Data and Data-intensive computing challenges in Earth and Universe Sciences

Jean-Pierre Vilotte

Institut des Sciences de l'Univers (CNRS-INSU)
Institut de Physique du Globe de Paris (IPGP)

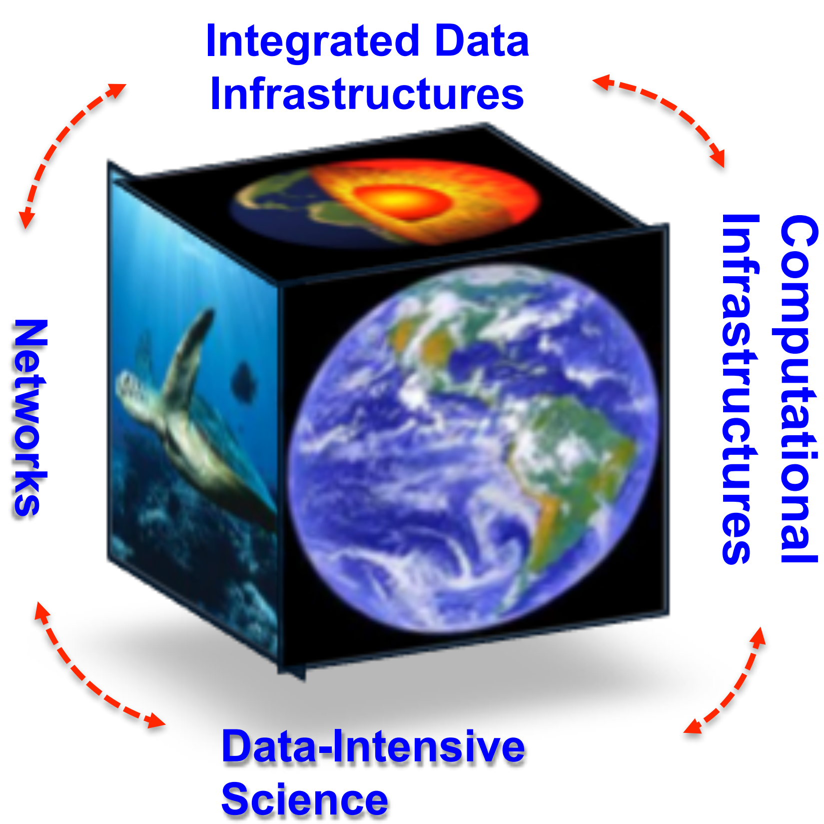

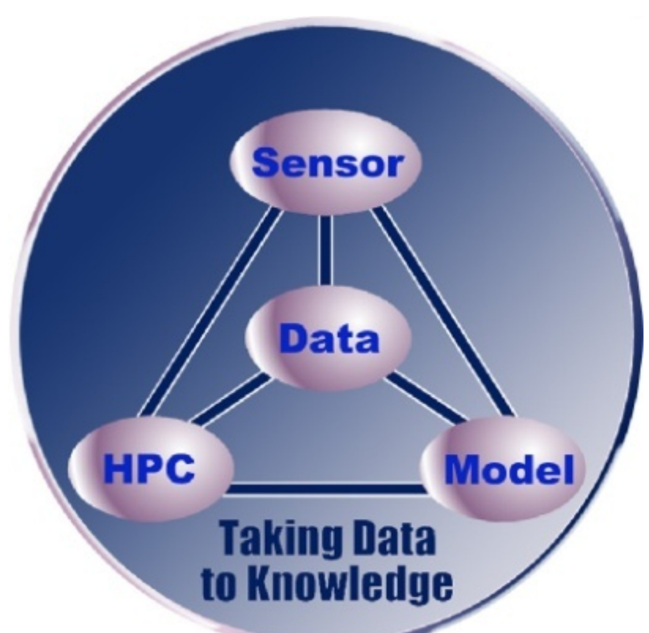

Scientific and Computing Challenges

Drive Scientific discoveries

- Observational data & simulation data
- High-end computational simulation
- Data inversion and Data assimilation
- Statistical data analytics

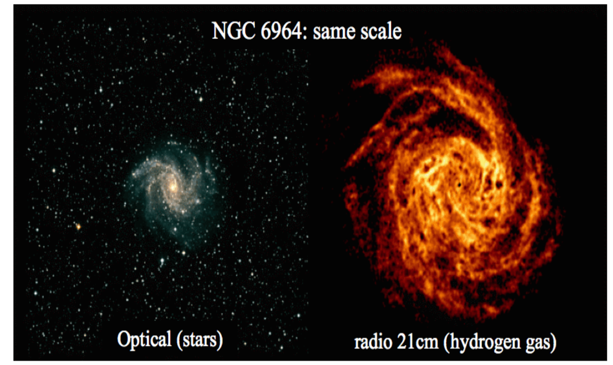

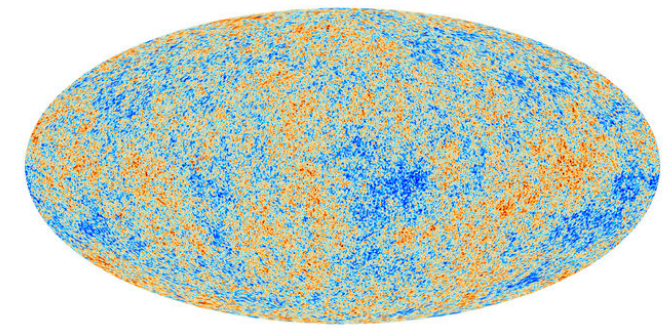

Across multiple disciplines

- Astronomy & Astrophysics
- Climate, Atmosphere, Ocean
- Solid Earth Sciences
- Continental surfaces and interfaces

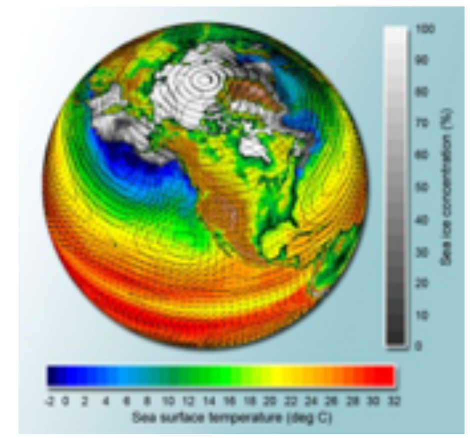

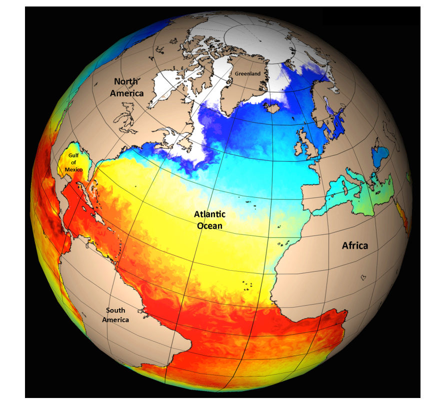

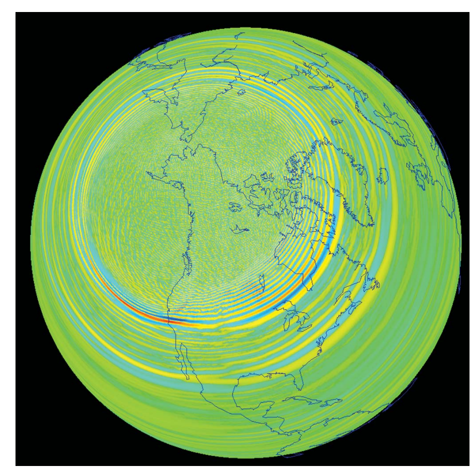

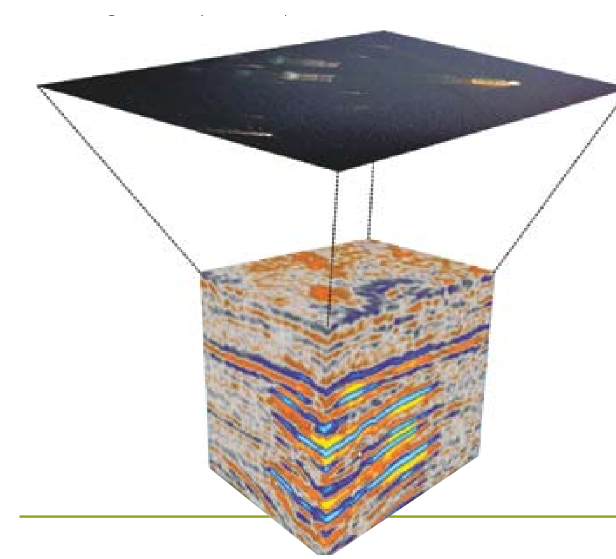

Socio-economical applications

- Climate evolution and forecasting
- Natural hazards (earthquakes, volcanoes, tsunamis, landslides, floods …)
- New energetic resources
- Environmental changes

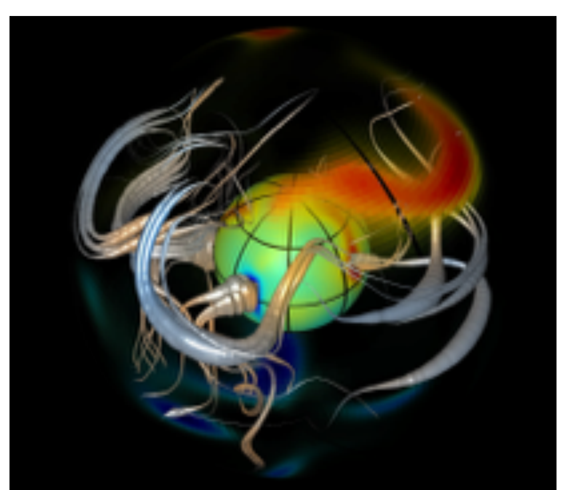

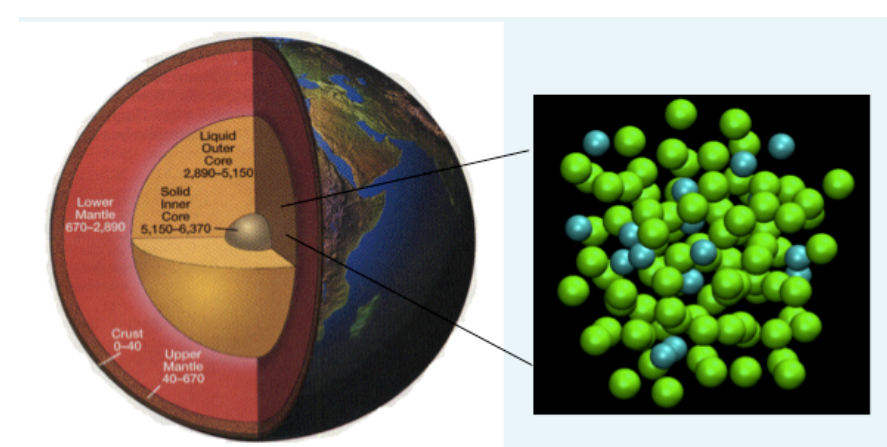

An increasing wealth of data

Ubiquitous data explosion: 100 PBs era

LOFAR/SKA
~10 exabytes/day

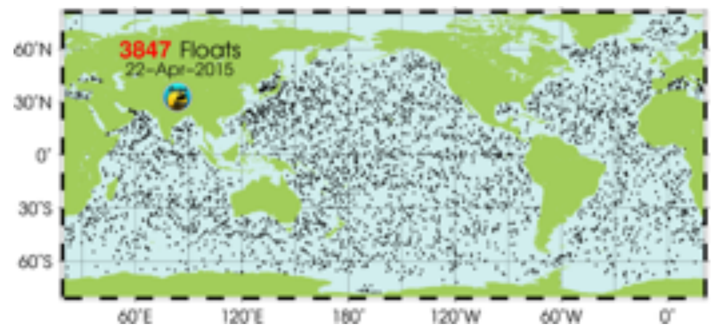

ARGO

Copernicus/Sentinel
SWOT
~4 PBs/day

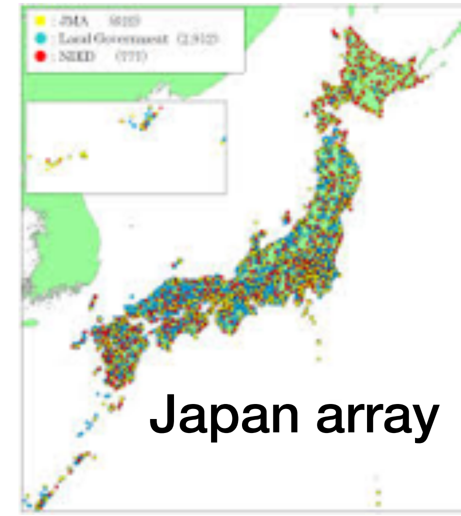

USArray

Large seismic arrays
~100 TBs/year

VIRGO/LIGO

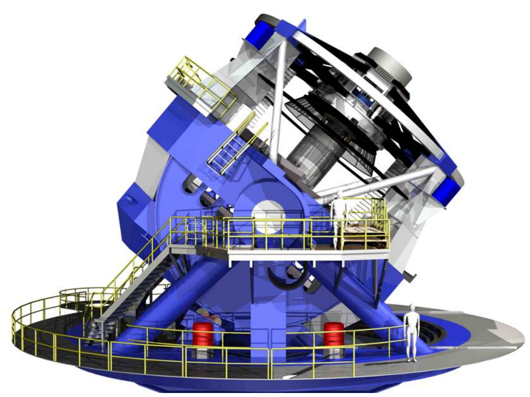

LSST/EUCLID
~20 PBs/night

Data explosion:

- Large throughput instruments; observation and monitoring systems (spatial, land,ocean and ocean bottom) at global and regional scales
- Large HPC simulations

Next generation discoveries:

- **Managing data**: streaming data processing, archiving, curation, metadata, provenance, distribution
- **Data analytics**: statistical streaming data analysis, machine learning methods of high-dimension data
- **Data-intensive simulation**: scalable, resilient large-scale, multi-physics, multi-scales simulations
- **Data-driven inversion and assimilation**: high-dimensional ``Bayesian" inference methods
- **Statistics and stochastic methods**: direct-inverse uncertainties, extreme events statistics

