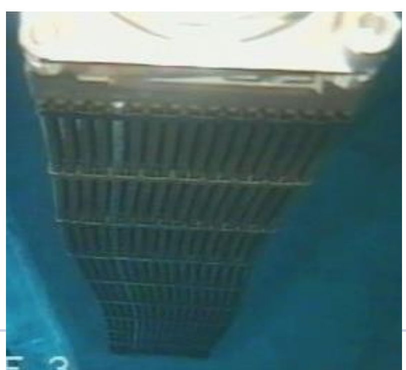

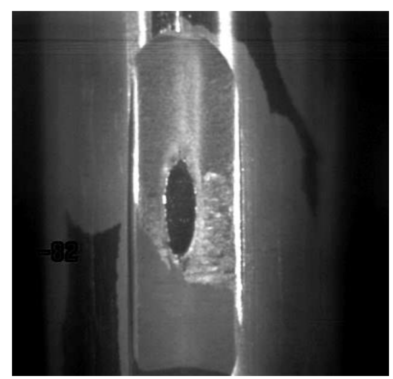

fretting-induced wear of cladding

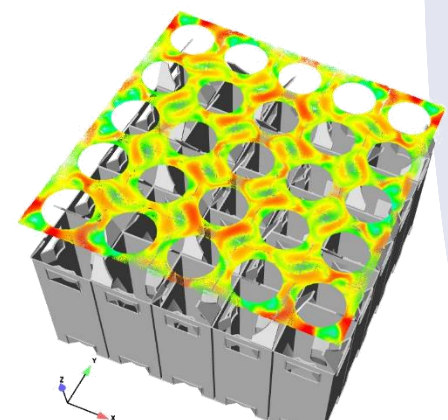

Some Exascale applications challenges

Target for addressing key science challenges in weather & climate prediction:

