The Abdus Salam International Centre for Theoretical Physics

High-Performance Computing at the ICTP: Challenges of Large Scale Scientific Simulations and Programs for Education

Ivan Girotto – igirotto@ictp.it

International Centre for Theoretical Physics (ICTP)

What is High-Performance Computing (HPC)?

- Not a real definition, depends from the prospective:
  - HPC is when I care how fast I get an answer
  - HPC is when I foresee my problem to get bigger and bigger
- Thus HPC can happen on:
  - A workstation, desktop, laptop, smartphone!
  - A supercomputer
  - A Linux Cluster
  - A grid or a cloud
  - Cyberinfrastructure = any combination of the above
- HPC means also High-Productivity Computing

Why use Computers in Science?

- Use complex theories without a closed solution: solve equations or problems that can only be solved numerically, i.e. by inserting numbers into expressions and analyzing the results
- Do “impossible” experiments: study (virtual) experiments, where the boundary conditions are inaccessible or not controllable

Why use Computers in Science?

- Reduce costs of experiments

Why use Computers in Science?

- Benchmark correctness of models and theories: the better a model/theory reproduces known experimental results, the better its predictions
- Predict complex theory applying techniques of AI/Deep learning

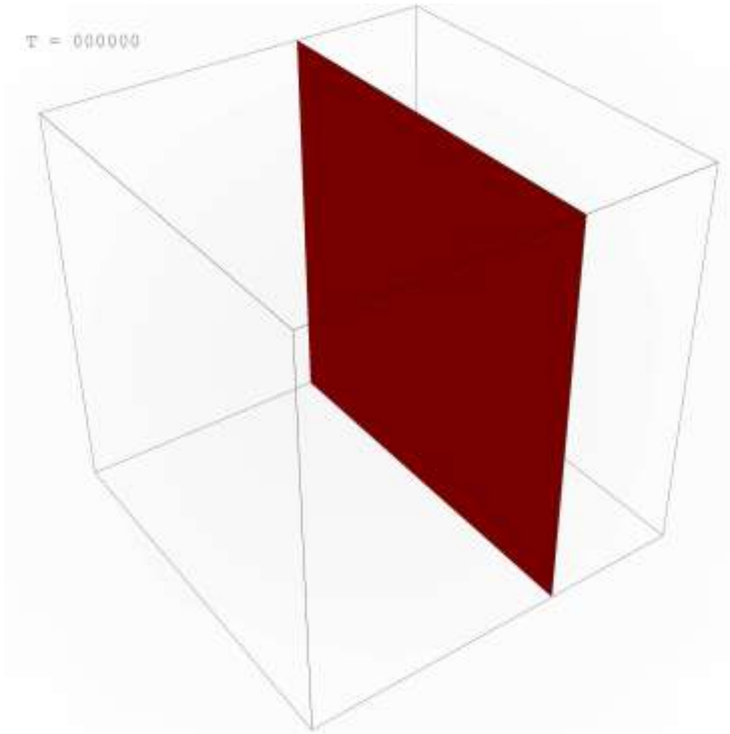

*** PRACE project, TurEmu – The physics of (turbulent) emulsions, lead by Prof. Toschi at TU/e**

The growing computational capacity

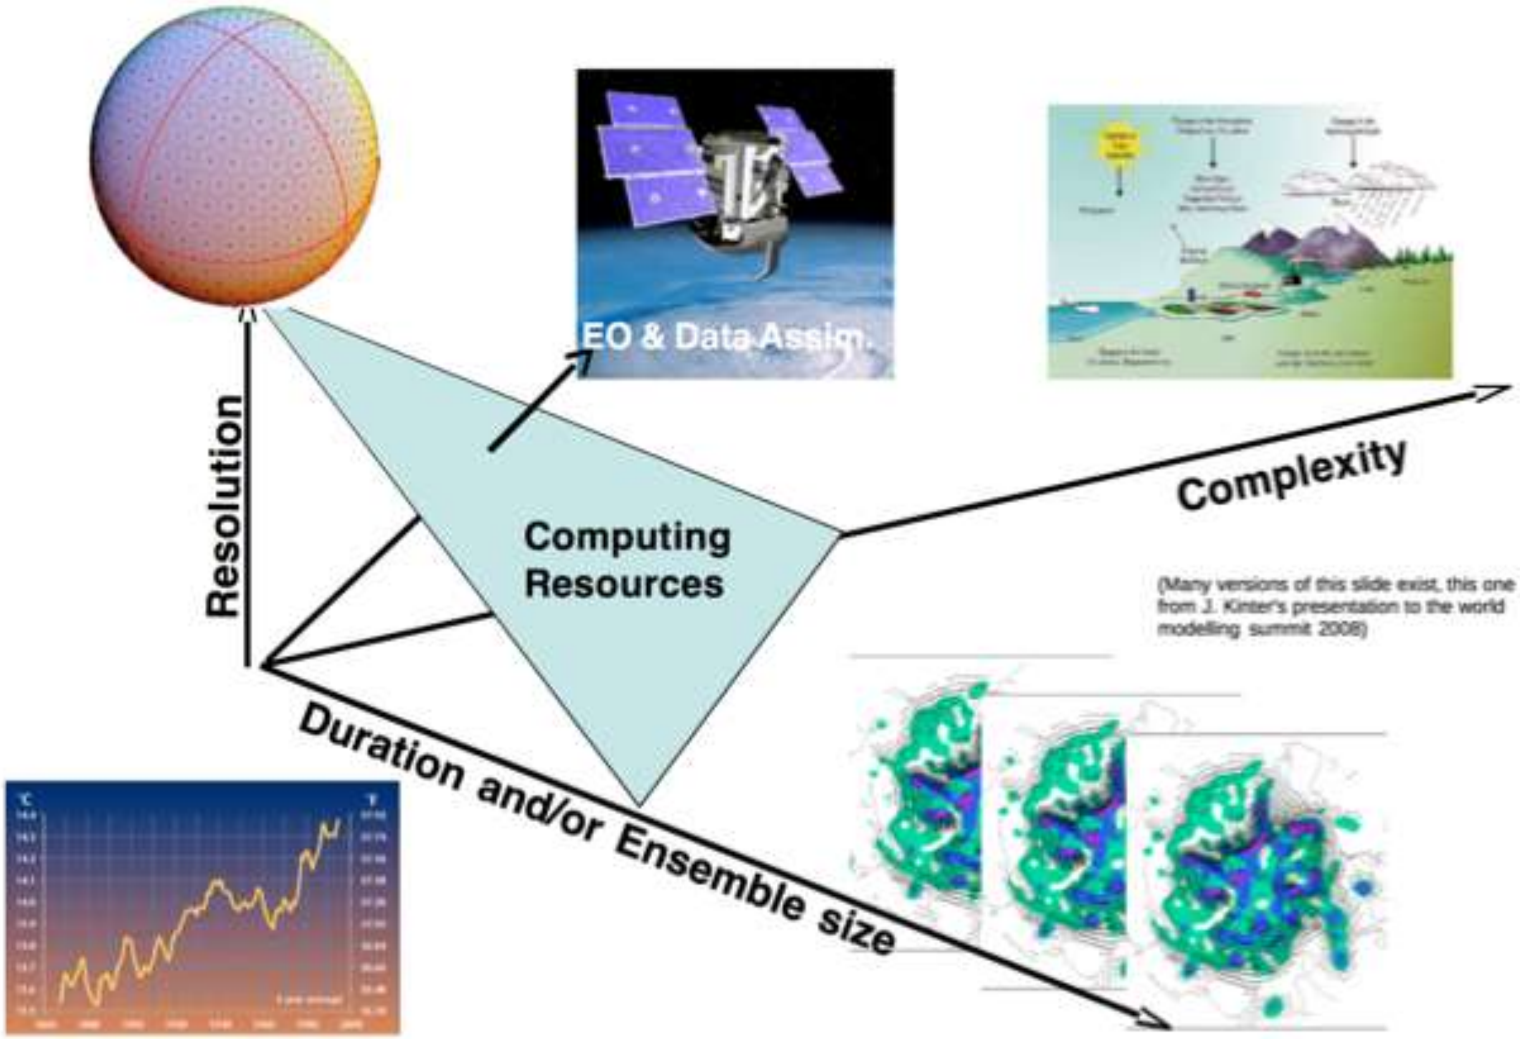

Impact of Using Computer is Science

- **A more competitive industry**
  - We could never have designed the world-beating Airbus A380 without HPC
  - Thanks to HPC-based simulation, the car industry has reduced the time for developing new vehicle platforms from 60 months to 24
- **Direct benefits to our health**
  - One day of supercomputer time was required to analyse 120 billion nucleotide sequences, narrowing down the cause of a baby's illness to two genetic variants. Thanks to this, effective treatment was possible and the baby is alive and well 5 years later
- **Better forecasting**
  - Severe weather costs 150.000 lives and **€**270 billion in economic damage in Europe between 1970 and 2012
- **Making possible more scientific advances**
  - Supercomputing is needed for processing sophisticated computational models able to simulate the cellular structure and functionalities of the brain
- **More reliable decision-making**
  - The convergence of HPC, Big Data and Cloud technologies will allow new applications and services in an increasingly complex scenario where decision-making processes have to be fast and precise to avoid catastrophes

*** from EU Digital Single Market Blog by Roberto Viola, Director-General, DG Communications Networks, Content and Technologies and Robert-Jan Smits, Director-General, DG Research and Innovation**

HPC as a Priority (in a nutshell)

GLOBAL CONTEXT: *Increasing Recognition of Impact of HPC Enabled Applications*

U.S. President Barack Obama issued an executive order establishing the National Strategic Computing Initiative with initial allocation of $ 450 M from the U.S. Department of Energy). (29/07/2015)

Former Vice President of the European Commission Neelie Kroes announced the launch of a contractual Public-Private Partnership (cPPP) on Big Data (€ 500M EC funding) and for the development of a European HPC Eco-system (€ 700 M EC funding). (17/12/2013)

French President François Hollande announced the launch of French National "Plan Supercalculateur" with € 50M of national funding. (07/05/2014)

Chinese President Xi applauds success of world's fastest supercomputer Tianhe-2 built by China's National University of Defense Technology with the fundamental purpose of