Not a single dimensional challenge

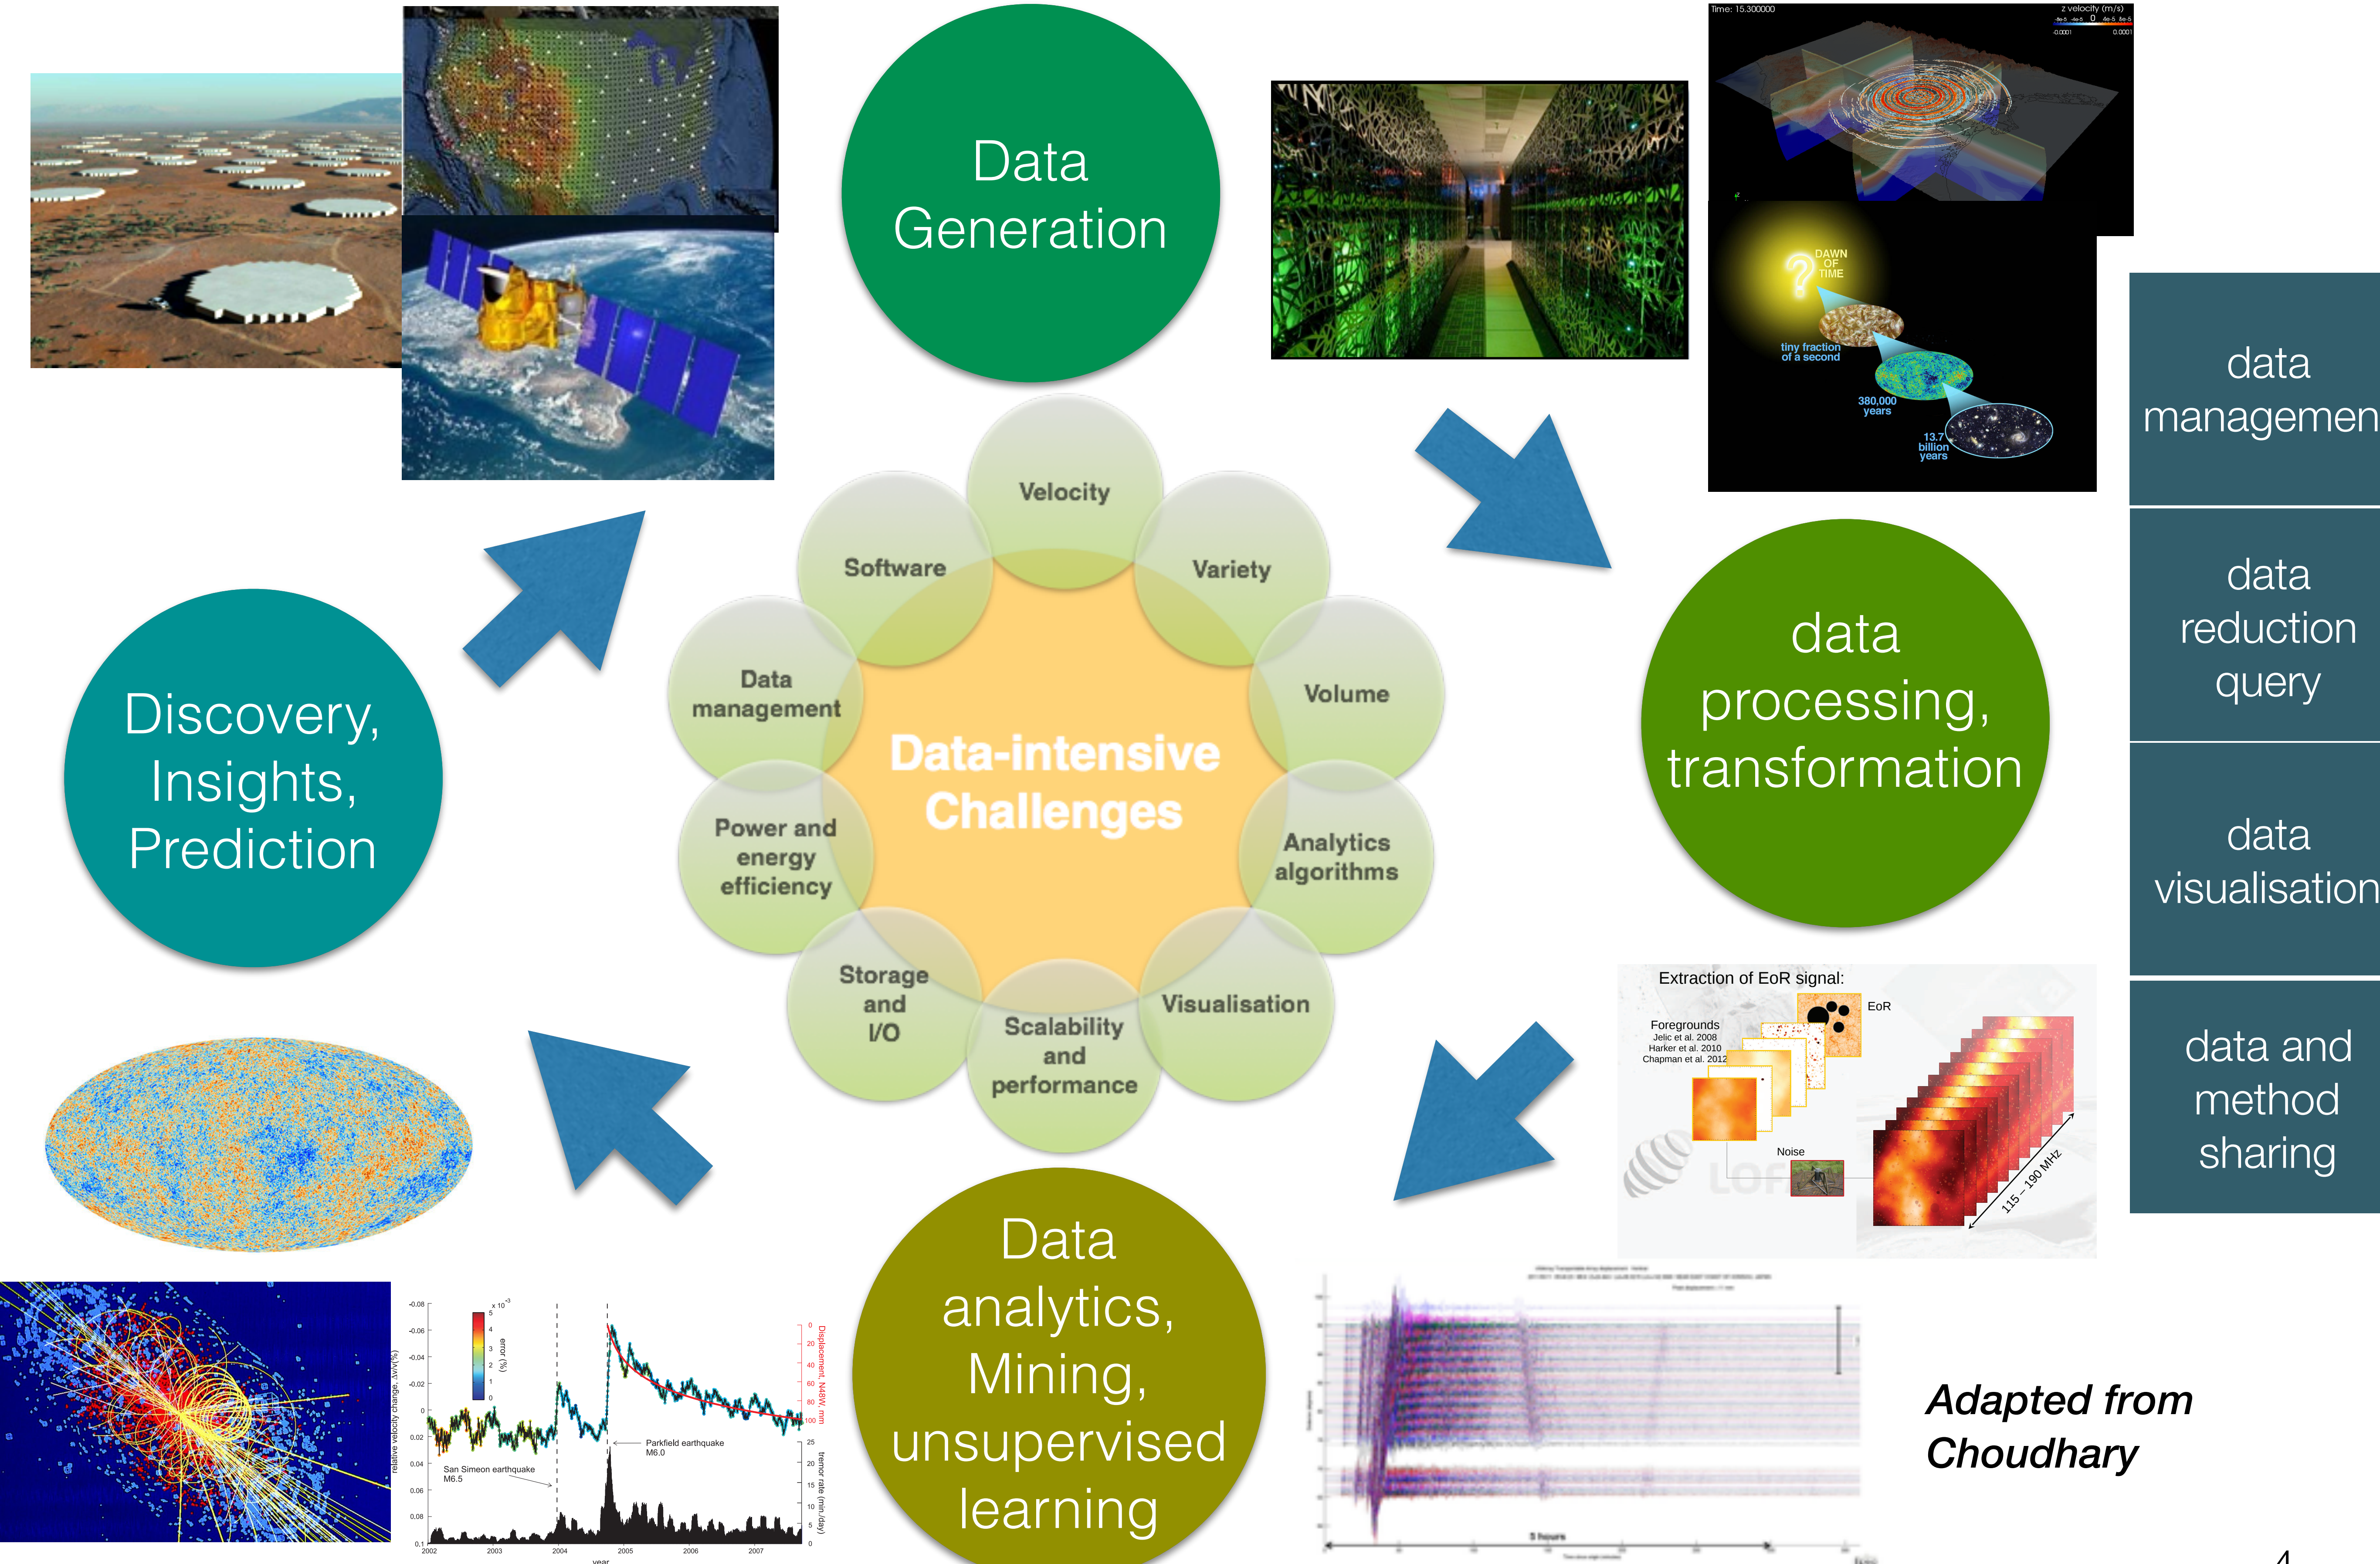

Adapted from Choudhary

Data -> Extraction/cleaning -> Integration/aggregation -> Learning models -> trigger / question -> predict

Big Data statistical analysis

Seismology: data-intensive analysis

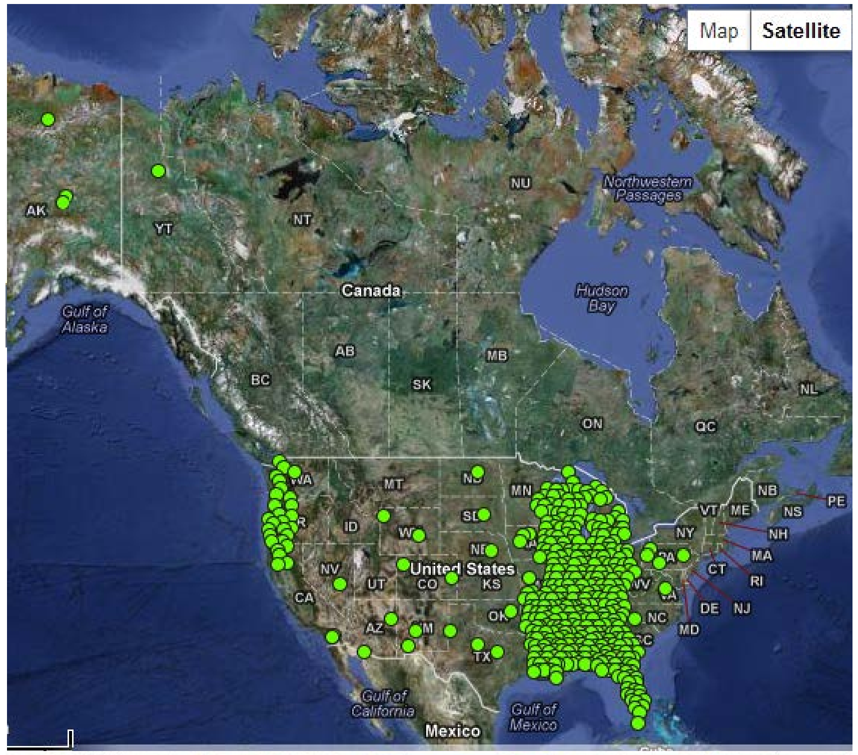

Continuous Waveforms Analysis

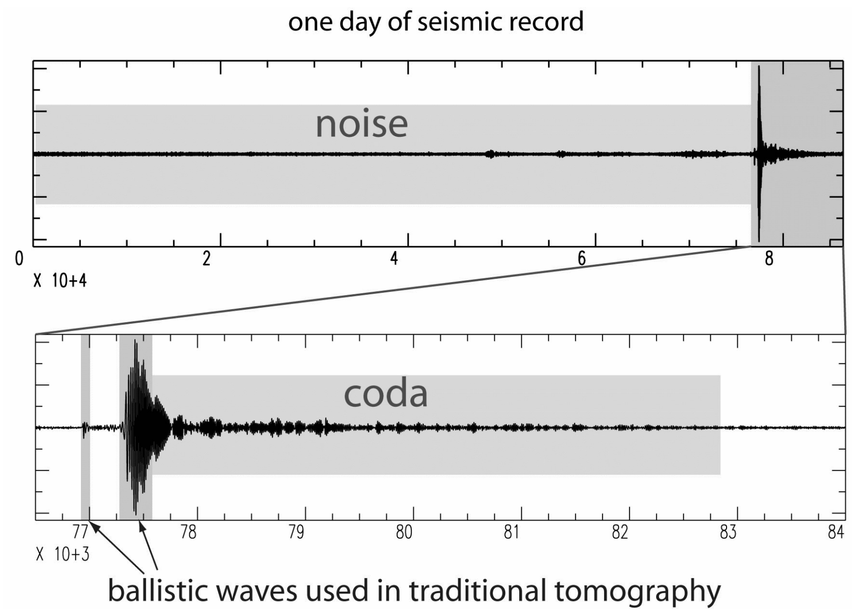

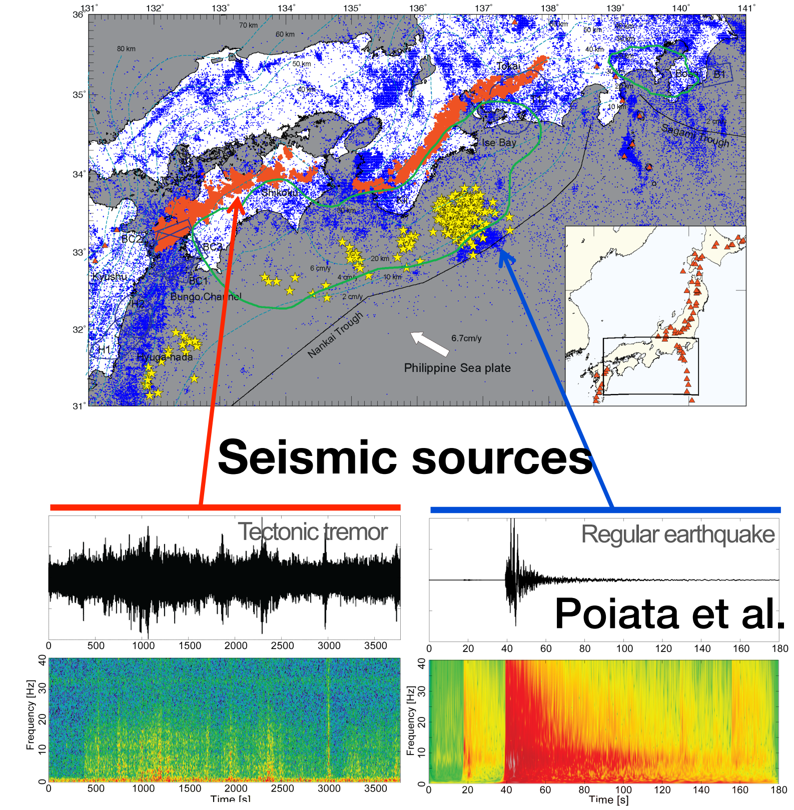

Seismic sources

Poiata et al.