Global 1-km Earth system simulations @ ~1 year / day rate

10 km

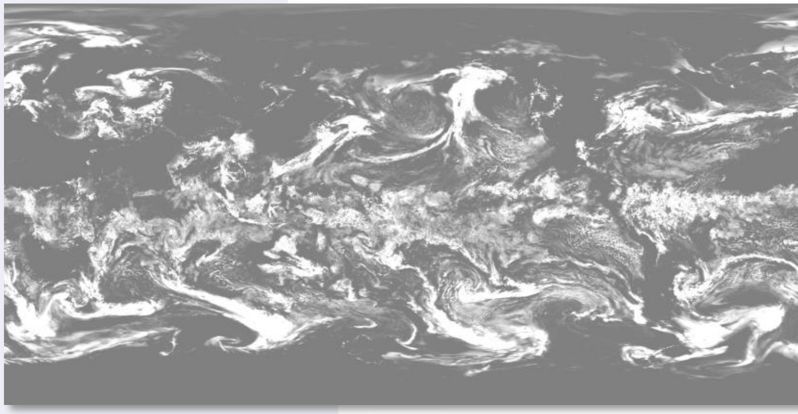

1 km

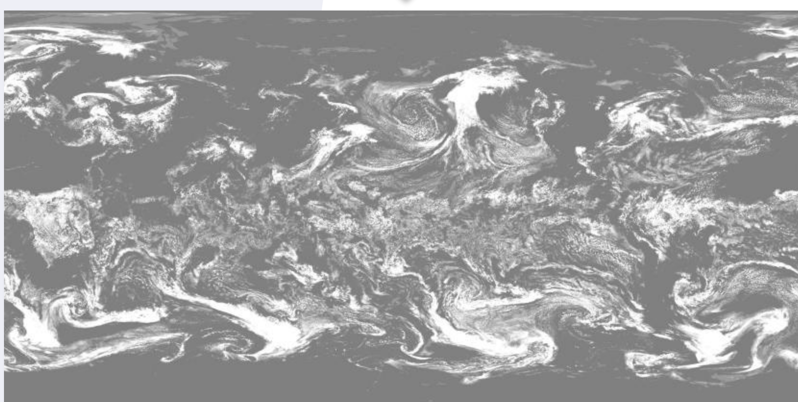

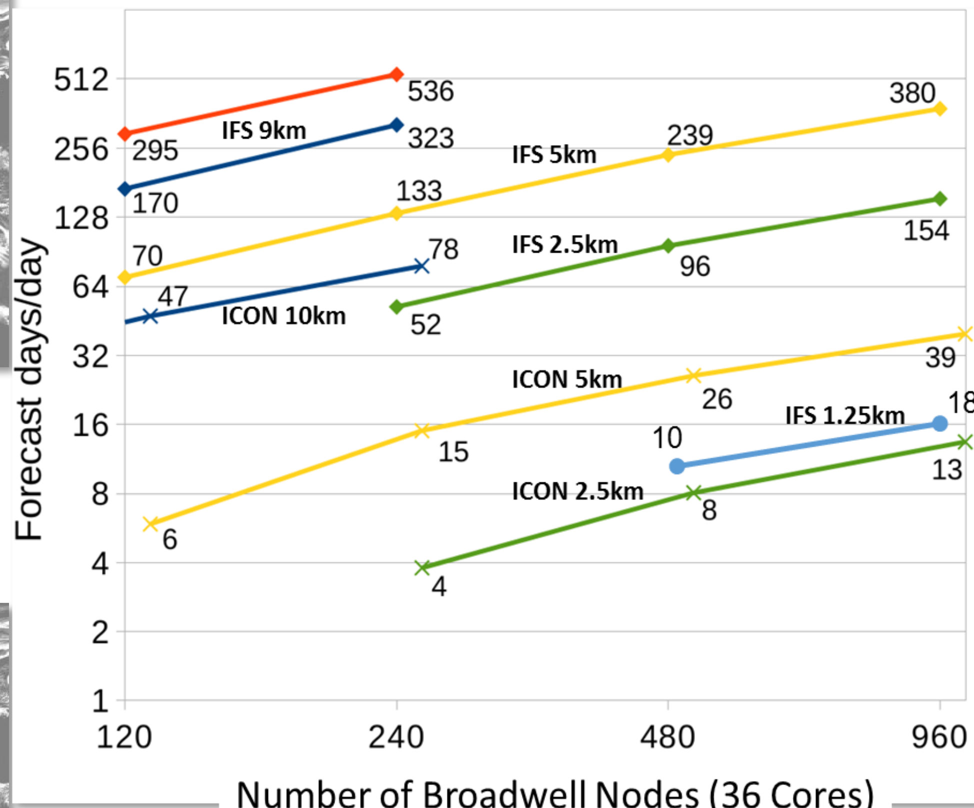

Some Exascale applications challenges

- Massive cosmological simulations toward the next Era of Galaxy Surveys
  - Next gen satellites (Euclid, WFIRST) and telescopes (LSST) for 2020-2022
  - Ongoing massive simulations to feed / optimise future observations
  - Understanding evolution Universe wrt initial dark matter / energy

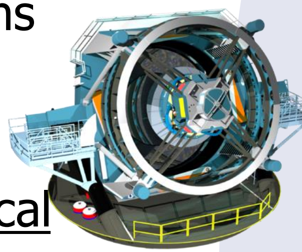

- May 2017 : Team from Univ. Zurich performed the biggest cosmological simulation with 2 billion particles using 4000+ GPU on PizDaint (CSCS)
  - PKDGRAV3 : N-Body simulation using Fast Multipole Methods (FMM)
  - Aggressive optimisations performed
    - reduced memory footprint, mix of single/double precision, …
    - load balancing between CPU and GPU, on the fly post processing of the data
  - Excellent scaling up to 18 000 GPU
    - Benchmark performed using 8 trillion particles on Titan (ORNL/USA)