"providing a driving force for the construction of an innovation-oriented country". (18/06/2013)

HPC Development Trend

List	Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
06/2018	1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,282,544	122,300.0	187,659.3	8,805.5
11/2017	1	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
06/2017	1	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
11/2016	1	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
06/2016	1	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
11/2015	1	Tianhe-2A - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P , NUDT National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808
06/2015	1	Tianhe-2A - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P , NUDT National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808
11/2014	1	Tianhe-2A - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P , NUDT National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808

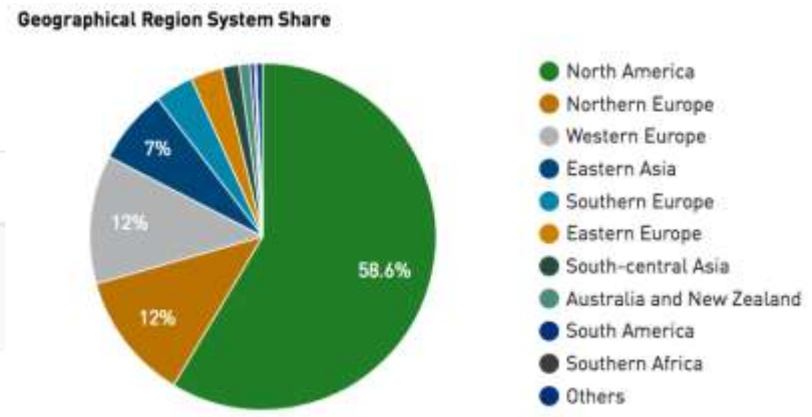

Nov 2008

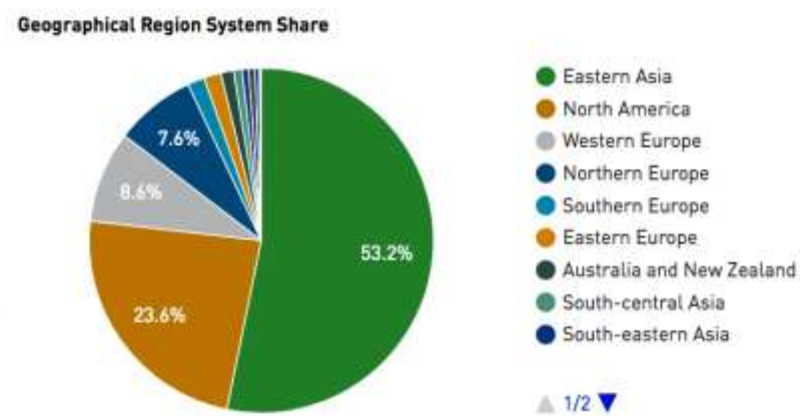

Nov 2018

data from www.top500.org

Collateral Consequences /1

- Growing of computer capability is achieved increasing computer complexity

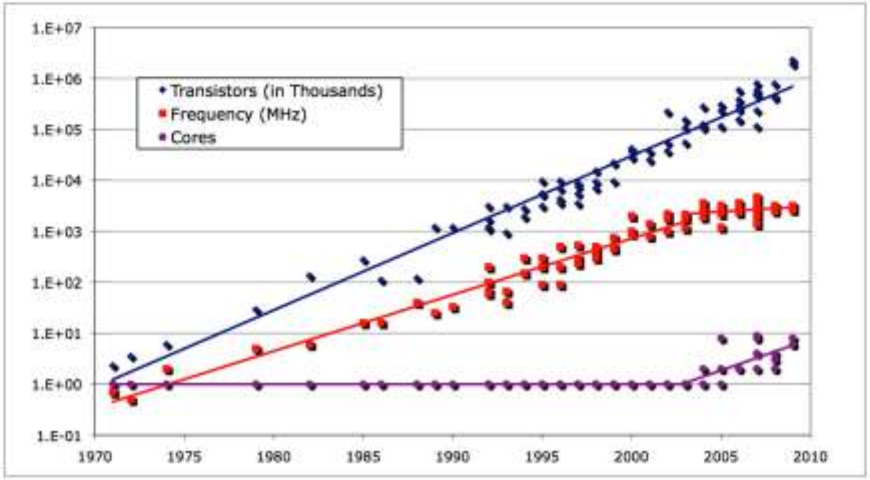

- CPU power is measured in number of floating point operations x second (FLOPs)
  - FLOPS = #cores x clock freq. x (FLOP/cycle)

#cores	Vector Length	Freq. (GHz)	GFLOPs
1	1	1.0	1
1	16	1.0	16
10	1	1.0	10
10	16	1.0	160

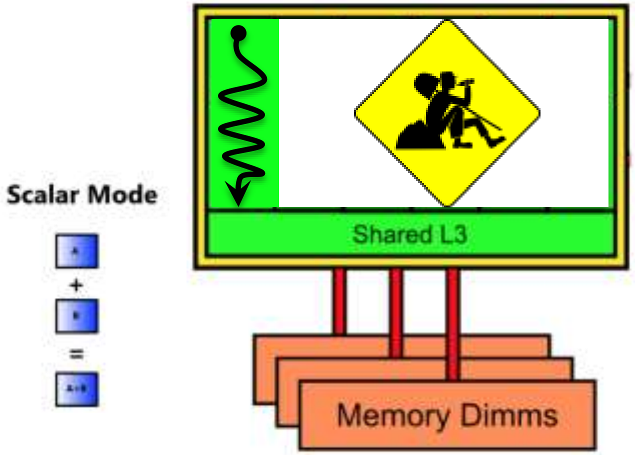

Collateral Consequences /2

- When all CPU component work at maximum speed that is called *peak of performance*
  - Tech-spec normally describe the theoretical peak
  - Benchmarks measure the real peak
  - Applications show the real performance value
- CPU performance is measured FLOP/s
- But the real performance is in many cases mostly related to the memory bandwidth (Bytes/s)
- The way data are stored in memory is a key-aspect for high performance

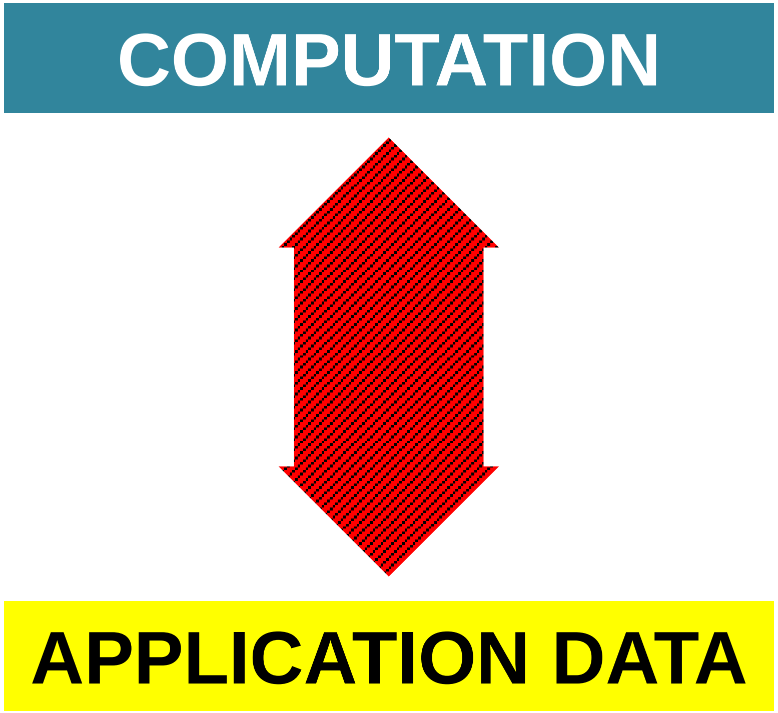

Collateral Consequences /2

- Complexity of physical models is directly proportional to the software complexity
- Number of operations aa well as the size of the problem (data) grows extremely quickly when increasing the size of a 3D (multidimensional) domain

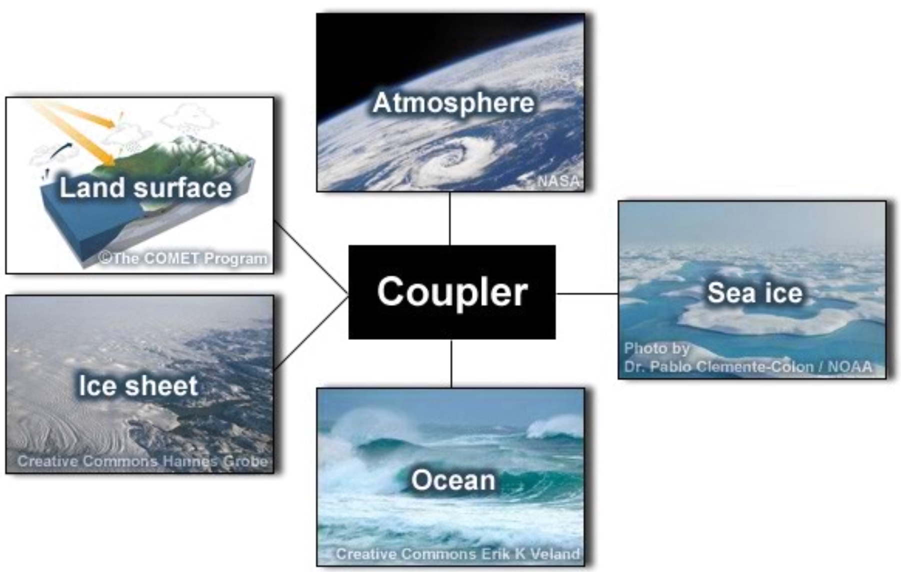

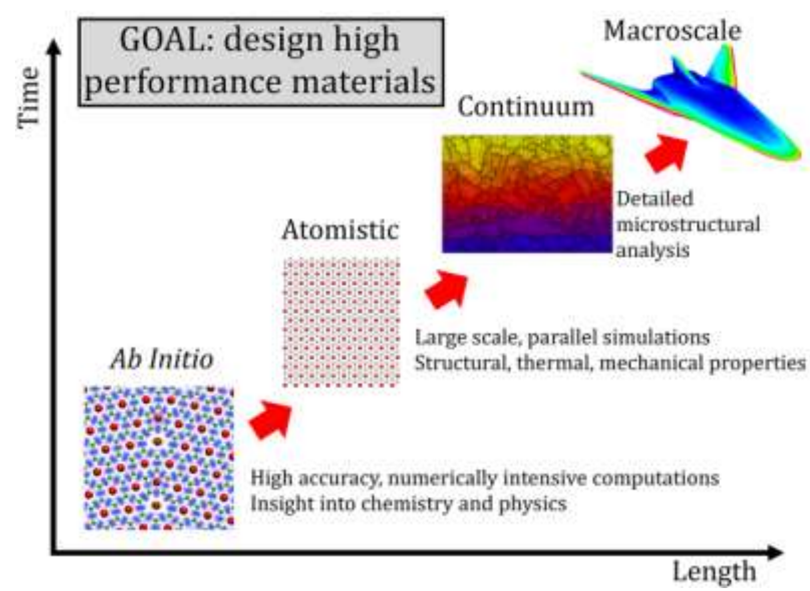

https://www.nas.nasa.gov/SC14/demos/demo26.html#prettyPhoto

Collateral Consequences /3

- No longer a stand-alone project of translating formulations in a computer code from scratch