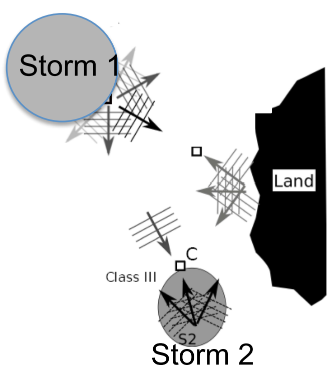

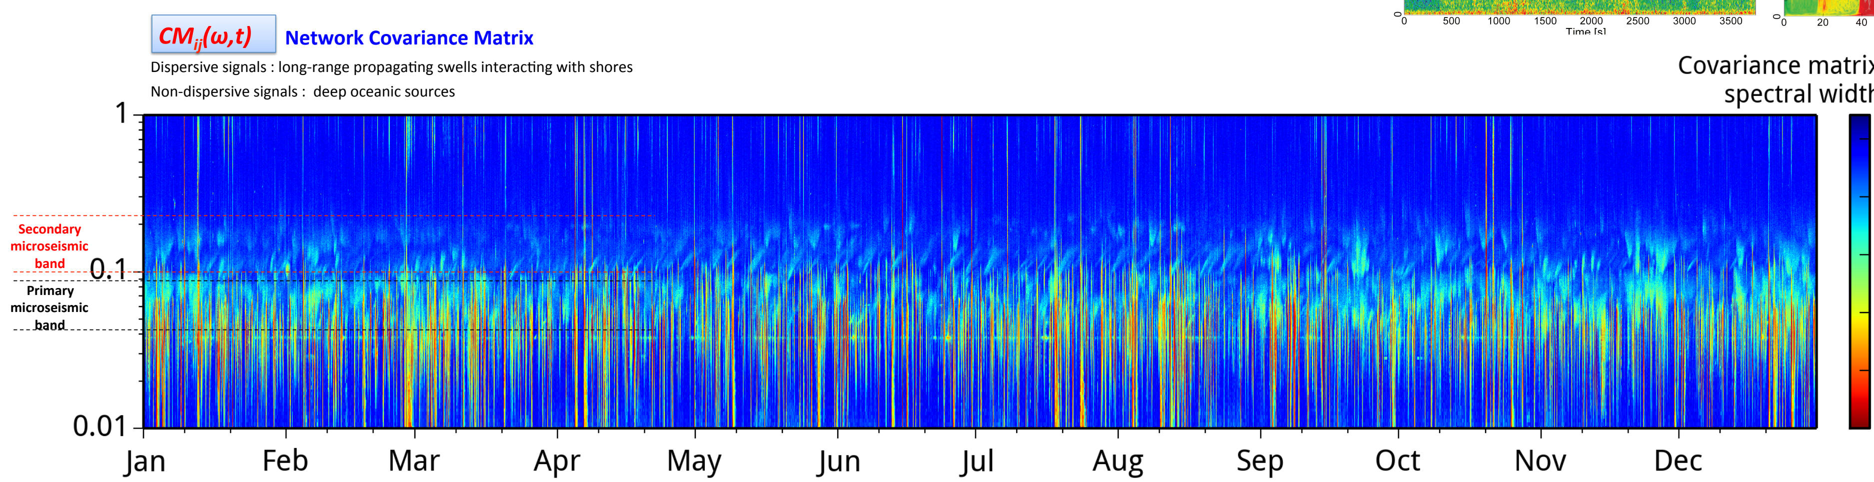

Ocean seismic sources

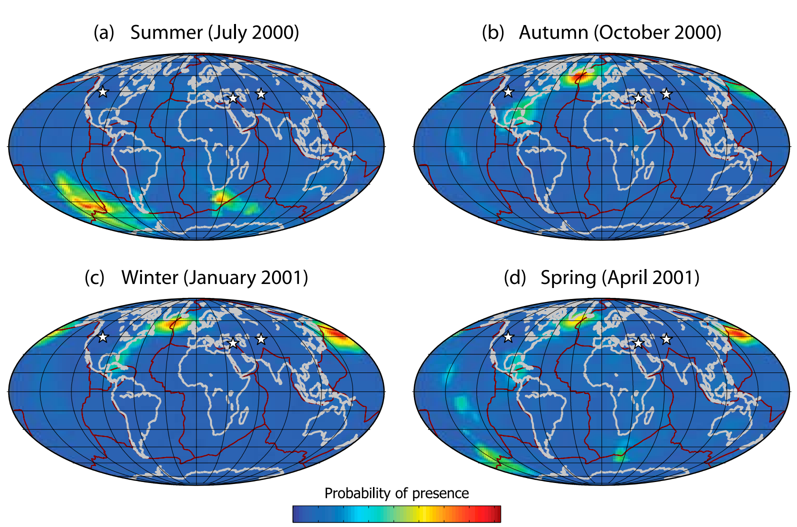

Landes et al

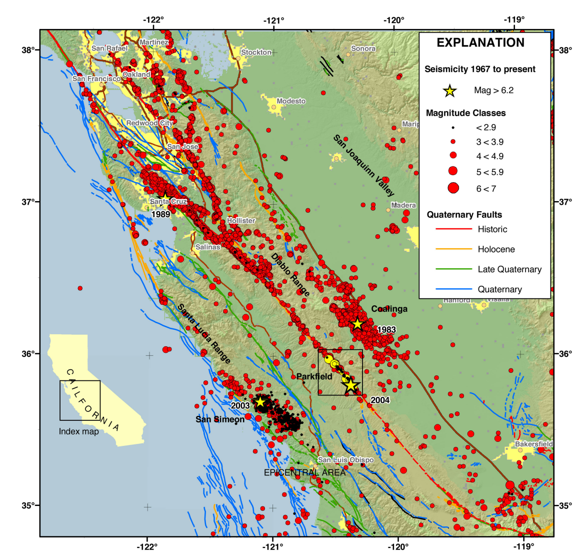

Earthquake-induced property changes

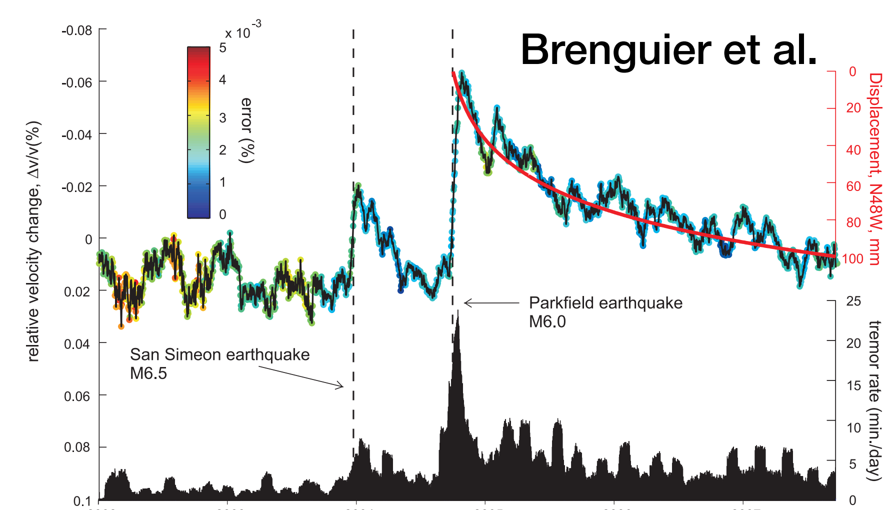

Brenguier et al.

Waveform analysis : data stream workflow

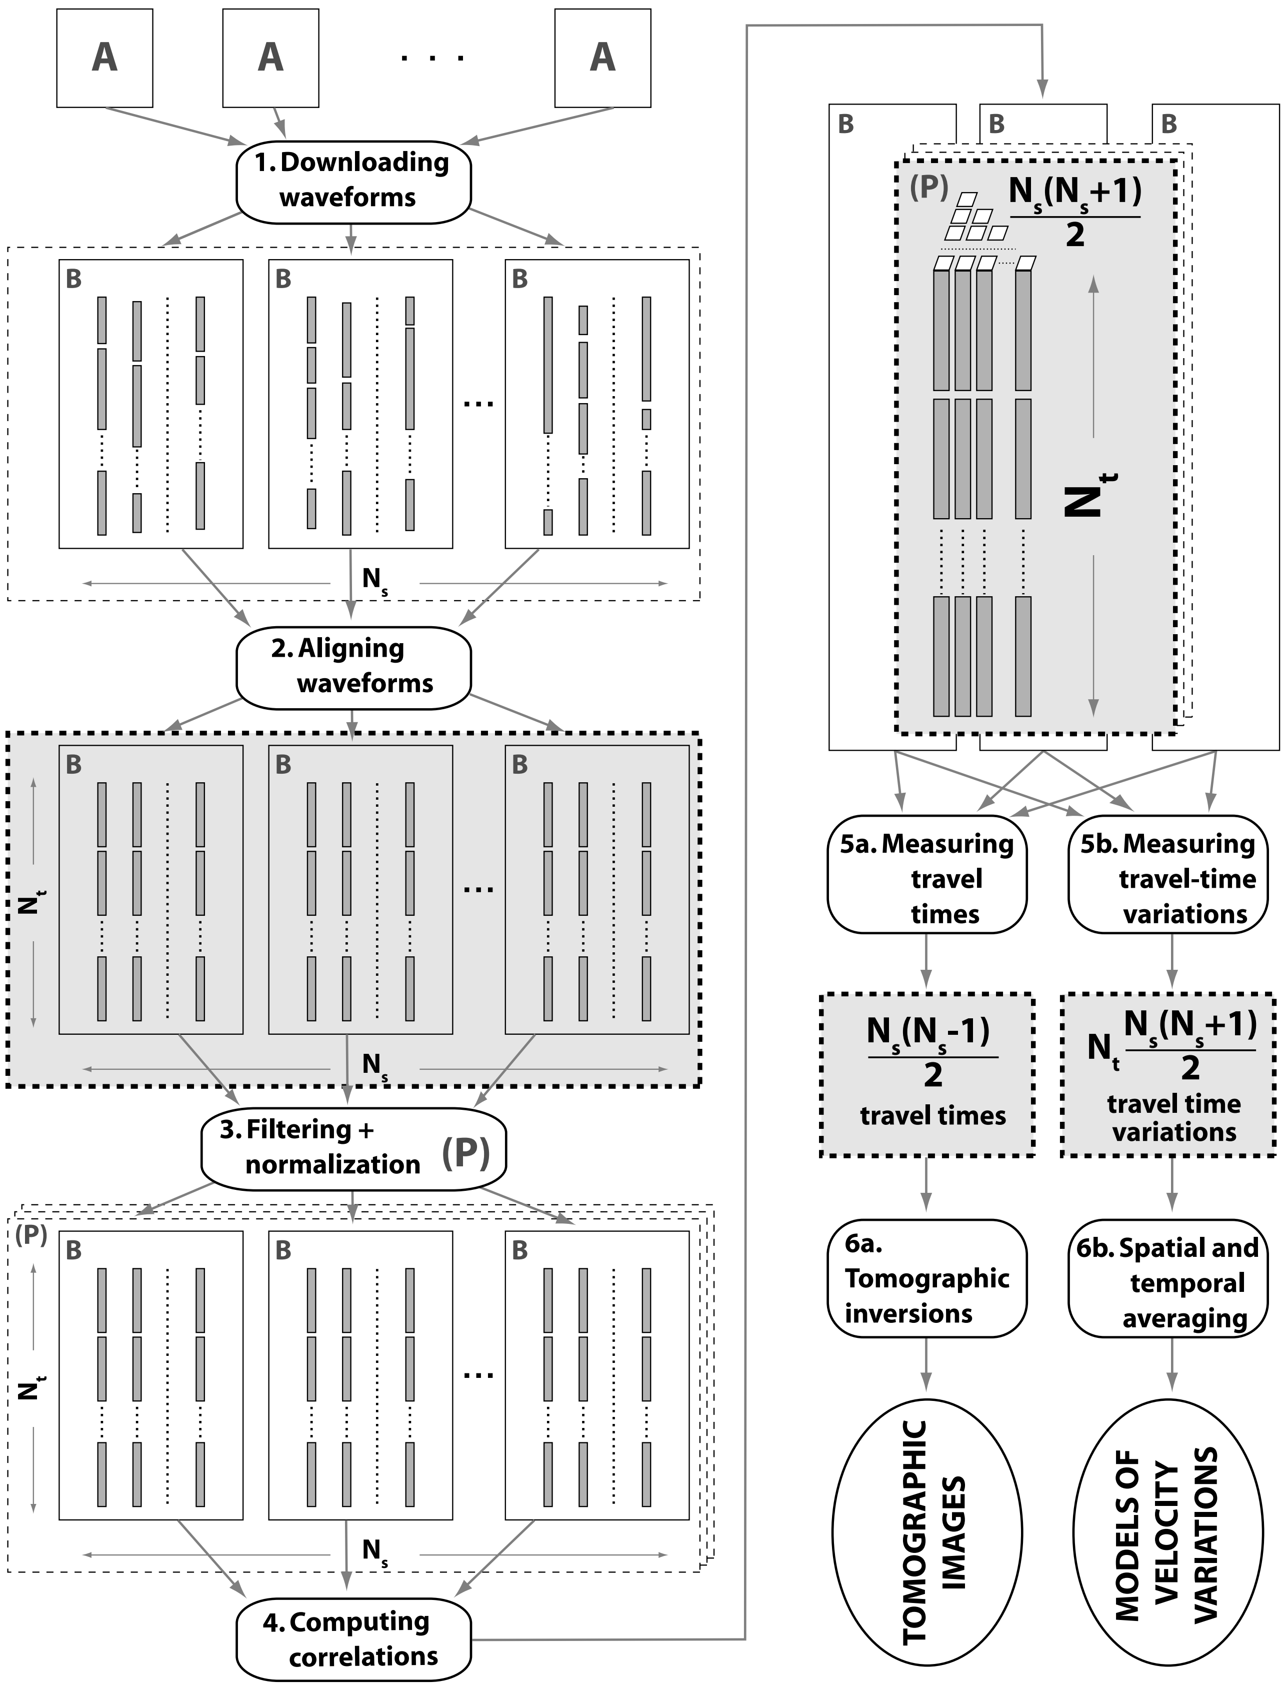

PBs scale …

Data ingestion / quality control

- N-dimensional *time (frequency) series*
- Binary large objects : > ~100 TBs
- *fine granularity (GBs)*
- Partitioning, indexing, replication

Data processing

- Low level data access pattern
- Linear complexity
- fine-grained streaming data workflow
- Provenance and metadata management

Correlators (time-frequency)

- Cross-correlation and higher order statistics
- Quadratic complexity
- Thread-blocks GPUs / MIC
- Secondary data : ~ 6 * N^2 * N_t (* N_f)
- Provenance and metadata management

Imagers (space-time-frequency)

- Convolution / projection
- High-order complexity
- Gridding
- Clustering - classification - machine learning
- Provenance and metadata management

From Shapiro, Vilotte et al.