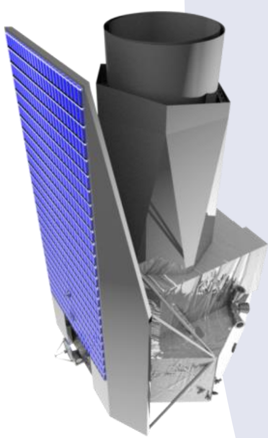

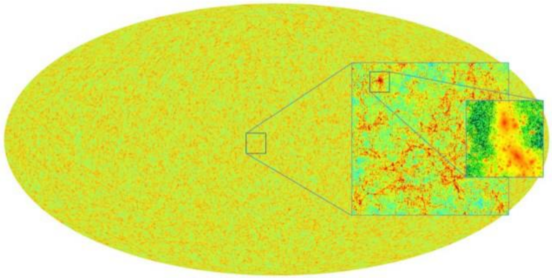

European Extreme Data & Computing Initiative

Some Exascale applications challenges

- Postprocessing of massive amount of data issued by the SKA radiotelescope

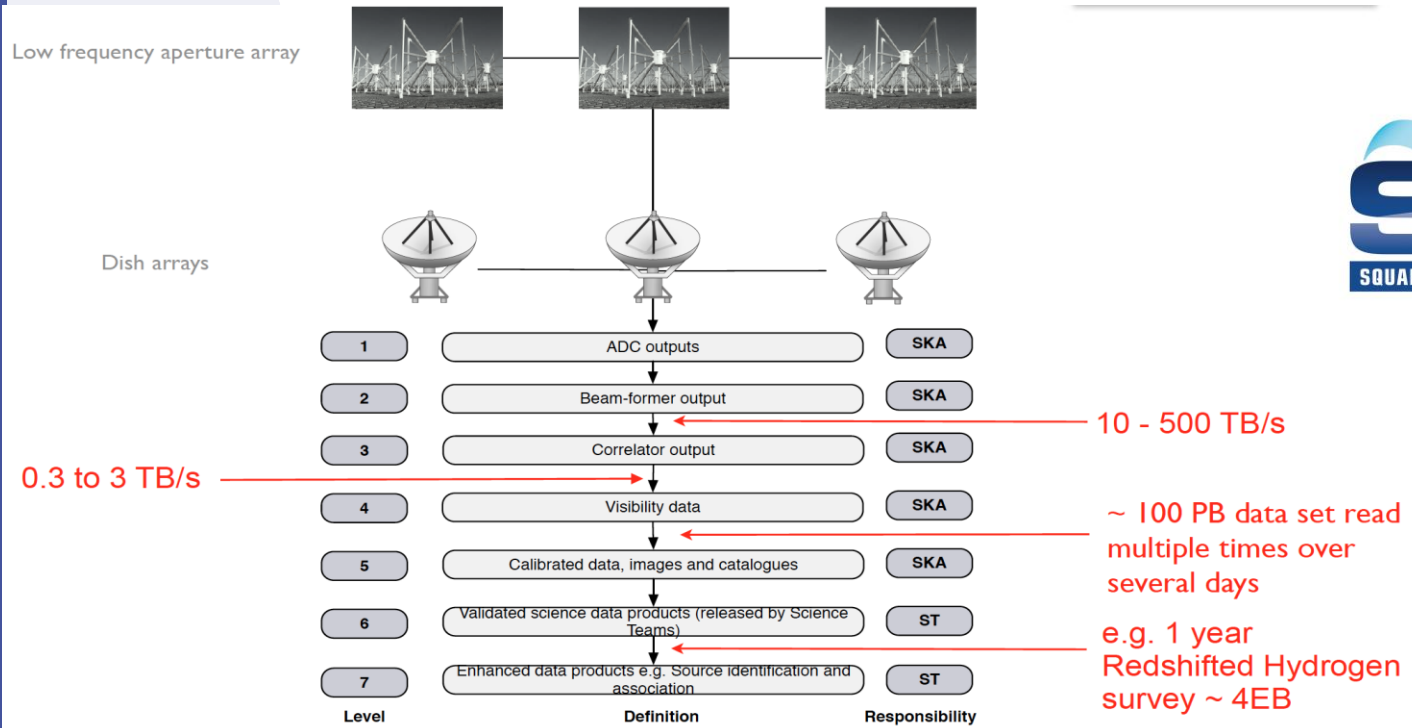

# 2013 estimate by SKA South Africa	MeerKAT Pre-Cursor 2014-15	SKA Phase 1 2017-19	SKA Phase 2 Est. 2020-24
Data into CSP	2 Tbps	50 Tbps	up to 5 Pbps
Data into SDP	0.4 Tbps	20 Tbps	up to 500 Tbps
Into Storage	35 Gbps	300+ Gbps	up to 2 Tbps
Computing load	200 TFlops	30+ PFlops	3+ EFlops