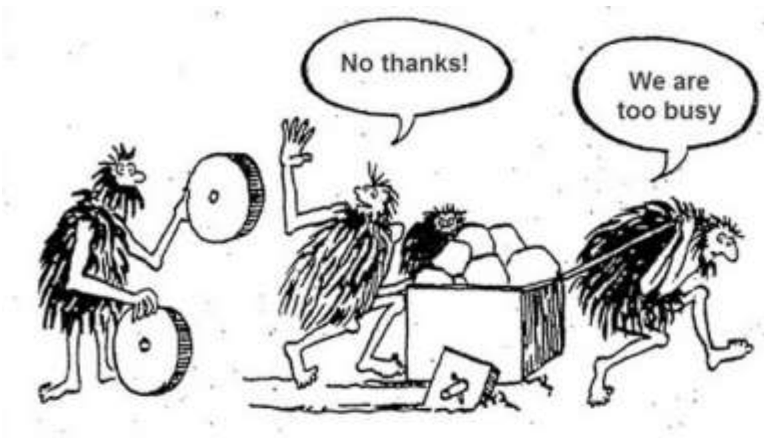

- A huge amount of software is freely available (mostly open source)
- It is matter to use it efficiently and/or make it better
- A collaborative effort of development

GitLab

Collaborative Development

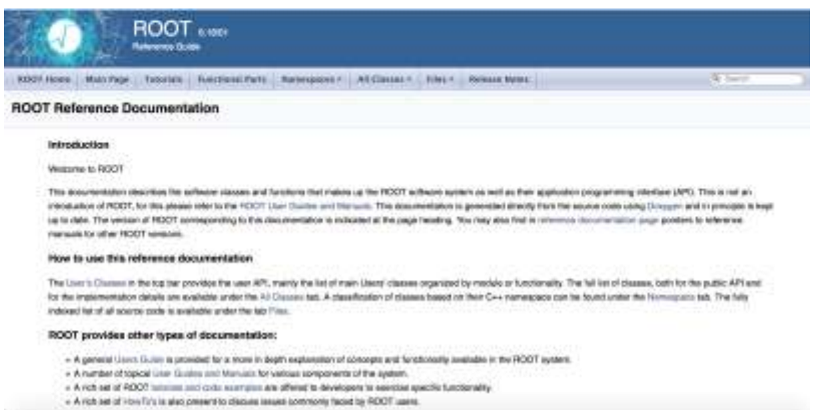

Collateral Consequences /5

- The components of a ecosystem that must grow in concert to make large scale scientific challenges suitable and doable

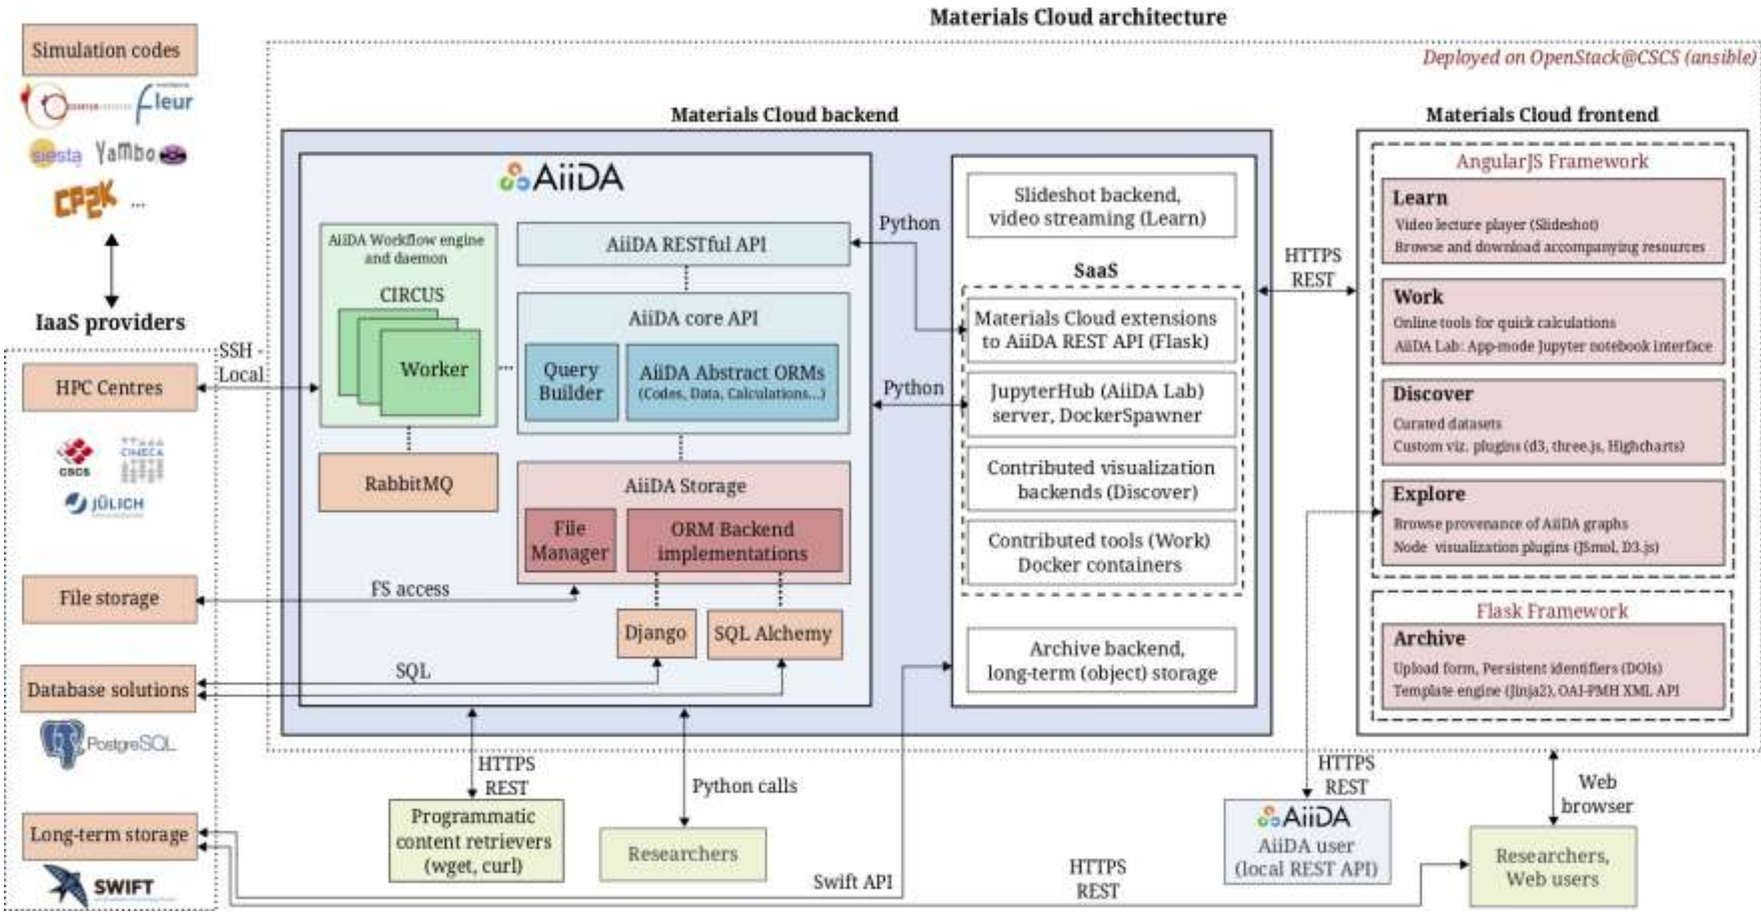

*** Curtesy of Prof. Nicola Marzari (EPFL)**

Workflow of Parallel Scientific Applications

- Data Assimilation
- Pre-processing
- Simulation
- Post-Processing
- Visualization
- Data Publication