LOFAR Epoch of Reionisation processing (> 100 MHz)

Omni-directional dipoles

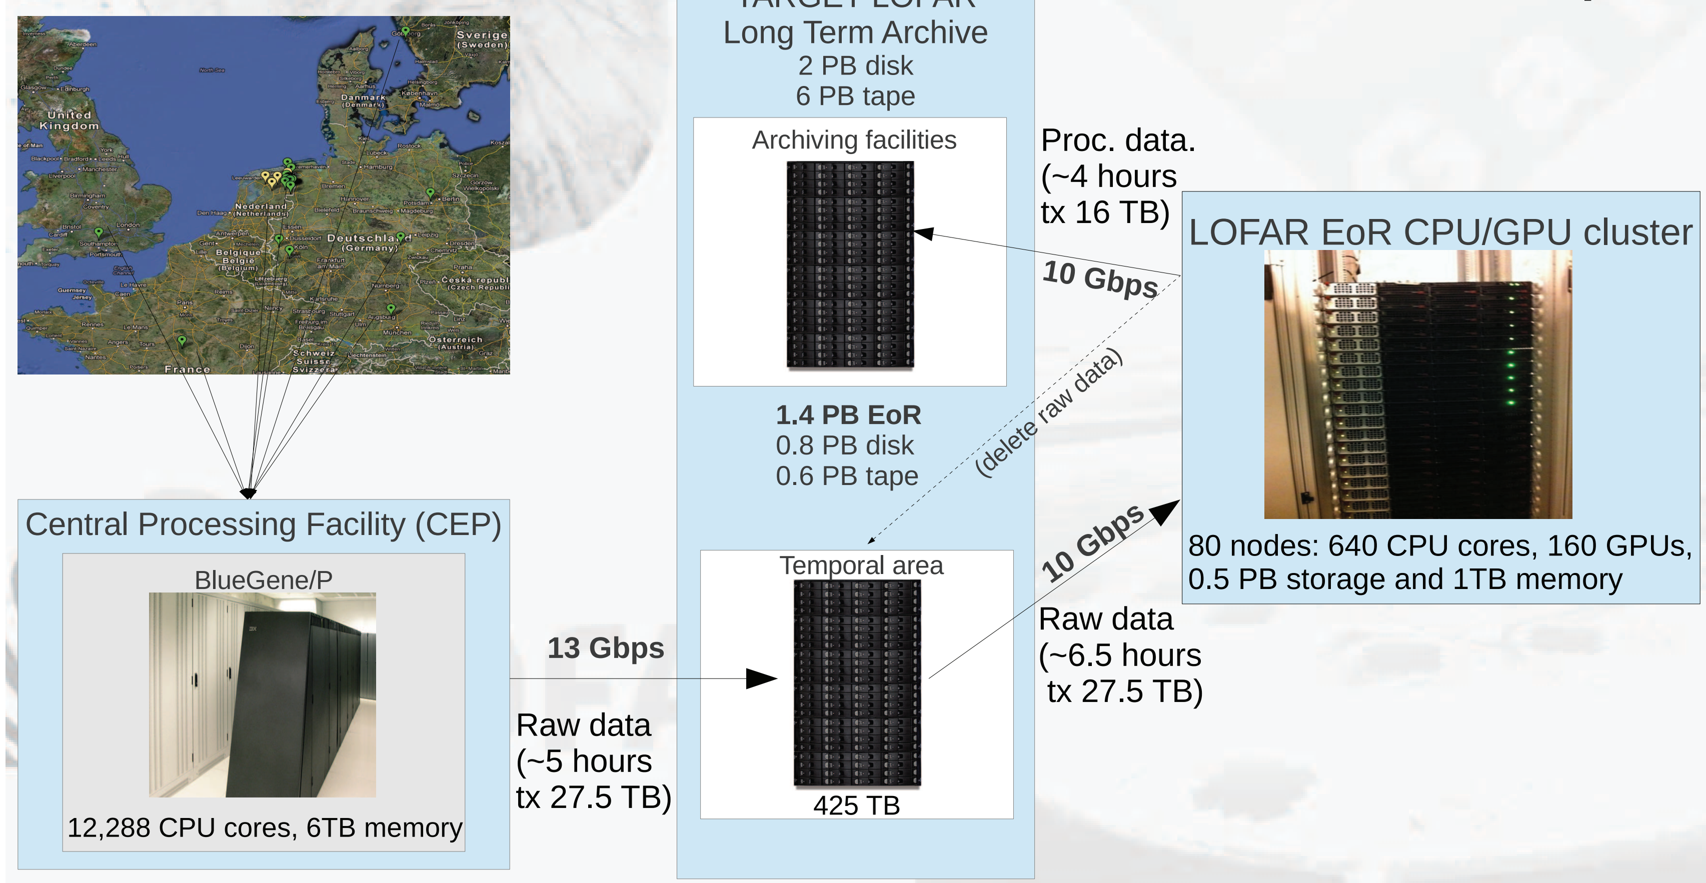

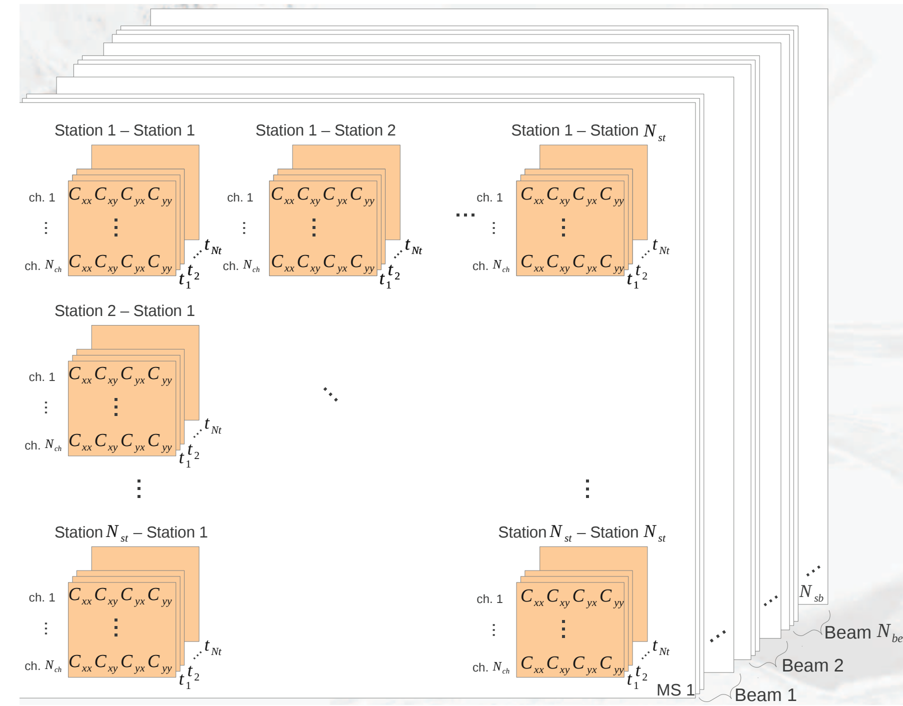

All-sky monitoring: detecting transient events and phase response-time

PB scale

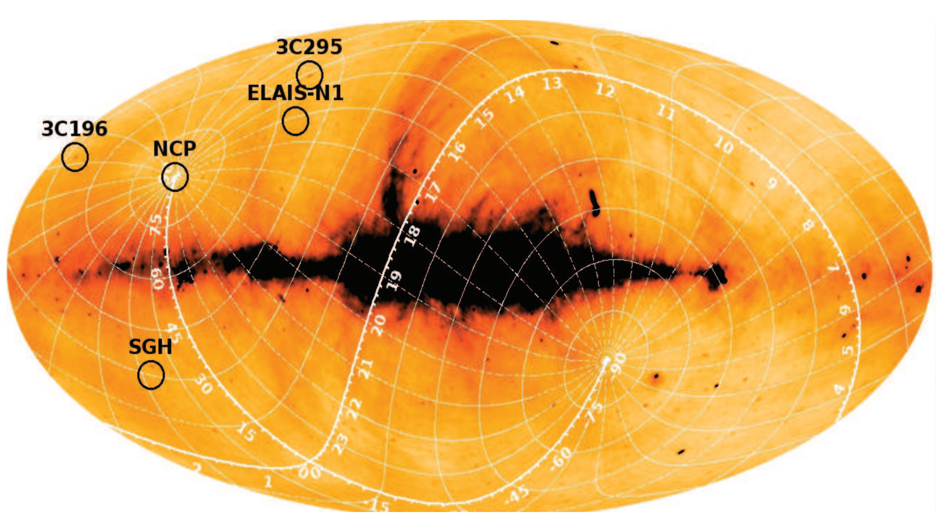

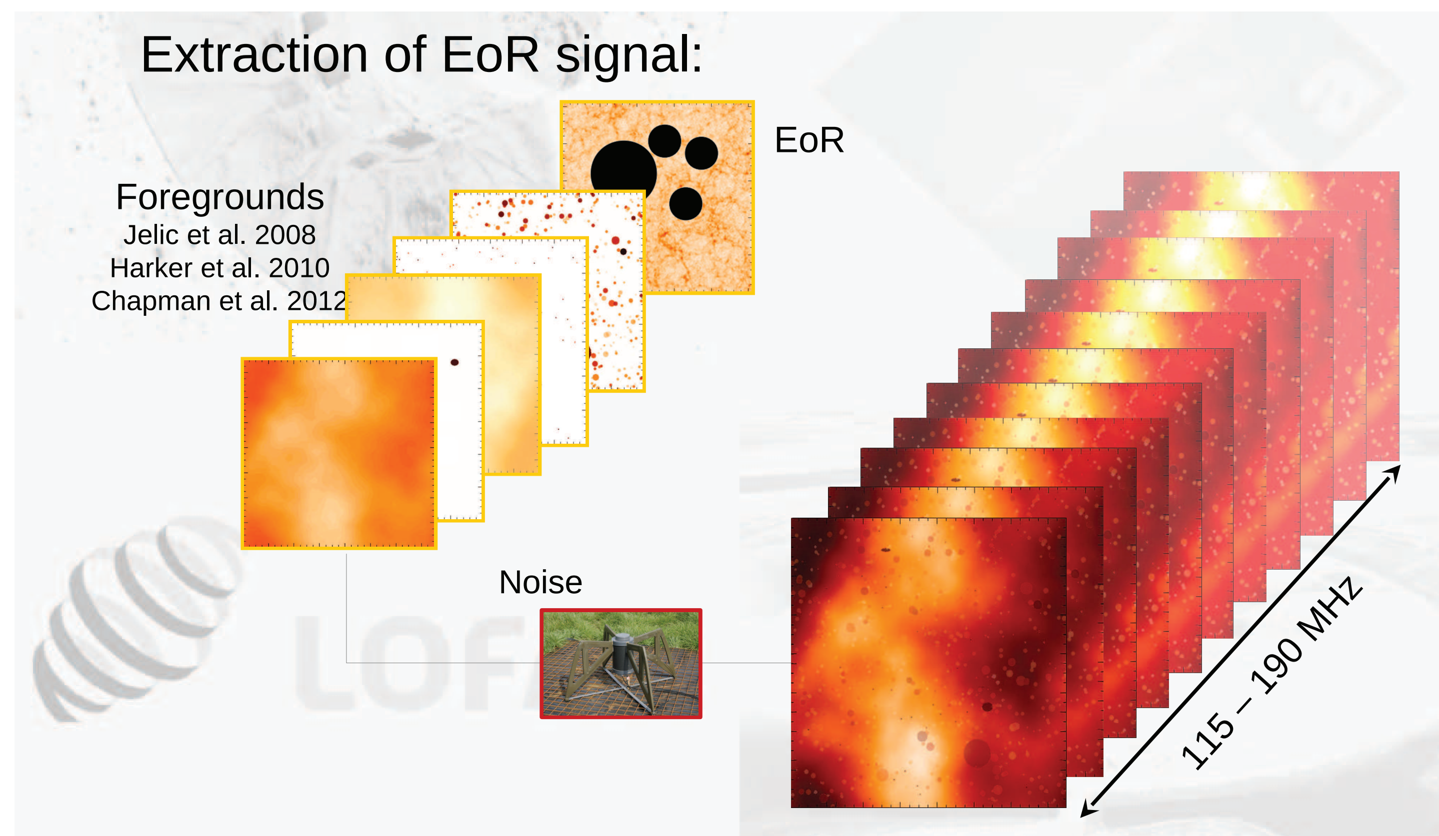

Challenges

What we value and experience

Researchers are part of the archiving process.

They know what is relevant to **understand** their **results**.

Automated system should provide support for a consistent and effective acquisition of **provenance metadata** - ***Selective and extensible Provenance***.

[A. Misra] [I. Foster.]