Exascale challenges in biomolecular sciences

bioexcel — Center of Excellence for Computational Biomolecular Research

Molecular Dynamics on the exascale

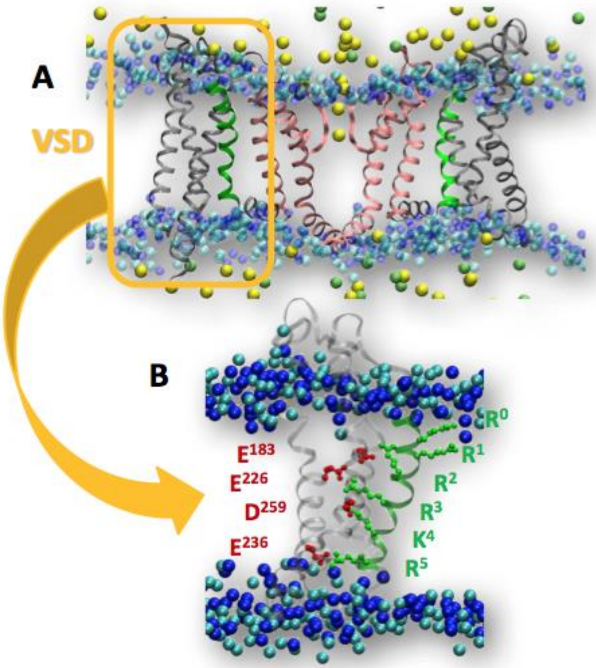

- Understanding proteins and drugs
- A 1 μs simulation: 10 exaflop
- Many structural transition: many simulations needed
- Study effect of several bound drugs
- Study effect of mutations
- All this multiplies to >> zettaflop
- Question: how far can we parallelize?

Example: ion channel in a nerve cell. Opens and closes during signalling. Affected by e.g. alcohol and drugs. 200 000 atoms

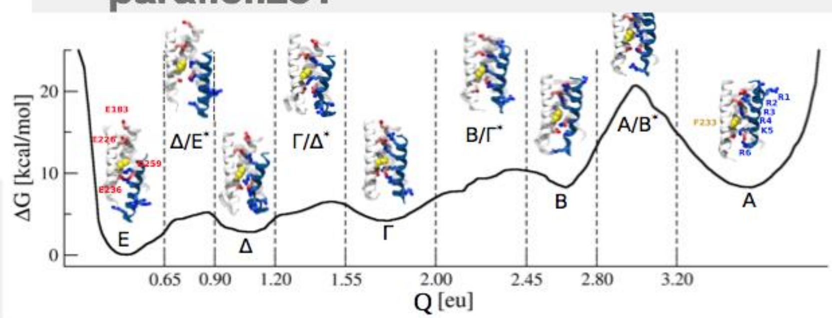

Partners

Horizon 2020 European Union Funding for Research & Innovation

Exascale challenges in biomolecular sciences

bioexcel Center of Excellence for Computational Biomolecular Research

Predicting interactomes by docking… a dream?