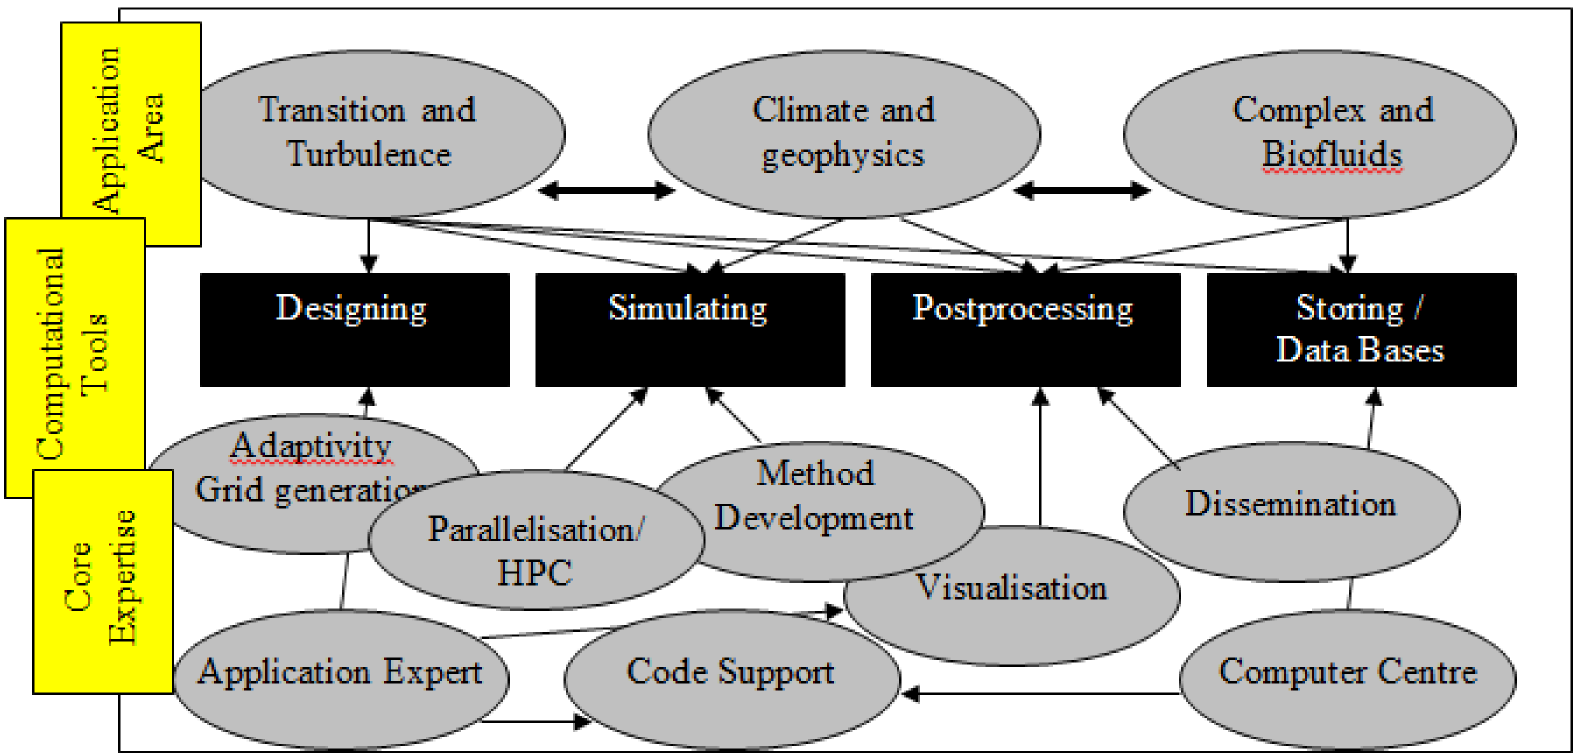

Conventional Software Development Process

- Start with set of requirements defined by customer (or management):
  - features, properties, boundary conditions
- Typical Strategy:
  - Decide on overall approach on implementation
  - Translate requirements into individual subtasks
  - Use project management methodology to enforce timeline for implementation, validation and delivery
- Close project when requirements are met

What is Different in the Scientific Software Development Process?

- Requirements often are not that well defined
- Floating-point math limitations and the chaotic nature of some solutions complicate validation
- An application may only be needed once
- Few scientists are programmers (or managers)
- Often projects are implemented by students (inexperienced in science and programming)
- Correctness of results is a primary concern, less so the quality of the implementation
- In most cases not driven by specific investments but part of the research activity

Complexity of software

Many scientific applications are several orders of magnitude larger than everything you have probably ever seen!

- For example, a crude measure of complexity is the number of lines of code in a package (as of 2018):
  - Deal.II has 1.1M
  - PETSc has 720k
  - Trilinos has 3.3M
- At this scale, software development does not work the same as for small projects:
  - No single person has a global overview
  - There are many years of work in such packages
  - No person can remember even the code they wrote
- Computers become more powerful all the time and more complex problems can be addressed
- Solving complex problems requires combining expertise from multiple domains or disciplines
- Use of computational tools becomes common among non-developers and non-theorists
  - many users could not implement the whole applications that they are using by themselves
- Current hardware trends (SIMD, NUMA, GPU) make writing efficient software complicated