Data stream processing engines

Data Intensive computation, present expensive requirements for provenance collection, either in terms of size or I/O [W. D. Pauw]

Data-intensive framework

Enable active researchers to invent, refine scalable, statistical data-intensive methods

Support diversity of methods and implementation in a single data-intensive framework with data-handling services

Researchers remain in full control in their familiar community tools and libraries

Collaborative developments: from theoretical research to proof of concept to sustained use

Python library used to describe **abstract workflows** for distributed data-intensive applications.

Support for composition: Processing Elements defined with their own internal workflows.

Abstract streaming data flows: can be map and automatically executed in a variety of **parallel environments.**

Fine-grained provenance system: analyse and understand data relationship with triggered actions

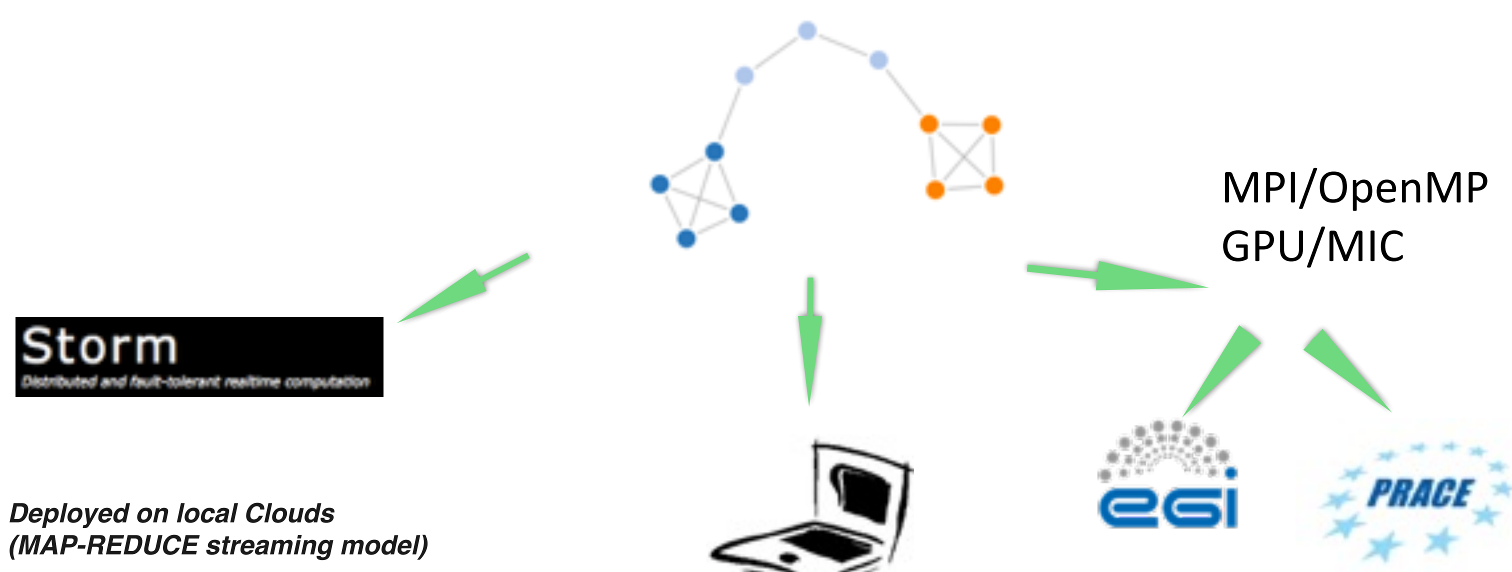

From EU VERCE project

Data-driven computing applications
Data inversion and assimilation

Data-driven applications: inversion and assimilation

Exploration and marine geophysics

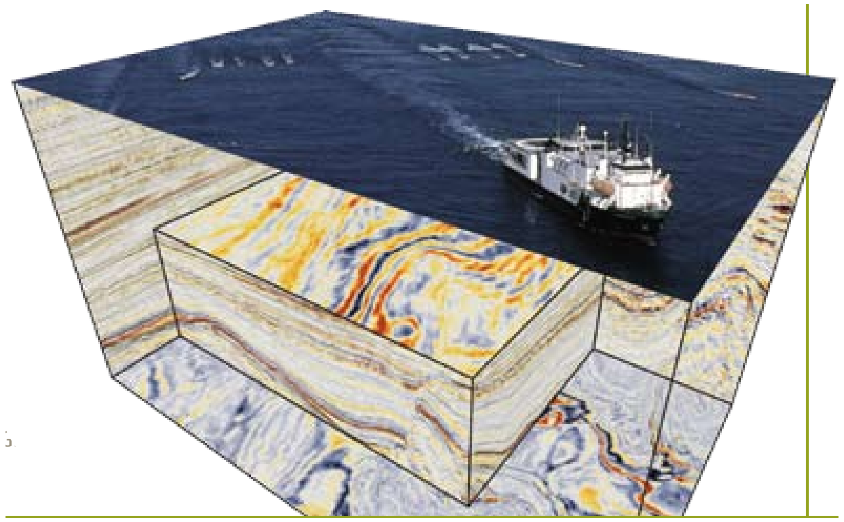

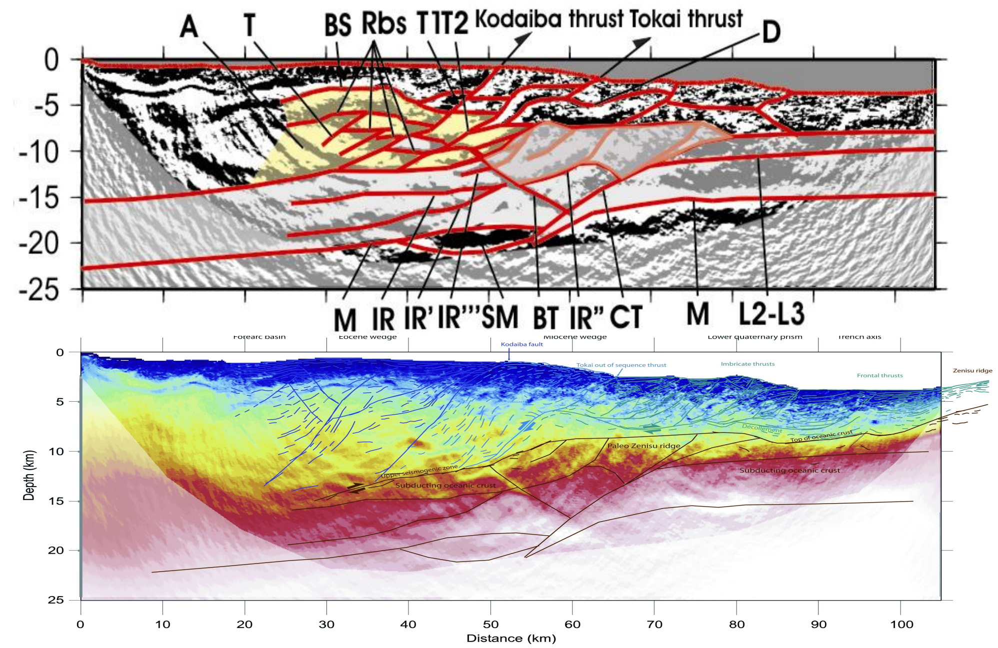

Global scale tomography

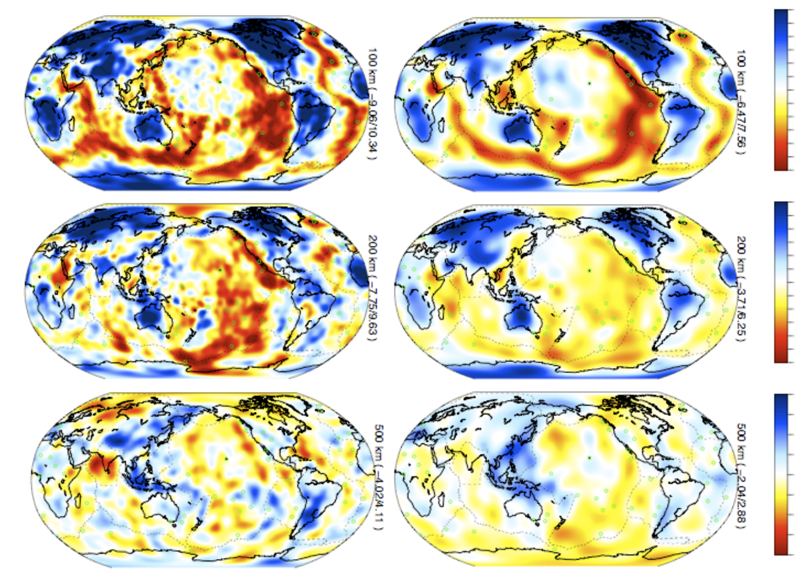

Geomagnetic secular variation

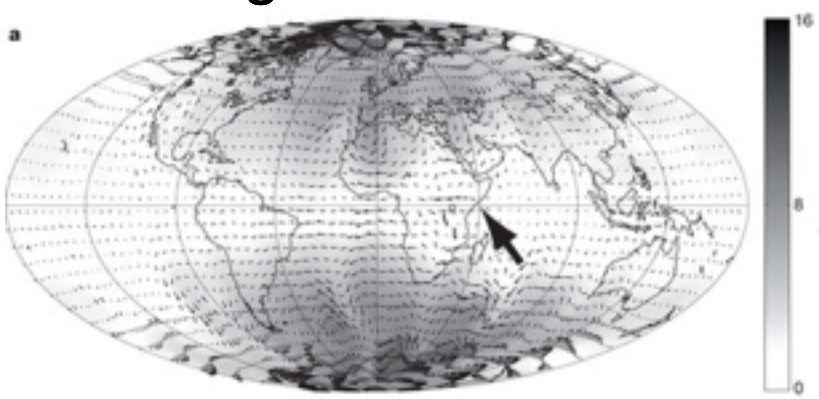

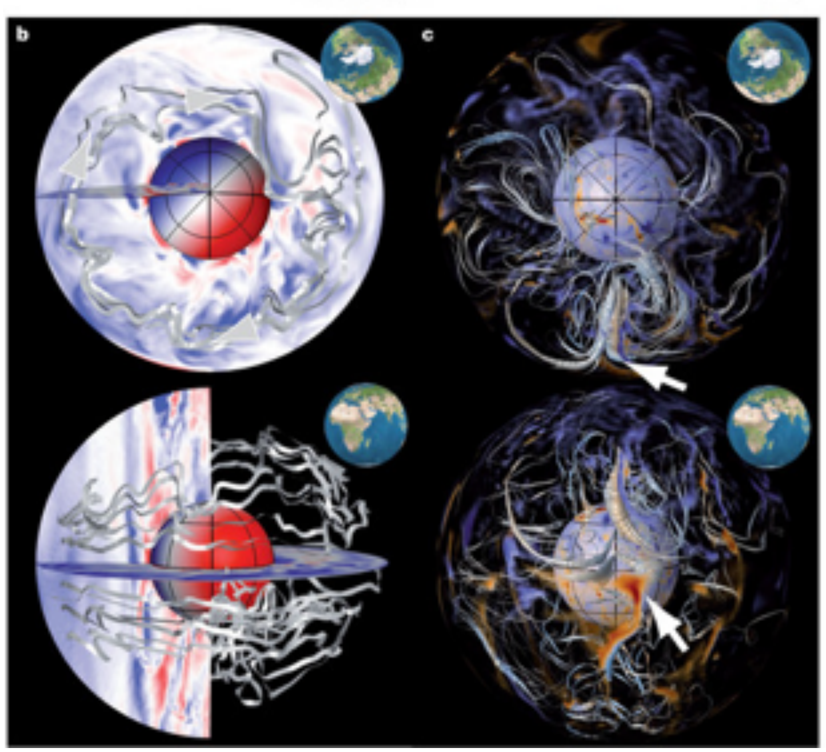

Earthquake source imaging

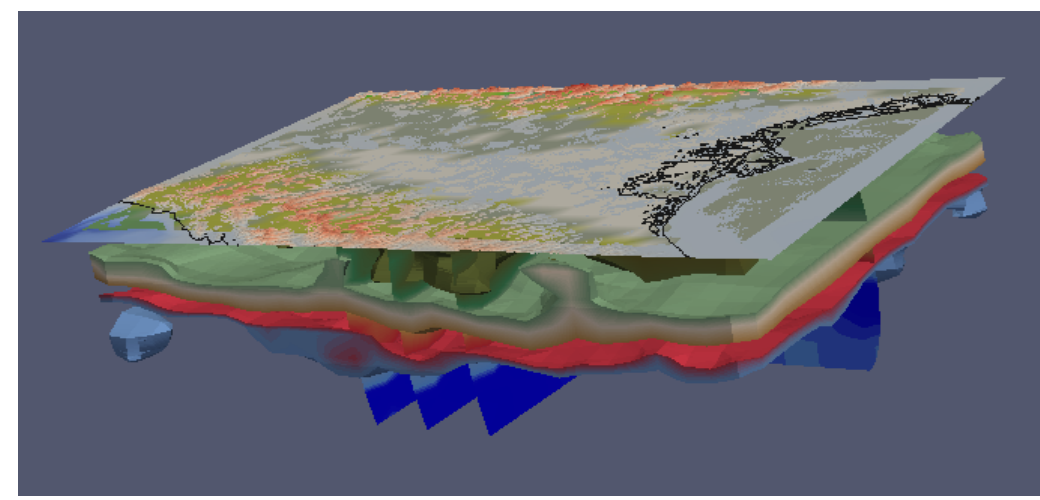

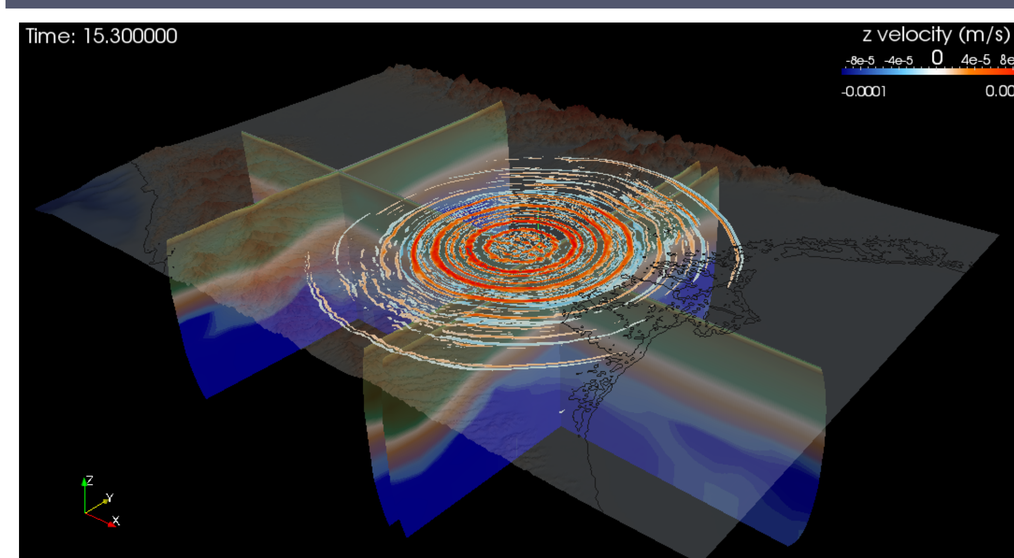