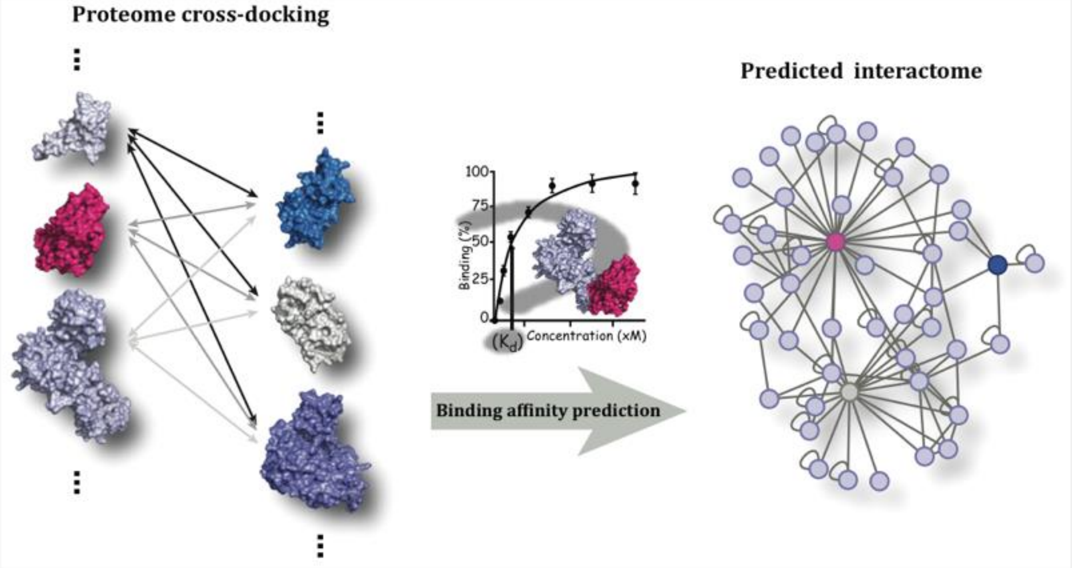

- ~20'000 human proteins
- Interactome prediction will require 20'000^2 docking runs
- Which will require > 10 billions CPU hours and generate about 100 exabytes of data
- Interest in simulating/understanding the impact of disease-related mutations that affect/alter the interaction network

Partners

Funding

Horizon 2020 European Union Funding for Research & Innovation

D3.1 – WP3 first global recommendations

- The context

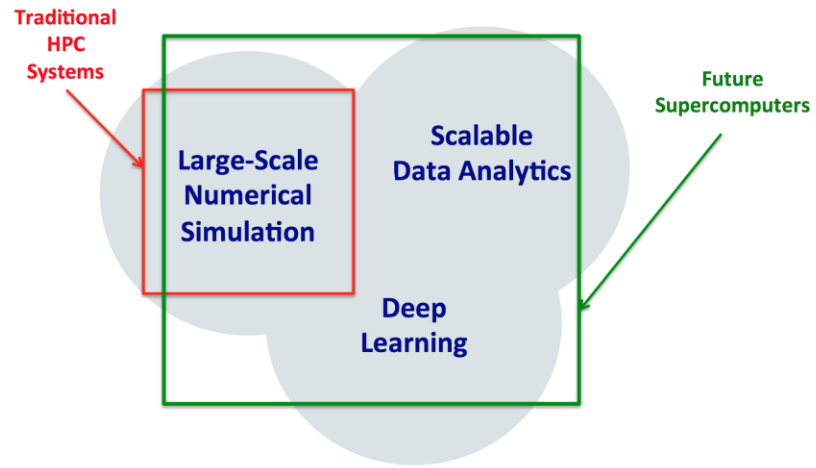

- **R1 : Toward the convergence of in-situ data analysis and deep-learning methods for efficient post processing of pertinent structures among huge scientific datasets**
- **R2 : New HPC services toward improved link with large scale instruments and urgent computing**
  - Interactive supercomputing, co-scheduling, smart app based checkpoint/restart, malleability of applications, …
- **R3 : Development at the EU level of new CoE**
  - Engineering and industrial applications
  - Opensource software industrialisation, promotion and long term support

Toward the convergence of in-situ data analysis and deep-learning methods for efficient post processing of pertinent structures among huge scientific datasets

S. Fiore[1], G. Aloisio[2], P. Ricoux[3], S. Brun[4], JM. Alimi[5], M. Bode[6], R. Apostolov[7], S. Requena[8]

[1]Euro-Mediterranean Center on Climate Change Foundation, Italy and ENES

[2]University of Salento, Italy

[3]Total, France

[4]CEA, France

[56]LUTh, CNRS, OBSPM, France

[6]RWTH Aachen University, Germany

[7]KTH and BioExcel CoE, Sweden

[8]GENCI, France

September 7, 2017

The EXDCI project has received funding from the European Unions Horizon 2020 research and innovation programmed under the grant agreement No 671558.