Complexity of software

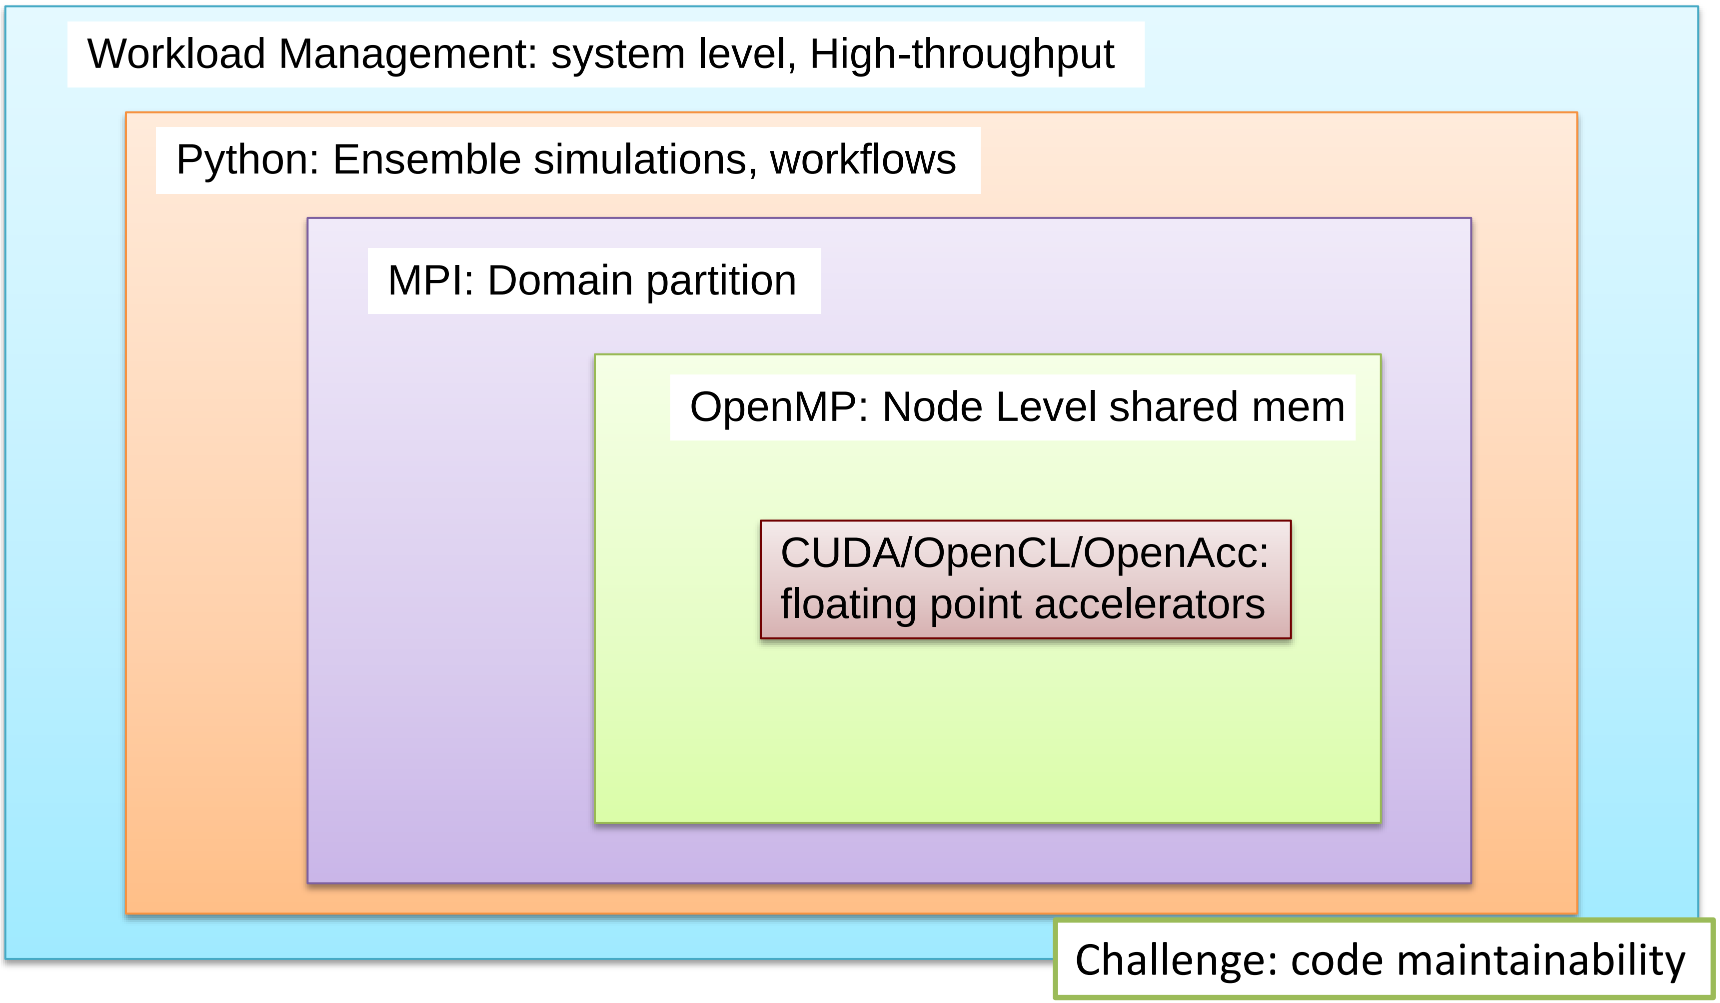

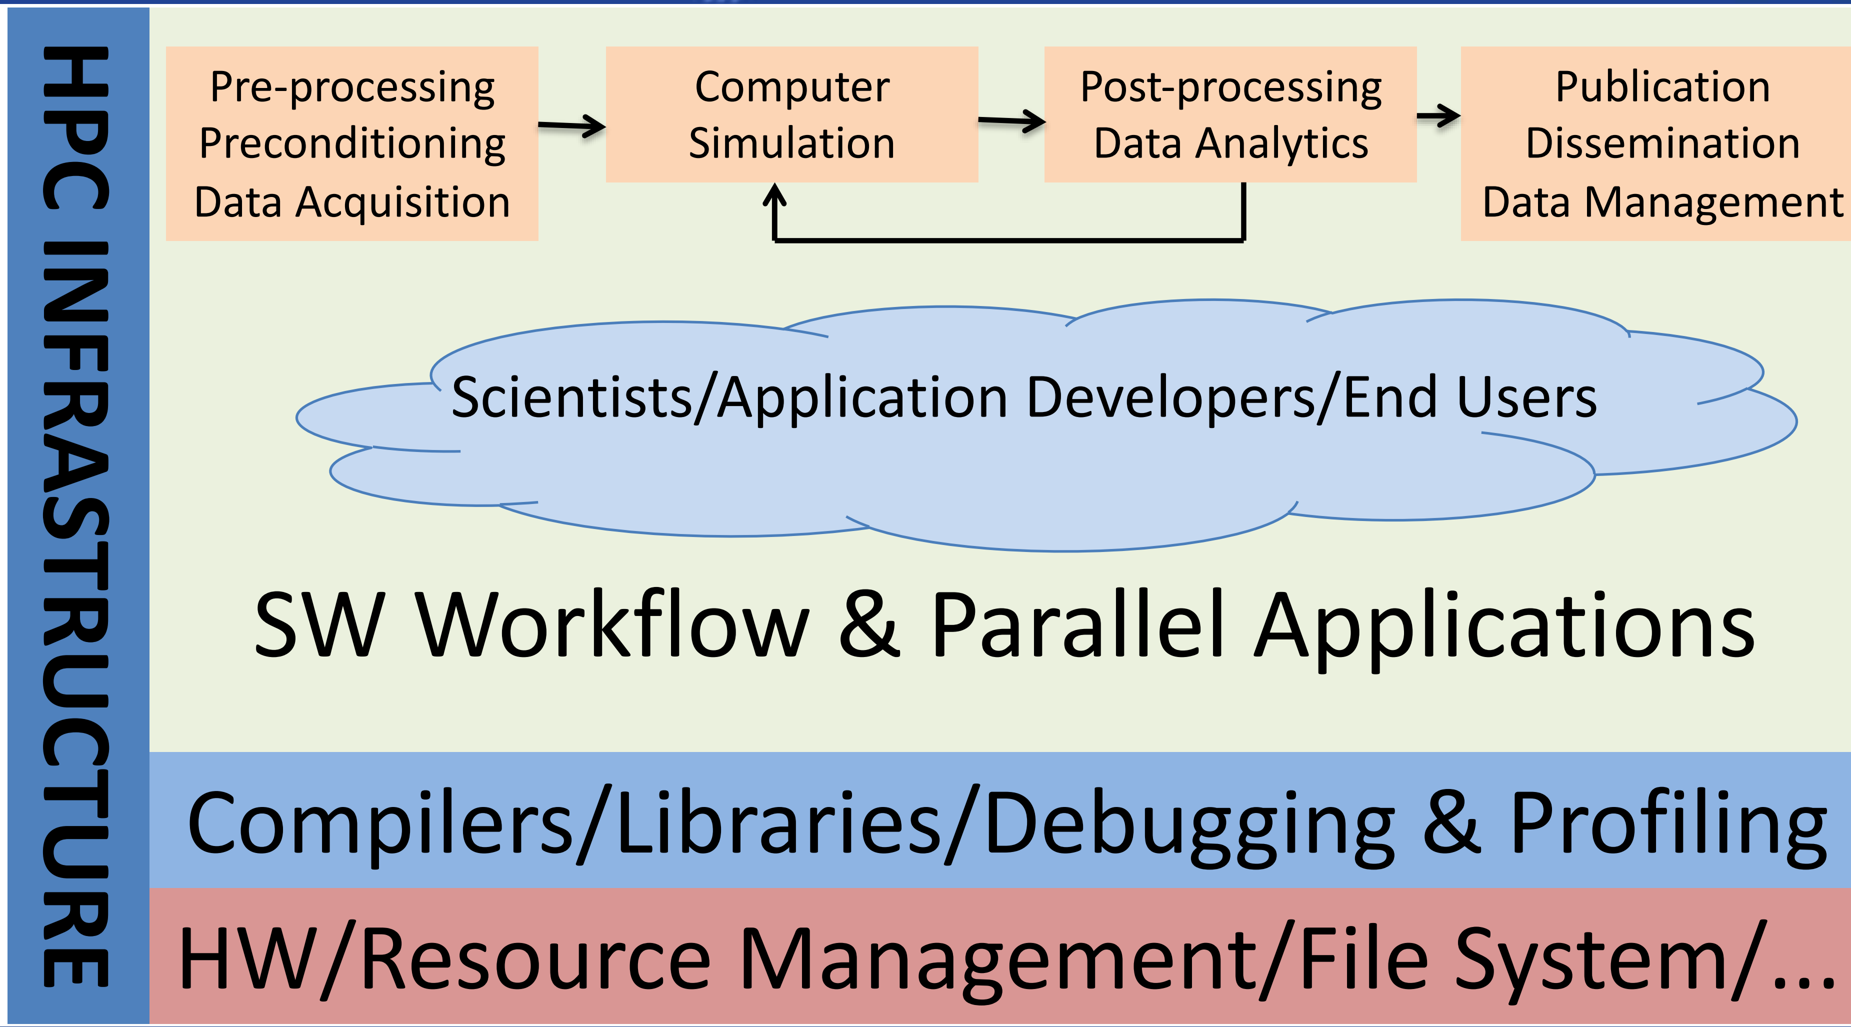
HPC INFRASTRUCTURE
Pre-processing Preconditioning Data Acquisition
Computer Simulation
Post-processing Data Analytics
Publication Dissemination Data Management
Scientists/Application Developers/End Users
SW Workflow & Parallel Applications
Compilers/Libraries/Debugging & Profiling
HW/Resource Management/File System/...

Tech Support to HPC Infrastructures

- No Users
  - Make HPC visible, Documentation, HPC Dissemination and Training
  - The community must first understand the benefit
- non-Expert Users
  - Specific support to software for the whole production, from software building to parallel simulations
  - Requires a really close collaboration and patience
  - At the frequent rising of problems they might give up
- Expert users
  - Drive the software environment
  - Require highly specialized support for:
    - large scale simulations
    - Software optimization, porting to high-end technology, perf. analysis

The Essential

- Documentation
  - How to access, software, job monitoring and execution, brief description of the infrastructure, quotas, tech. contacts
- Compilers and MPI library
- Scripting tools: Python, R
- Building tools: cmake, autotools
- Scientific tools: gnuplot

Math Libraries

- Scalable Parallel Random Number Generators Library (SPRNG)
- Parallel Linear Algebra (ScaLAPACK)
- Parallel Library for Solution of Finite Elements (dealii)
- Parallel Library for FFT (FFTW)
- Parallel Linear Solver for Sparce Matrices (PETSc)

Formatted data libraries

- Most scientific communities have defined today a protocol to describe their data (formatted data)
- Based on generic libraries: HDF5, NetCDF, etc…
- But also more specific (i.e., SEG-Y)
- Most implement parallel I/O
- Formatted data provide the opportunity to scientific data visualization and publication

Task Farming

- I am working on an embarrassing parallel problem
- Work is divided in independent tasks (no communication) that can be performed in parallel
- The same program (set of instructions) among different data: same model adopted by the MPI library
- A parallel tool is needed to handle the different processes working in parallel
- The MPI library provides the *mpirun* application to execute parallel instances of the same program
- Quite common in Computer Graphics, Bioinformatics, Genomics, HEP, anything else requiring processing of large data-set, sampling, ensemble modeling