Seismology

- Full Waveform Inversion
- Extended Earthquake source

Geomagnetism

- Inversion of secular variation
- Variational data assimilation

Gravimetry

- Inversion of gravity field
- Geoid and Earth shape

Toward Bayesian-inference reconstruction

Orchestrated workflow: data-intensive & HPC

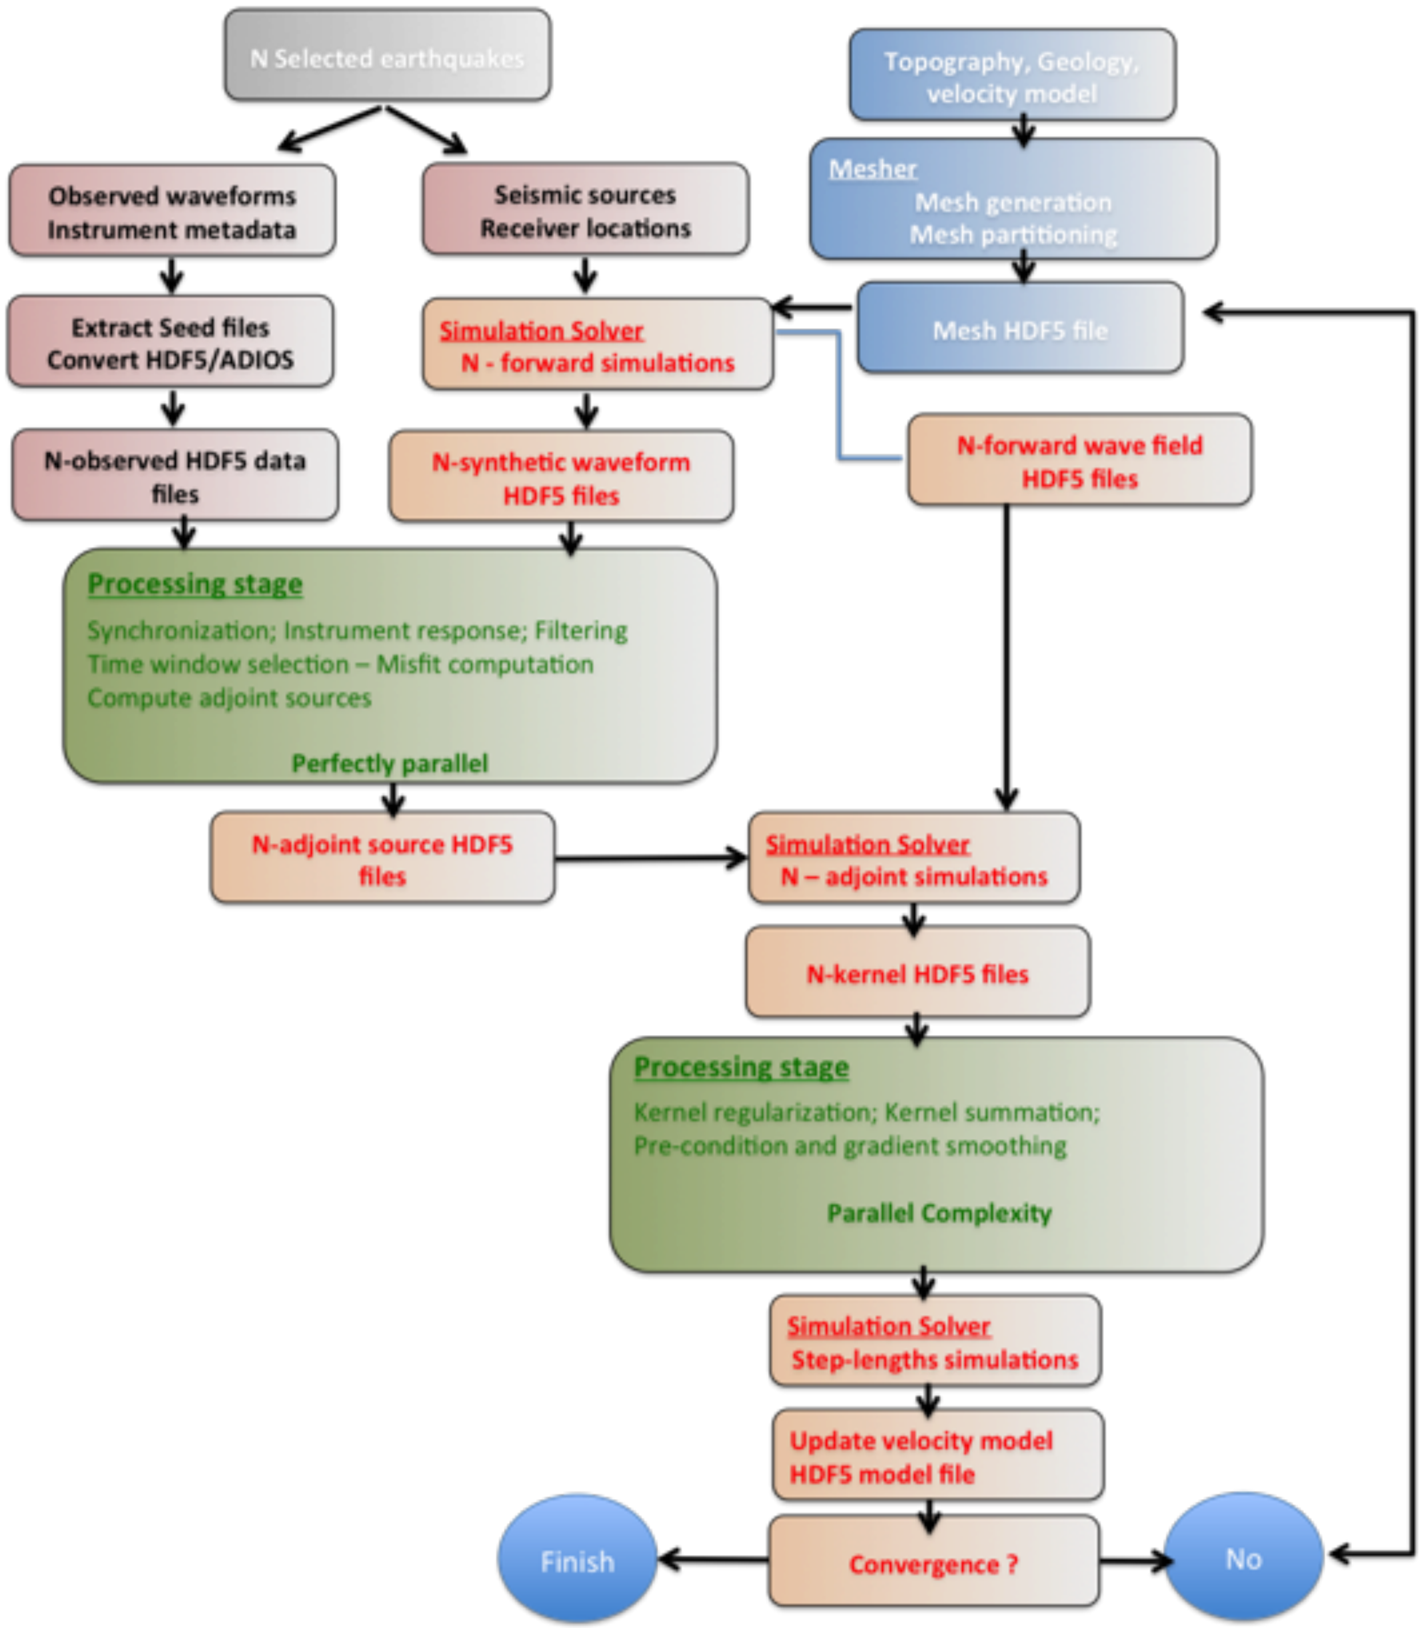

Full Waveform Inversion (FWI)

- non-linear Bayesian inversion
- adjoint-based inversion

High-performance parallel codes

- forward and adjoint wave simulations
- billion of coresOrchestrated workflow
- data-intensive analysis and HPC
- CPU and Data-intensive architecture

Big N

- synthetics and observed wave forms
- Earth model and wave propagation
- I/O and CPU balance (~10s Gb/s, 100Tb per iteration)
- higher-oder abstract file format (HDF5)
- indexing and Data Bases

Adapted from Tromp, Komatitsch et al.

FWI compute and data analysis

Convergence of data with computation

Federating autonomous diverse resources

Handling independent data sources

Fluent path from development production

Hiding complexity

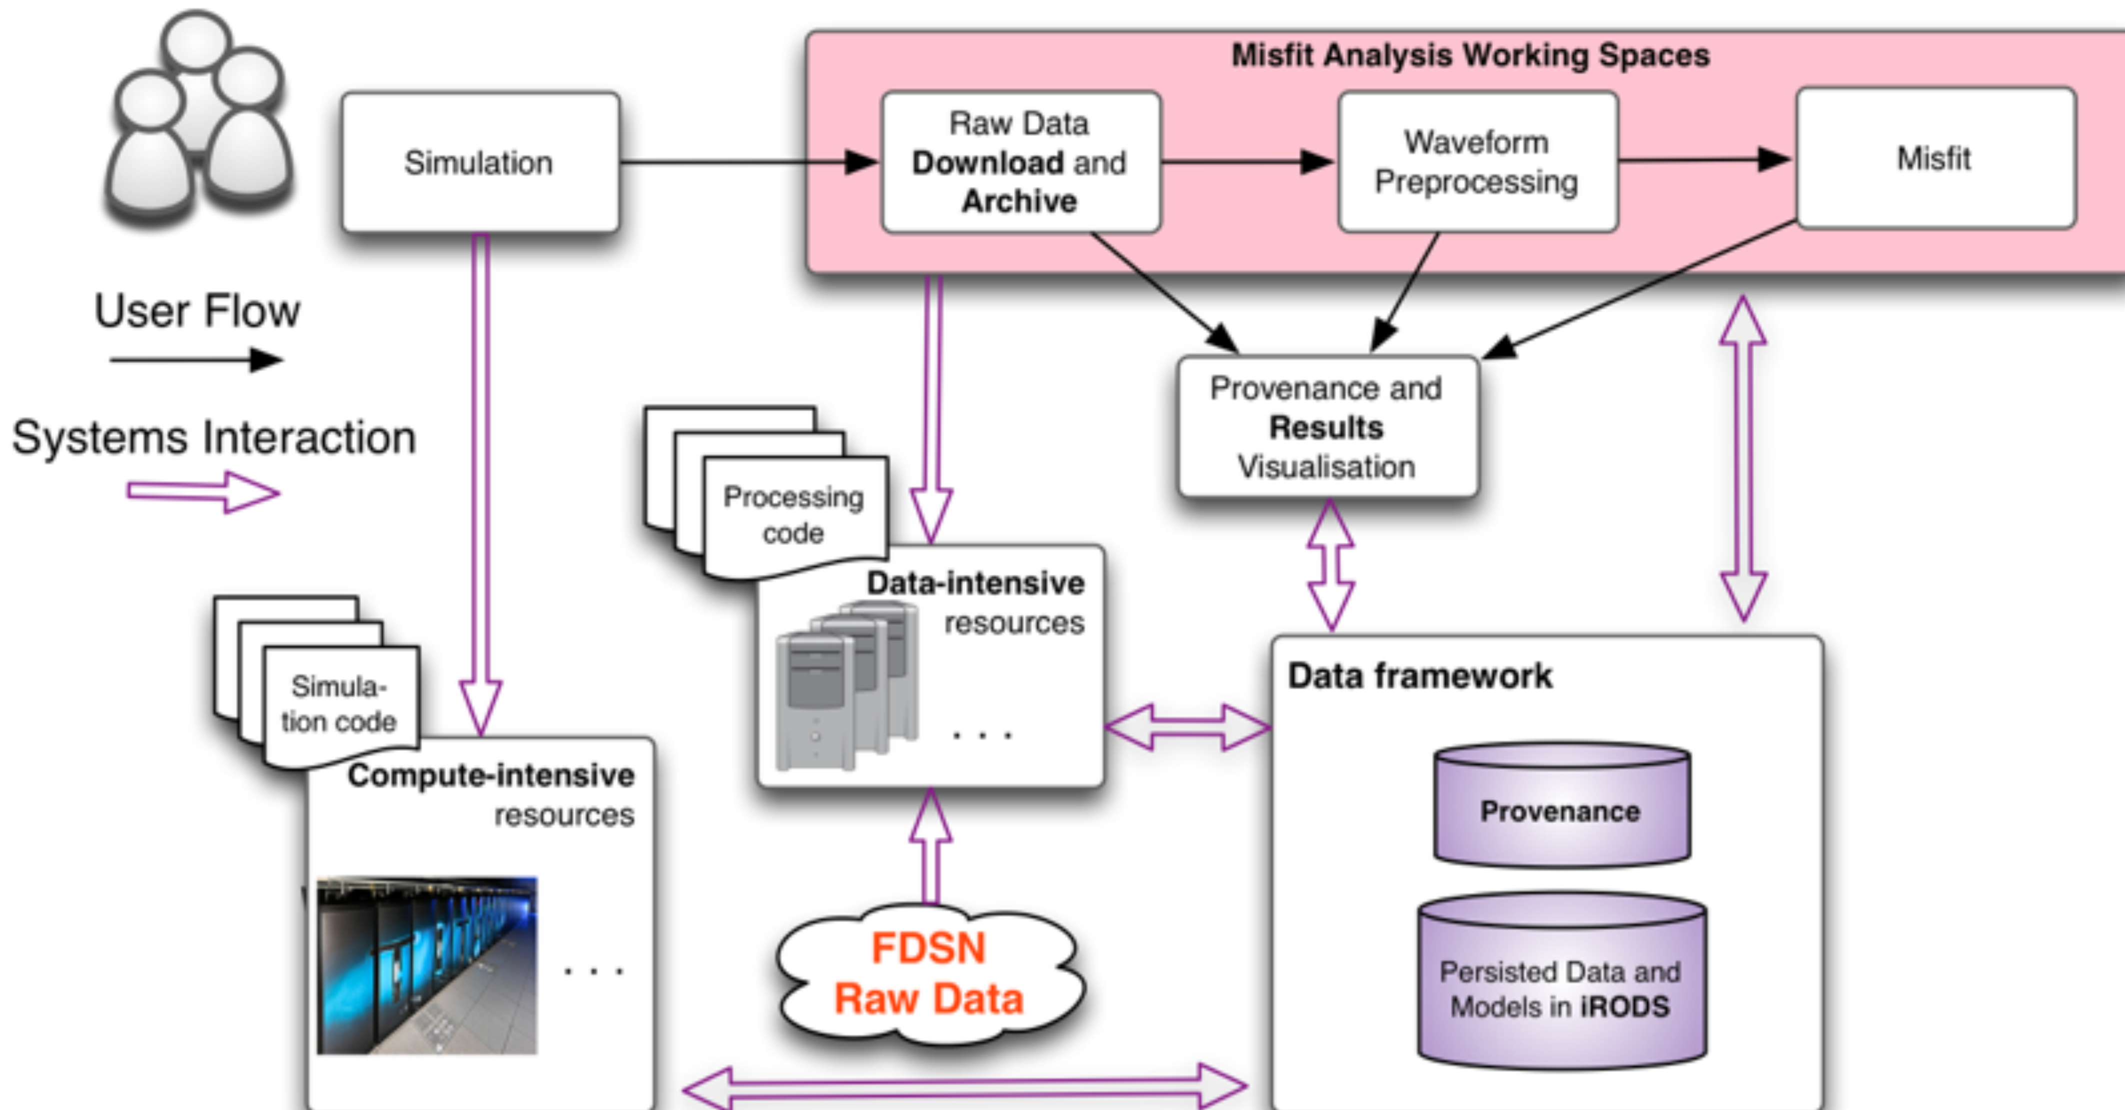

- ✳ **Federation of independent autonomous organisations**
  - data and computing infrastructures providers.
- ✳ **Services/access policies**: data-transfer, job control, task-oriented workflows
- ✳ **Transient storage** for users' work in progress and intermediate data.
- ✳ **Shared persistent and caching storage:** optimise costs of data movement, assembly, processing, distilling and simulations over multiple investigations