The context (1/2) – Big Data is a today problem !

Explosion of the data from the computational side :

- High spatial and temporal resolutions
- Multiscale and multiphysics simulations
- Ensemble and UQ simulations

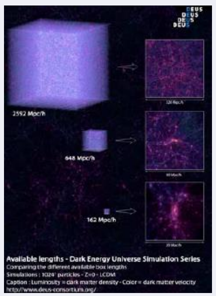

Cosmology
DEUS project
150 PB raw data

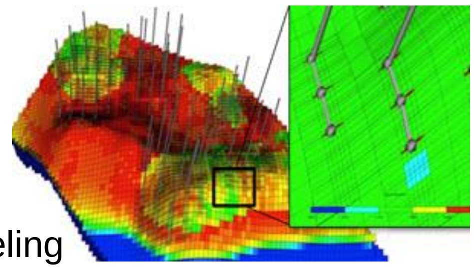

Reservoir modeling
of gigamodels 350 TB/run

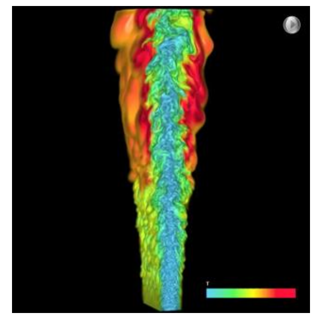

HiFi turbulent
DNS combustion
S3D : 1PB / 30mn

Ongoing convergence between HPC and HPDA platforms

- Data centric architectures
- Tiered storage
- Prg language support for data analysis, DSLs, …

BUT :

Simulation time : from days to months

Post processing time : from months to years

The context (2/2)

Explosion of the data from the computational side :

- High spatial and temporal resolutions

Necessity to reduce the time to discovery :
amount of data impossible to process in a reasonable time for humans !
New approaches are needed !

EESI2 recommendations toward in-situ/in-transit post processing

- EESI2 FP7 project (sept 2012 – May 2015) - vision and roadmaps toward Exascale Eu applications and software tools
- 15 scientific recommendations proposed to EC
  → « In Situ Extreme Data Processing and better science through I/O avoidance in HPC systems »

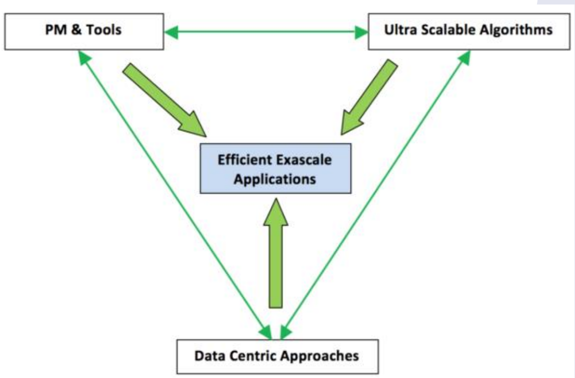

- Objective : Funding of appropriate R&D programs
  - Design and implement in-situ and in-transit post-processing frameworks
    → maximize data utilisation while loaded and minimize raw data movements to remote storage units in order to save energy
  - Perform data reduction, ordering, filtering, compression, …
  - And feature extraction through topological segmentation or trajectory based flow feature tracking, ...