Task Farming

```
$ mpirun -np 12 my_program.x
```

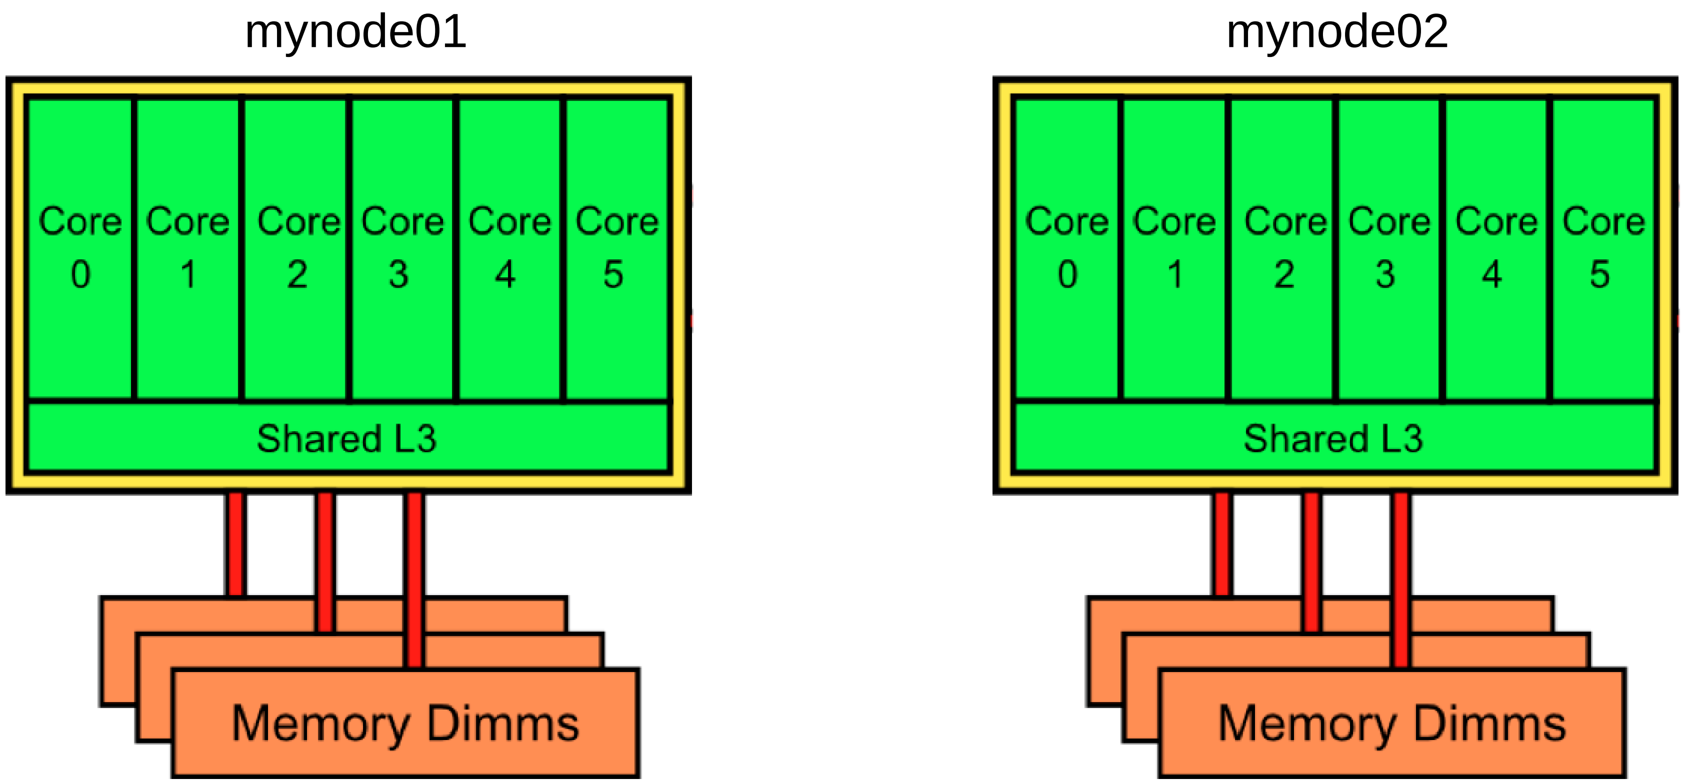

Task Farming

```
[igirotto@mynode01 ~]$ mpirun -np 12 /bin/hostname
mynode01
mynode02
mynode01
mynode02
mynode01
mynode02
mynode01
mynode02
mynode01
mynode02
mynode01
mynode02
```

PATH name common to all processes !!

Task Farming

In Python

```
import os
myid = os.environ['OMPI_COMM_WORLD_RANK']
[...]
```

In BASH

```
#!/bin/bash
myid=${OMPI_COMM_WORLD_RANK}
[...]
```

```
[igirotto@mynode01 ~]$ mpirun ./myprogram.[py/sh...]
```

Task Farming: conclusion

- Executing multiple instances on the same program with different inputs/initial cond.
- Reading large binary files by splitting the workload among processes
- Searching elements on large data-sets
- Other parallel execution of embarrassingly parallel problem (no communication among tasks)
- Task Farming is a simple model to parallelize simple problems that can be divided in independent task
- The *mpirun* application aids to easily perform multiple processes, includes environment setting

Argo

ARGO: used CPU hours x month

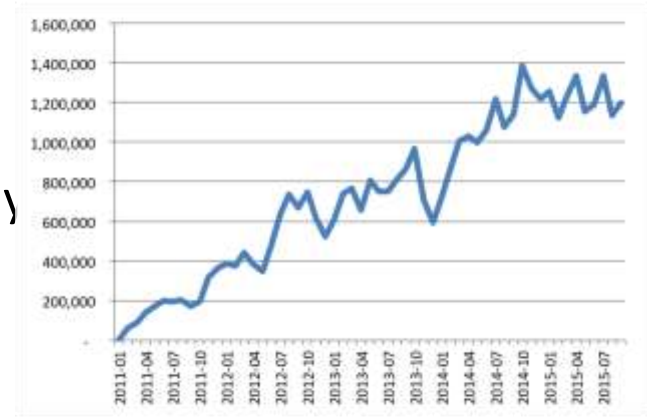

In-house HPC cluster “ARGO”

- Heterogeneous system with about 200 nodes
- ~3000 processor core of 4 different Intel Prod. Family
  - Westmere, SandyBridge, IvyBridge, Broadwell
- Nominal power: ~100 TeraFLOP
- Storage: ~1.5PB of Usable storage

Objective:

- Provide in-house resources and manage sw configuration and user access directly
- Maintain in depth HPC knowledge in house
- Flexible usage for the ICTP needs
- Open for selected scientists from the developing world
- ICTP schools and workshops

Tier-0 world-class HPC resources

Marconi @ CINECA, is equipped with thousands of compute nodes. ICTP scientists access through an agreement between the institutions, national grants and PRACE grants.

Large Scale Computer Simulations @ ICTP

GFDL/MOM global ocean model as part of the international program FAFMIP. Simulations runs on 2400 cores (~50 nodes) on the SKL partition

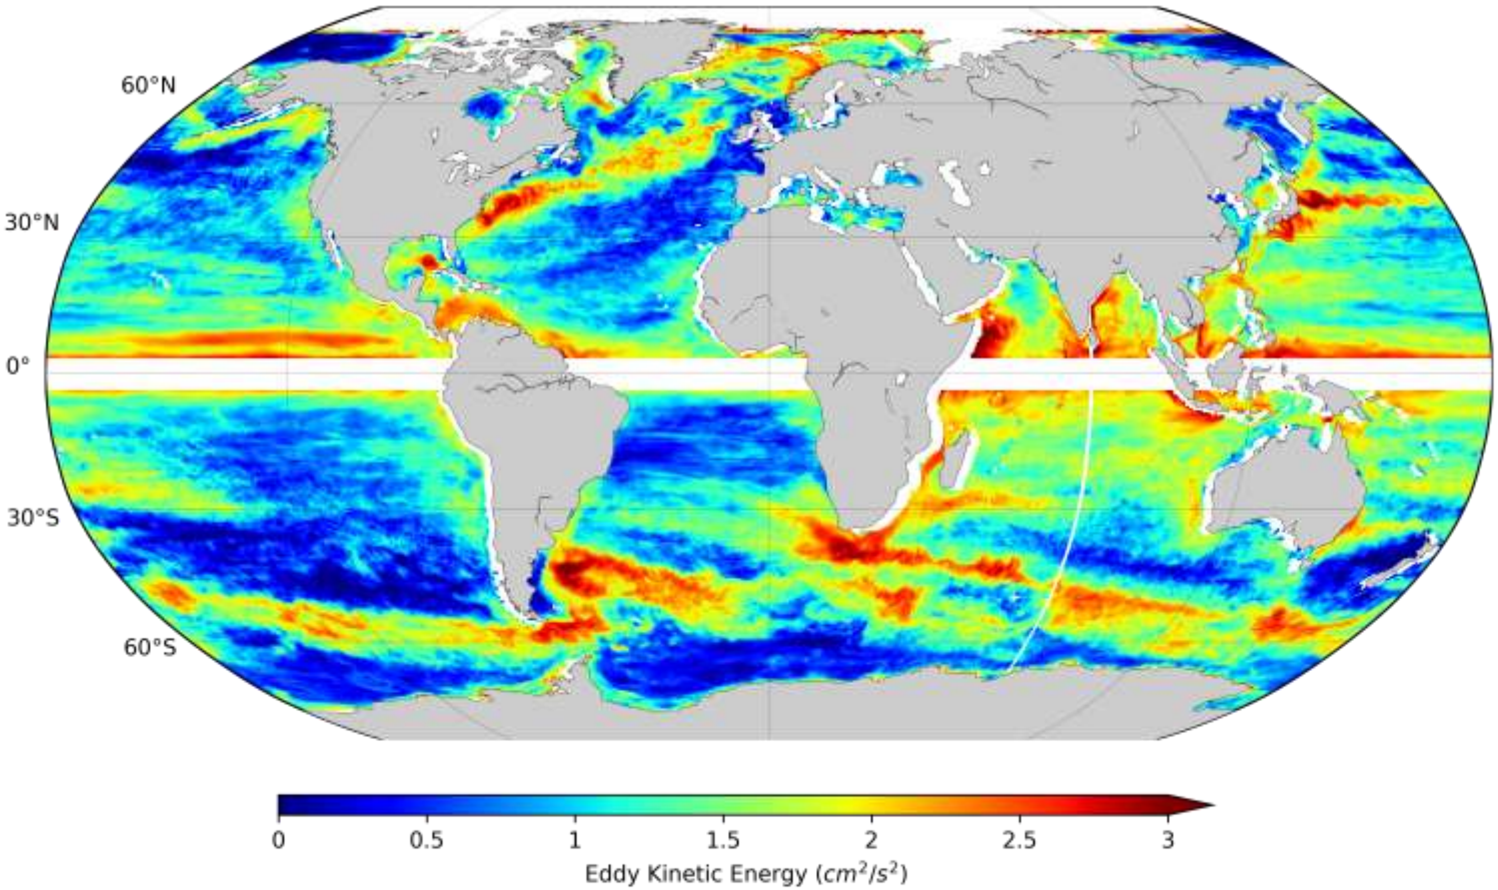

*** Courtesy of Dr. Ricardo Farneti (ESP group)**

Large Scale Computer Simulations @ ICTP

Solving numerically a wave-like PDE on a large 3d grid (up to 4096^3)