Challenges

What we want: flexibility and reactive systems and users

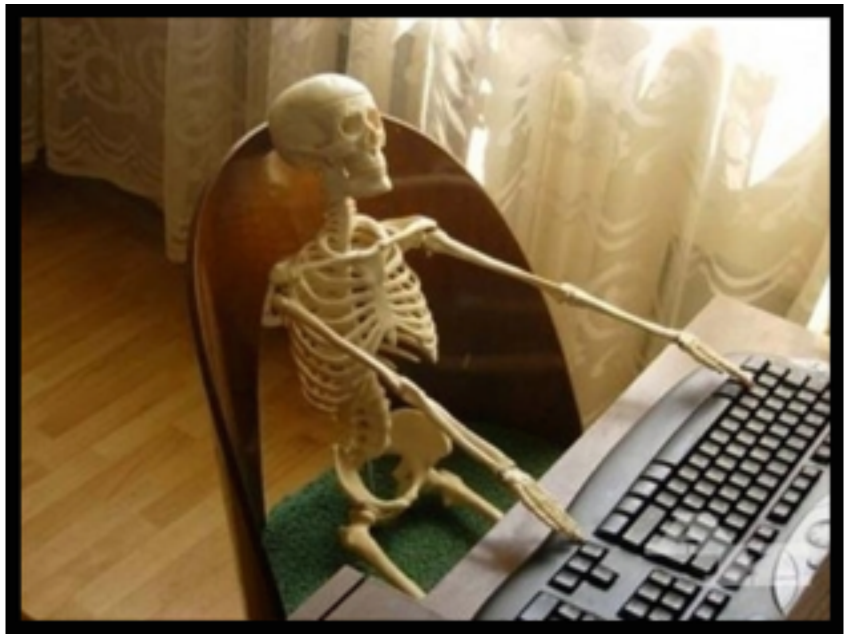

- **Provides** Run-Time **feedback** on the process with **tuneable metadata** and **controlled data movements**
- **Avoids** useless waits **for long and unfruitful runs**
- **Fosters** Dynamic Steering, Diagnostics**, saving computing cycles, storage ($$) and energy!**

Turning large simulations into numerical laboratories

Turning large simulations into numerical laboratories

Mantle convection geodynamics

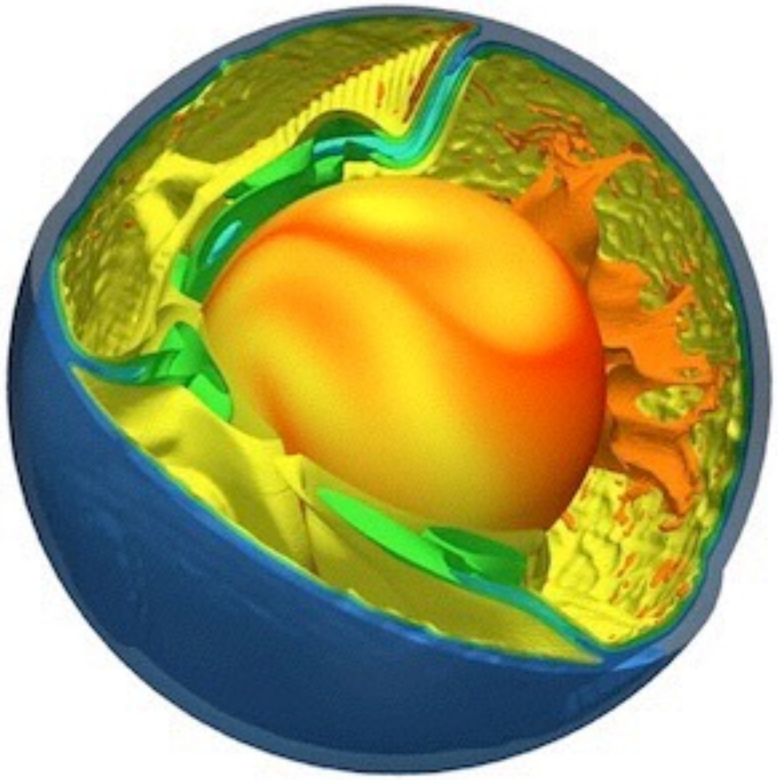

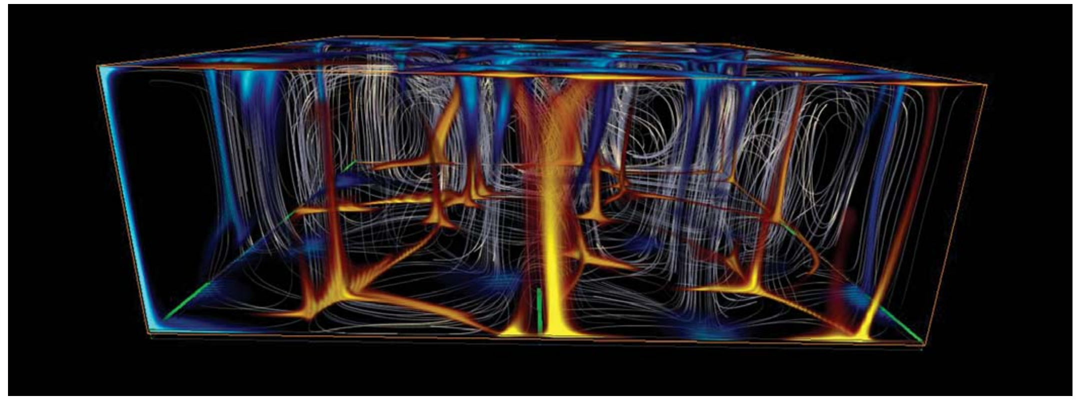

Strong motion prediction

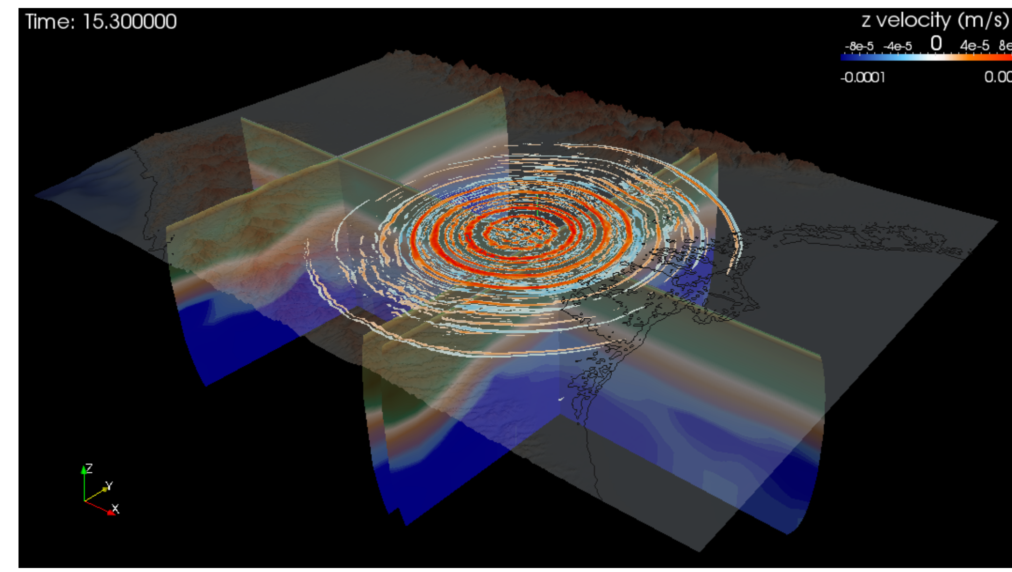

Geodynamo and Earth's core dynamics

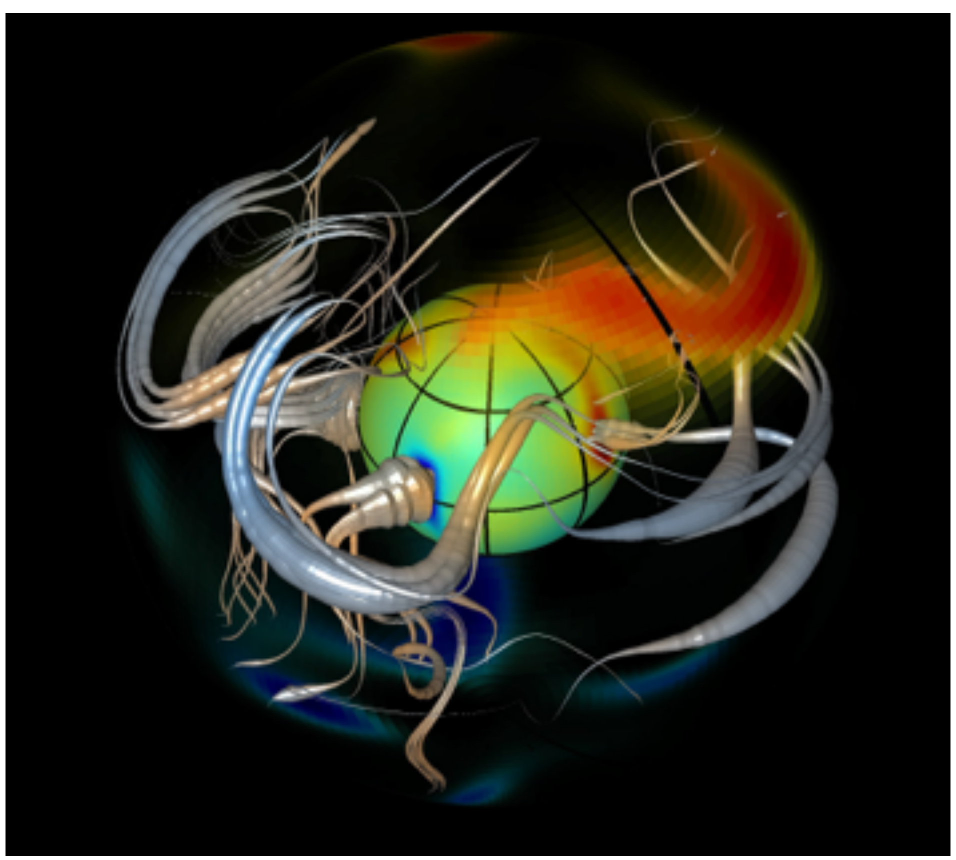

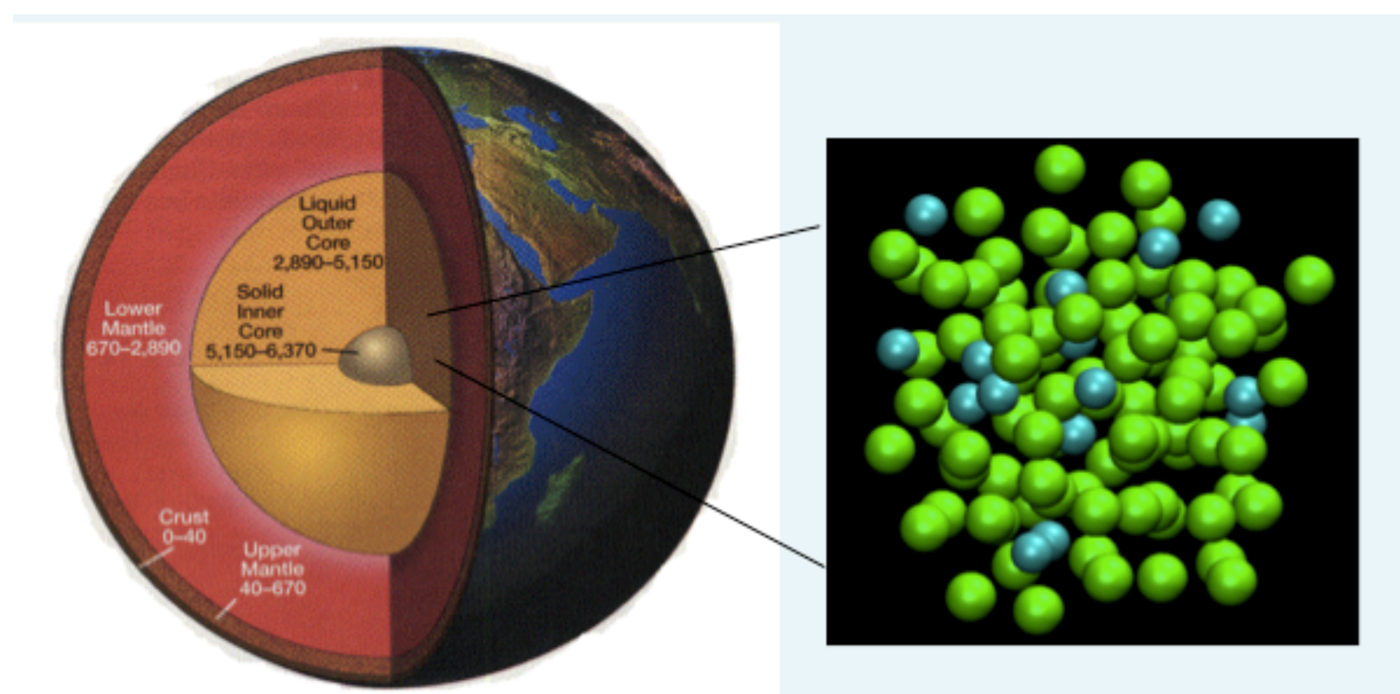

Molecular dynamics : High Pressure and High Temperature Physics

The AMA-DEUS application: N-Body simulation

A TGCC-CURIE grand challenge

- 550 billion particles
- 2.5 trillion computing points
- 50 million CPU hours (> 5700 years)
- 76 032 cores & 300 Tb memory
- > 50 Gb/s data throughput (PFS)
- 1 500 Pbs reduced on fly to 1 500 Tbs