And now comes machine/deep learning

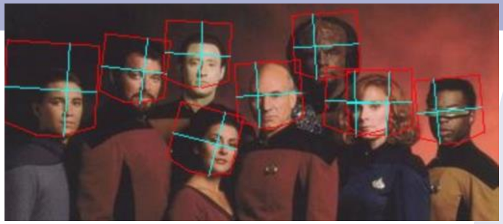

- Machine/Deep learning : the reborn of AI
  - use of HPC and neural networks for automatic (face, picture, speech, form, fraud, …) detection, medical diagnostic, drug design,… in big data analytics

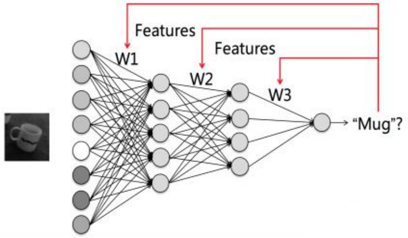

- Idea : <u>to use ML/DL methods for improving in-situ, perform AI assisted computational steering, replace models with learned functions …</u>
  - smart feature extraction into massive amount of data
  - already some initial work done in Europe
  - and funding available in US : 17.7M$ from NSF for 12 data science projets including ML/DL ones
- Strong expertise to pool, both in academia and industry, in Europe
  - France (Inria, CNRS, UPMC), Spain (BSC), Germany (Tum), UK (Alan Turing Inst.), ...
  - IBM Zurich, Sony CSL, Facebook FAIR, ...

Proposed recommendation

- Extend current EESI2 recommendation toward in-situ/in-transit methods with machine/deep learning features
  - Assess concretely the potential of these techniques to 10 to 12 scientific pilots
  - <u>Bridge together communities of scientific computing and ML/DL</u>

BDV BIG DATA VALUE ASSOCIATION

- Organisation of a (urgent) call for proposals funded by EC or national agencies
  - Experts from domain science : climate, astrophysics, fusion, combustion, seismic, biology, chemistry, …
  - Experts in applied maths and machine/deep learning methods
  - Experts from HPC centres
  - Target of 10 M€ for a first call, could be scaled out to a joint BDVA – ETP4HPC joint project ?

EXDCI WP3 and SRA3 : 5yr - specific challenges per communities

	WG3.1 Industrial & engineering apps	WG3.2 WCES	WG3.3 Fundamental Sciences	WG3.4 Life Sciences & Health
HPC System Architecture and Components	move to GPU or others accelerators virtualisation of resources	co design, GPU/MIC move hierarchical storage	GPU, MIC, FPGA, specialised HW ok with opt. libs, memory per core (astrophysics, fusion), mem & network BW	Large width vector units, low-latency networks, high-bandwidth and large memory; fast CPU<-> accelerator transfer rates, Heterogeneous acceleration, floating-point
System Software and Management	smart runtime systems, scalable couplers, data-aware schedulers	scalable couplers, smart scheduling	scalable couplers, mix of capacity/capability	urgent computing, link with instruments, co-location of compute & data, dynamic scheduling, support for workflows
Programming Environment	use of standards, sustainability	use of DSL, solution & performance portability	task programming / runtimes, use of DSL, standards	portability, fast code driven by python interfaces
Energy and Resiliency	scalable monitoring tools, data compression, application based FT	mixed precision	reduced precision, data compression	Distributed computing techniques to handle resiliency/fault tolerance
Balance Compute, I/O and Storage Performance	in-situ/in transit post processing	in-situ/in transit post processing active storage techniques	dynamic load balancing in mesh refinement and spectral element techniques	I/O driven, in memory database, Data-focused workflows, handling lots of small files in bioinformatics
Big Data and HPC usage Models	compression of data, remote viz, UQ/optimisation, ML/DL, mix of capacity/capability	UQ, compression of data, multi-site experiments to support multi-model analysis experiments, in-memory analytics, HPC-through-the cloud, ML/DL	remote data viz, ML/DL	convergence HPC-HTC, integration of HPDA tools (Hadoop, Spark, …) inc ML/DL Cloud Computing security / privacy
Mathematics and algorithms for extreme scale HPC systems	ultra-scalable solvers, // in time, automatic/adaptive meshers, model order reduction , meshless and particle simulations, coupling stochastic and deterministic methods	parallelisation in time, ensemble simulations	scalable solvers, adaptive meshers, // in time, FMM and h-matrices	multiscale/physics workflows tools, ensemble simulation, model order reduction