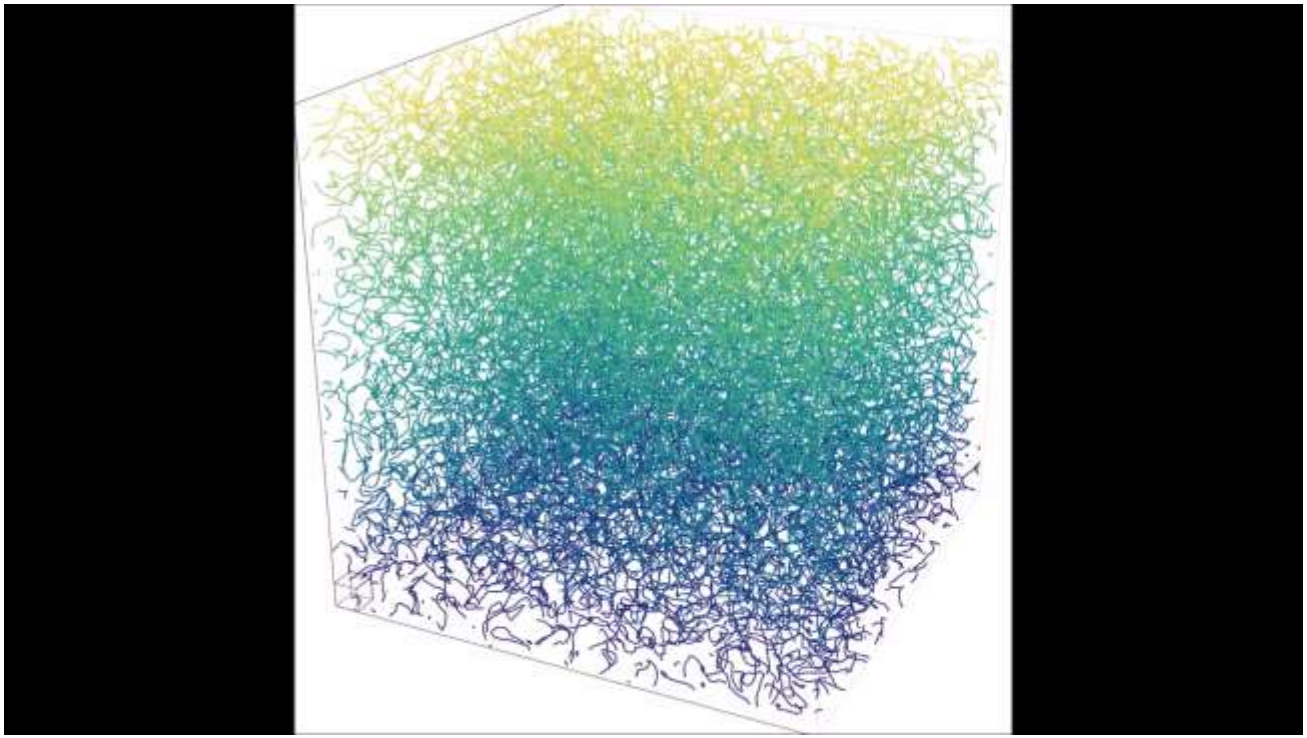

Journal of High Energy Physics
July 2018, 2018:151 | Cite as

Axions from strings: the attractive solution

Authors | Authors and affiliations

Marco Gorghetto, Edward Hardy, Giovanni Villadoro

Large Scale Computer Simulations @ ICTP

Understanding the electrochemical double layer at the hematite/water interface: A first principles molecular dynamics study

Cite as: J. Chem. Phys. **150**, 041707 (2019); https://doi.org/10.1063/1.5047930
Submitted: 10 July 2018 . Accepted: 06 September 2018 . Published Online: 29 November 2018

Kanchan Ulman, Emiliano Poli, Nicola Seriani, Simone Piccinin, and Ralph Gebauer

heavy ab-initio MD calculations!!

QUANTUMESPRESSO

nature COMMUNICATIONS

ARTICLE

DOI: 10.1038/s41467-018-07190-1 OPEN

Vibrational fingerprint of localized excitons in a two-dimensional metal-organic crystal

M. Corva[1,2], A. Ferrari[1,5], M. Rinaldi[1], Z. Feng[1,6], M. Roiaz[3], C. Rameshan[3], G. Rupprechter[3], R. Costantini[1,2], M. Dell'Angela[2], G. Pastore[1], G. Comelli[1,2], N. Seriani[4] & E. Vesselli[1,2]

Diagonalization of really big matrixes!

Computer Physics Communications
Volume 235, February 2019, Pages 477-488

Massively parallel implementation and approaches to simulate quantum dynamics using Krylov subspace techniques ☆

Marlon Brenes[a, b], Vipin Kerala Varma[a, c, d, e], Antonello Scardicchio[a, f], Ivan Girotto[a, g, h]
⊞ Show more

https://doi.org/10.1016/j.cpc.2018.08.010 Get rights and content

arXiv.org > cond-mat > arXiv:1902.09236

Condensed Matter > Strongly Correlated Electrons

Non-Abelian symmetries and disorder: a broad non-ergodic regime and anomalous thermalization

I. V. Protopopov, R. K. Panda, T. Parolini, A. Scardicchio, E. Demler, D. A. Abanin

(Submitted on 25 Feb 2019)

PHYSICAL REVIEW LETTERS

Many-Body Localization Dynamics from Gauge Invariance

Marlon Brenes, Marcello Dalmonte, Markus Heyl, and Antonello Scardicchio
Phys. Rev. Lett. **120**, 030601 – Published 19 January 2018

Large Scale Computer Simulations @ ICTP

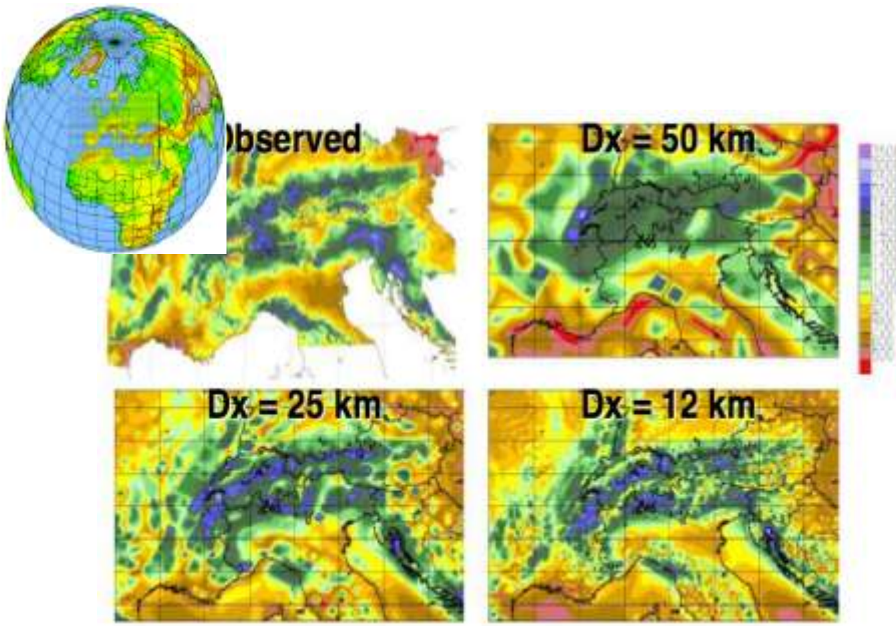

- ICTP contributes to the IPCC report
- The RegCM code is developed and maintained by the ESP group @ ICTP
- The model is performed on several domains covering most of the world’s land

[…] for Europe, we plan to run seven 140 year simulations. Each month of simulation, with 500 processors, is expected to take about 2.9 hours of compute time and produce 60GB of raw output files

<u>The on-going simulations, planned to end in early 2020, is expected to finally require about 100M cpu-h and approximately 2PB of data will be publicly made available</u>