D3.2 – WP3 final recommendations

- Ongoing process planned for end of 2017
  - Use and development of "**standard tools**" (à la MPI+X) including domain specific libraries and languages (DSL) to abstract underlying complex HW
  - **Ultra-scalable solvers** inc. communication/synchronisation avoiding, parallelisation in time, high order discretisation, H-matrices, …
  - Smart and scalable **meshers and code couplers** for supporting next gen multi-scale and multiphysics simulations
  - Importance of EU-wide **optimisation and UQ frameworks**
  - **Training** toward data analysis and AI
- Final deliverable **D3.2 : Update of the Scientific Case end 2017**
  - Specific and global recommendations provided
  - EXDCI WP3 to feed PRACE SSC effort

European Extreme Data & Computing Initiative

Download the D3.1 report www.exdci.eu

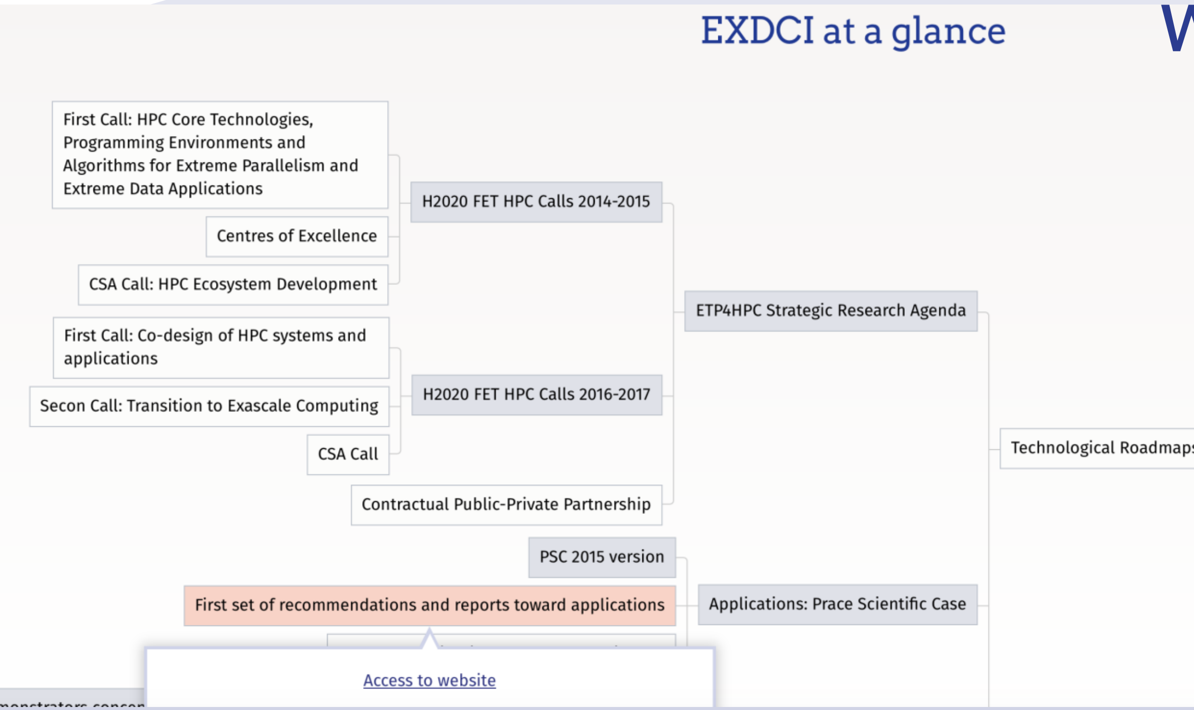

More info : exdci-wp3@prace-ri.eu

September 7, 2017

The EXDCI project has received funding from the European Unions Horizon 2020 research and innovation programmed under the grant agreement No 671558.