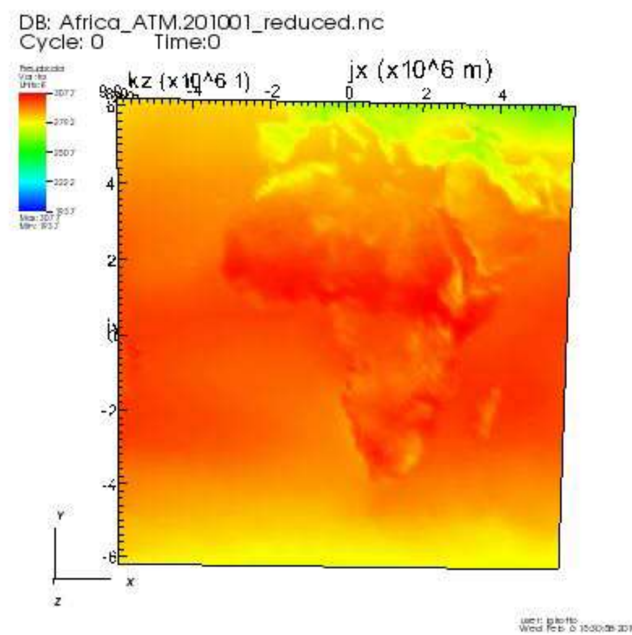

RegCM4.7 tested domain available and long scenario simulations available

*** Courtesy of Dr. Erika Coppola (ESP group)**

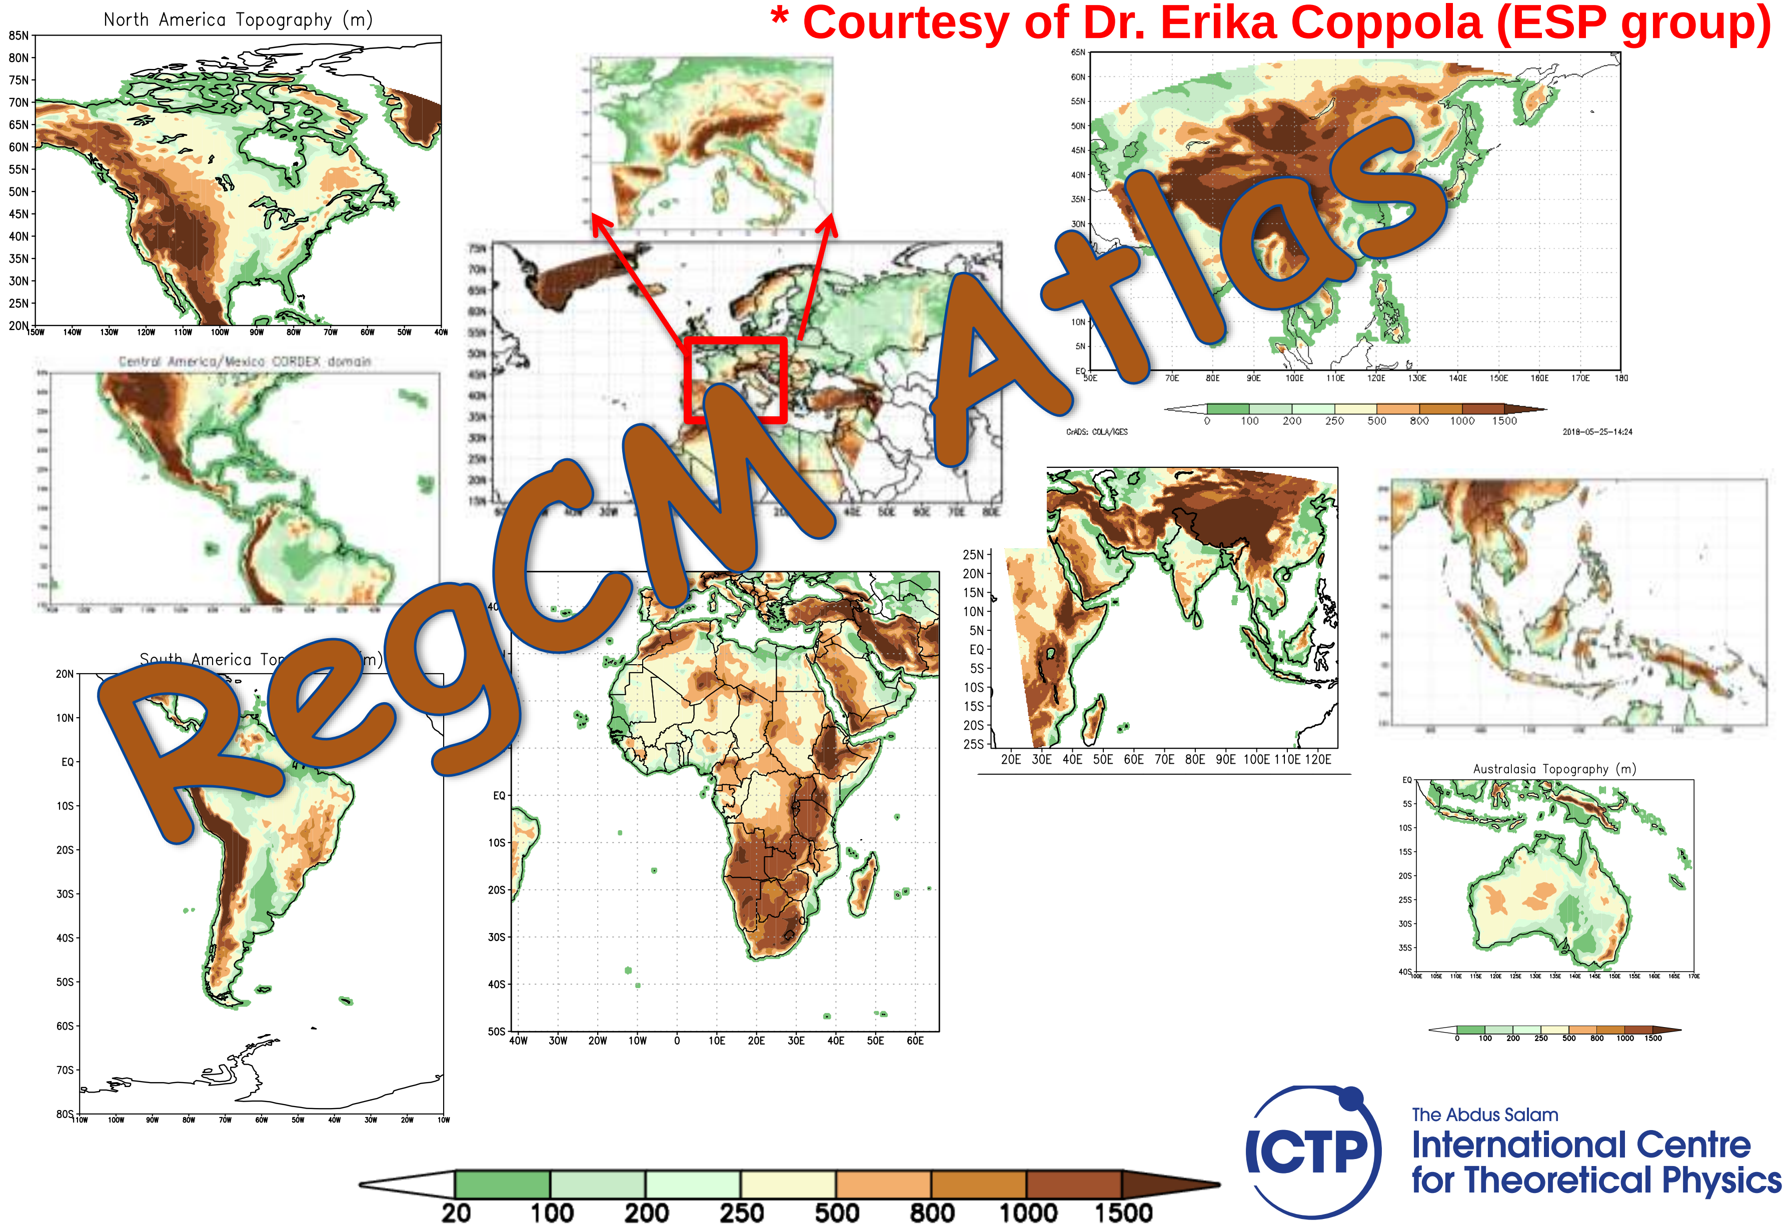

The Abdus Salam International Centre for Theoretical Physics

Data Research Management

- Handling and management of scientific data
- The strategy to make data, secure, safe and available in the long term
- Managers and administrators are responsible to define and implement the plan
- Scientists are responsible to follow the rules
- The only way to ensure sustainability when considering really large amount of data

Data Research Management

- Data flow
  - Workflow of the data
  - How frequent and how fast we need the data
  - Critical aspect for efficient productivity, costs reduction, data accessibility and usability
- Data life
  - how long do we need to keep a data alive

Workflow Mapping on Infrastructure

- High-Perf storage
  - for production and intense data analysis
  - DDN, NetAPP, Panasas, etc…
- Low perf-storage
  - For post production, low frequency access, archiving
  - Normally a huge bunch of disks raked somewhere
  - Synology, …
- Costs Vs Performance Vs Capacities
  - there are several orders of magnitude difference among the two categories among all aspects

At the ICTP

Name	Storage Capacity	Data life	Kind	File System
Home	Few tens of GBs	User account	NetAPP	NFS
User-data	5 to 10 TBs	User account	NetAPP	NFS
Scratch	Unlimited (up to storage capacity)	1 month	NetAPP	NFS
Archive	Few tens of TBs	Several years	Synology	NFS

Conclusions

- Scientific Computing is NOT an IT service
  - Few standards, often requires ad-hoc solutions
  - Scientific software for high-performance computing evolves daily
- Most scientists do not master the complexity of modern scientific software
  - There is high probability things do not work as expected
  - A problem on the application becomes also your problem
- The infrastructure must work smoothly and efficiently
  - High number of possible point of failures (inexperience, infrastructure, building, run time configuration, etc…)
- There are no customers, but mostly seen by users as a service to scientific production
- No users => no science => waste of money

The MHPC in pills: www.mhpc.it

- High-level educational program, beyond M.Sc.
- Intensive training aimed to build knowledge in solving complex problems with an HPC approach
- Innovative, hands-on based training
- 15 students x year
- since the first edition 100% employees after the program

MHPC
Master in High Performance Computing

Background Requirements

- Candidates must have some experience in programming and a competence in at least one of the languages between C, C++ and/or Fortran
  - Python knowledge is a plus
- A sound knowledge of Linux operating system
- Master level of a scientific degree is required
- No prior HPC knowledge is assumed
- Enthusiasm is a must

The Curricula : www.mhpc.it

1 year program divided in 6-8 months courses and 6 month project (some overlap)

Mandatory

- Scientific Programming Environment
- Introduction to Computer Architectures for HPC
- Object Oriented Programming
- Parallel Programming
- Introduction to Numerical Analysis
- Advanced Computer Architectures and Software Optimizations
- Parallel Data Management and Data Exchange
- High Performance Computing Technology
- Best Practices in Scientific Computing

Optional Choice

- Data Structures & Sorting and Searching Electronic structure: from blackboard to source code
- The Finite Element Method Using deal.II
- Reduced Basis Methods Fast Fourier Transforms in Parallel and Multiple Dimensions
- Cluster Analysis
- Monte Carlo methods
- Supervised & Unsupervised Machine Learning
- Machine Learning
- Deep Learning
- Approximation and interpolation of simple and complex functions
- Spatial locality algorithms
- Lattice Boltzmann
- Molecular dynamics

HPC Training Scholarship Winners

Alejandra Foggia
& Rajat Panda
CMSP group @ ICTP

Muhammad Owais, Pakistan
Junior Research Engineer
BSC-CNS (Spain)

Marlon Brenes Navarro, Costa Rica
PhD Student
Trinity College Dublin (Ireland)

James Vance, Philippines
PhD Student
Johannes Gutenberg University Mainz and the Max Planck Institute for Polymer Research(Ghana)

Michael O. Atambo, Kenia
Last month PhD in "Physics and Nanosciences"
CNR-NANO (Italy)

Jimmy Aguilar Mena, Cuba
PhD Student
BSC-CNS (Spain)

Elliot Menkha, Ghana
Last Month PhD in “Computational Chemistry”
Kwame Nkrumah University of Science and Technology (Ghana)

Anoop Chandran, India
PhD Student
Institute for Advanced Simulation, Julich (Germany)

Fernando Posada, Colombia
Assistant Professor
Temple University (USA)

DATA SCIENCE &
SCIENTIFIC COMPUTING

ABOUT US LEARN WITH US STUDENT SERVICES PARTNERSHIPS CONTACTS

MASTER DEGREE IN

Data Science and Scientific Computing

SISSA

ICTP

ABOUT US

ABOUT DATA SCIENCE AND SCIENTIFIC COMPUTING

The Master degree in Data Science and Scientific Computing (DSSC) aims to train professionals of the future: experts that are highly requested because of the digitalization of the society and the evolution to industry 4.0.

READ MORE→

Data Science and Scientific Computing: https://dssc.units.it

- International M.Sc. joint SISSA/ICTP/UNITS/UNIUD program (2 years)
- Final certificate delivered by the University of Trieste
- One Curricula in Data Science
    - training in the fields of **data management** and **data analysis**, with a particular focus on Big Data
- One Curricula in Computational Science and Engineering
    - computational modelling, optimization, scientific programming, and simulation in the areas of CFD, computational physics, computational chemistry

DATA SCIENCE & SCIENTIFIC COMPUTING

Other Programs

In summary: HPC @ ICTP

1) support to the ICTP scientific community on HPC related projects for delivering world-class scientific research

2) contribute with number of initiatives of education and training, to support those subjects that in the next years to come can drive the research in computational sciences as well as the HPC development, in the developing world