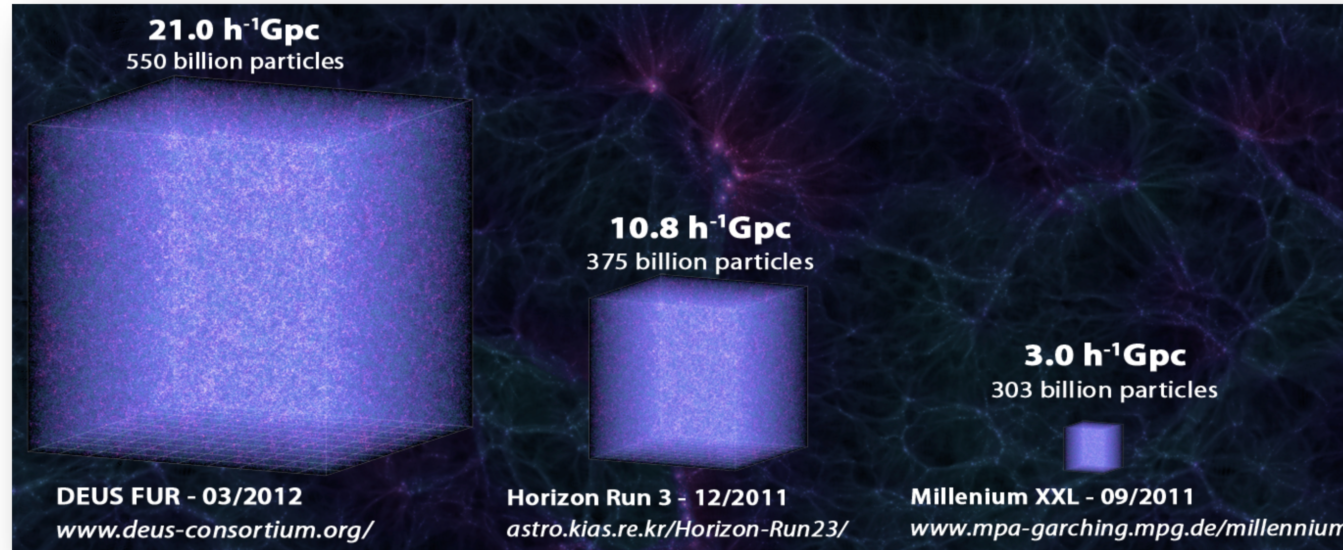

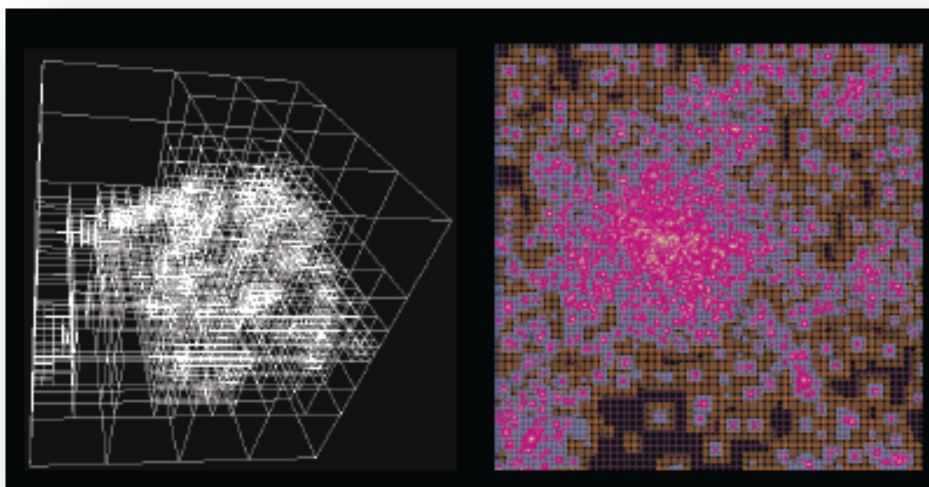

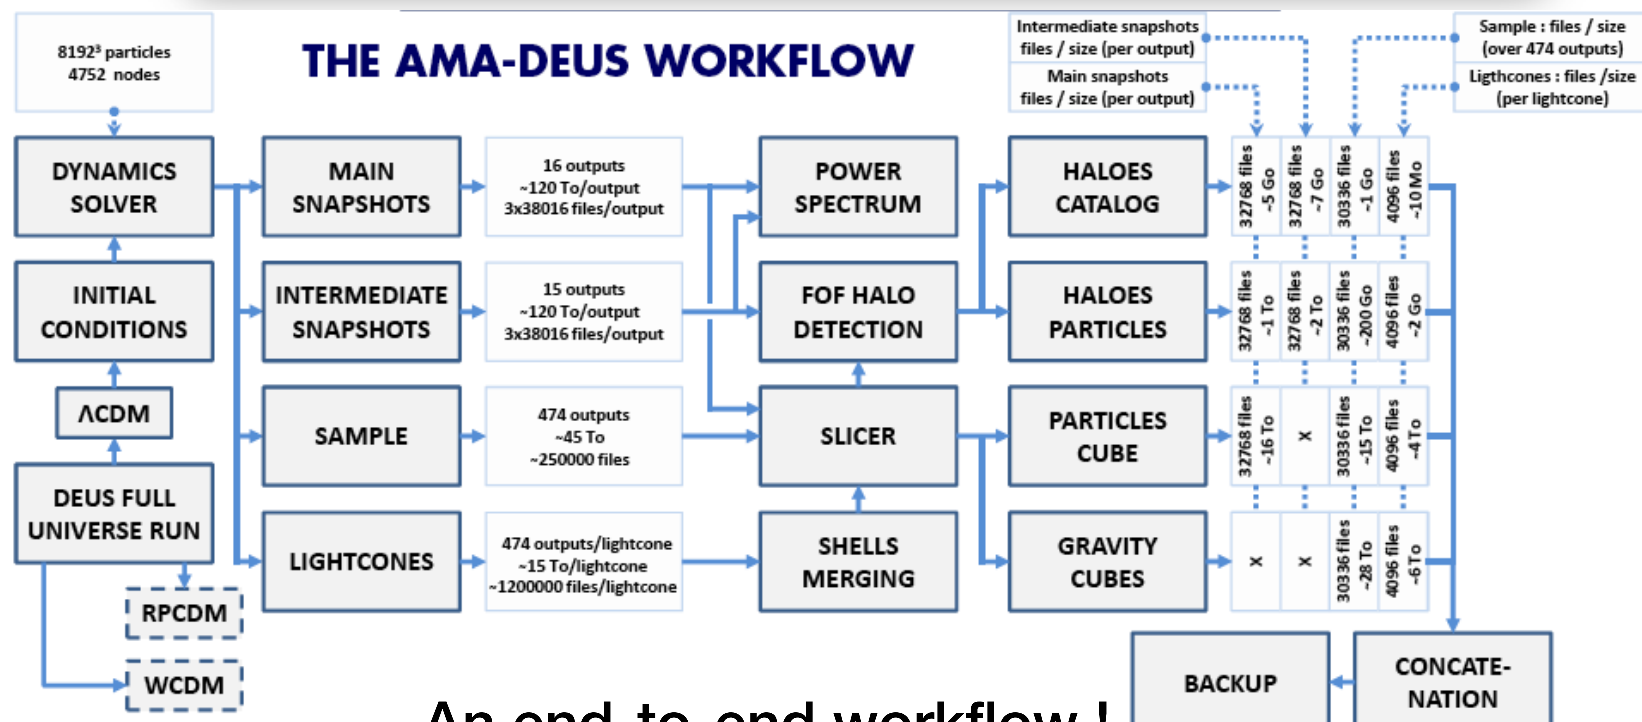

An end-to-end workflow !

Challenges

- dynamic load balancing
- smart parallel I/O optimisation
- reduction of raw data (time) -> direct post-processing
- physical objects -> on-the-fly processing workflow

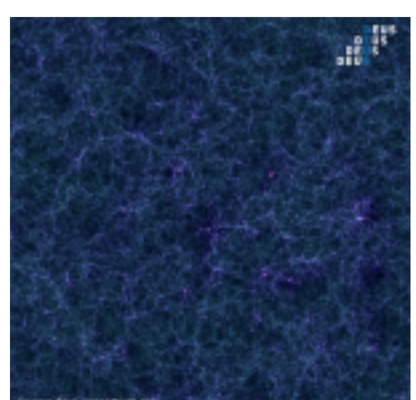

Snapshots
~16 x 16 TB

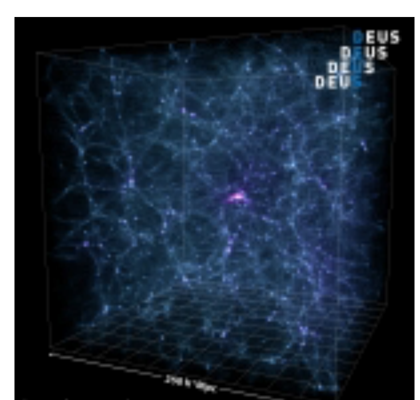

Samples
~40 TB

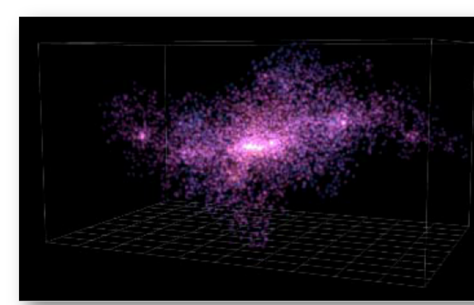

Halos/catalogs
~50 TB

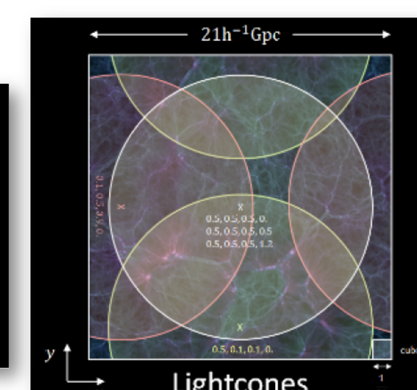

Lightcones
~ 5x10 TB

IAP, courtesy of J.-M. ALIMI

Numerical laboratory: Shared Data Analysis

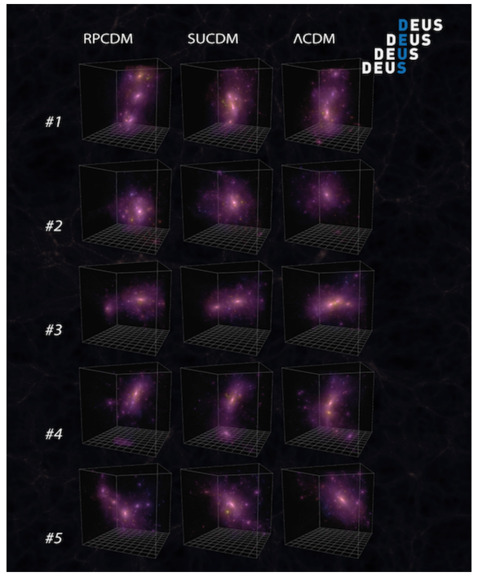

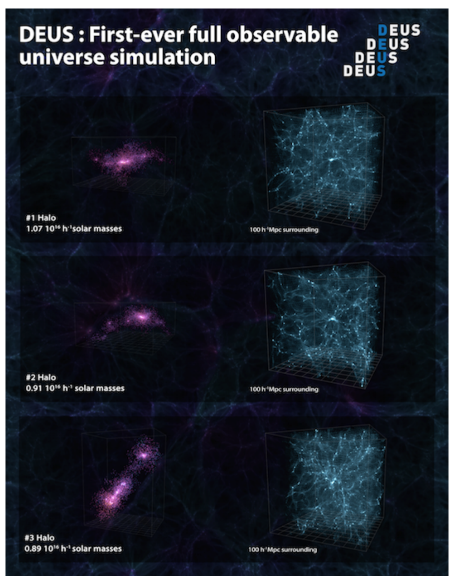

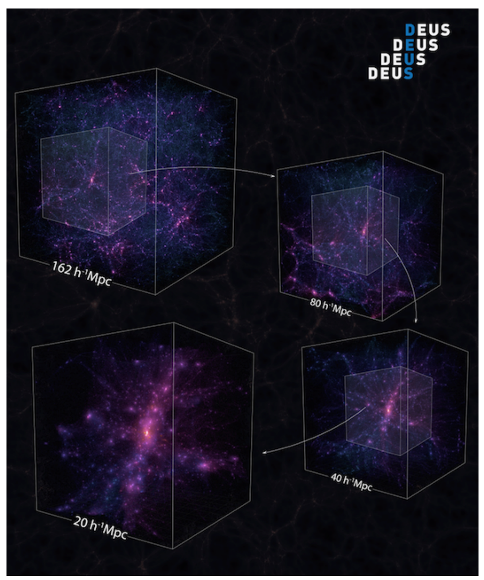

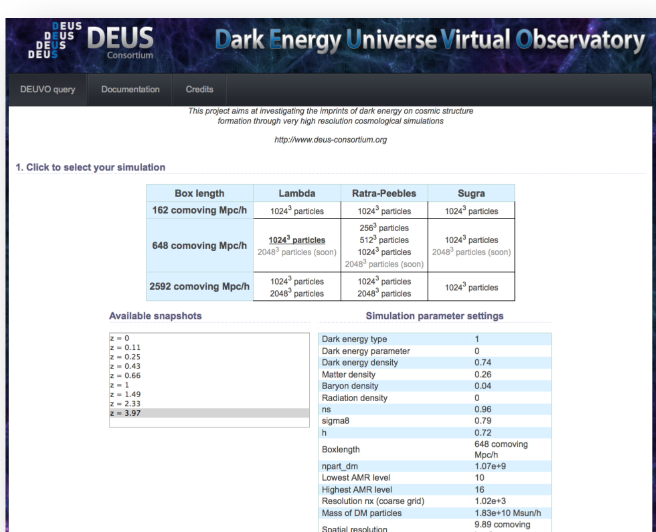
DEUS Consortium — Dark Energy Universe Virtual Observatory

DEUVO query | Documentation | Credits

This project aims at investigating the imprints of dark energy on cosmic structure formation through very high resolution cosmological simulations

http://www.deus-consortium.org

1. Click to select your simulation

Box length	Lambda	Ratra-Peebles	Sugra
162 comoving Mpc/h	1024^3 particles	1024^3 particles	1024^3 particles
648 comoving Mpc/h	1024^3 particles 2048^3 particles (soon)	256^3 particles 512^3 particles 1024^3 particles 2048^3 particles (soon)	1024^3 particles 2048^3 particles (soon)
2592 comoving Mpc/h	1024^3 particles 2048^3 particles	1024^3 particles 2048^3 particles	1024^3 particles

Available snapshots

z = 0
z = 0.11
z = 0.25
z = 0.43
z = 0.66
z = 1
z = 1.49
z = 2.33
z = 3.97

Simulation parameter settings

Dark energy type	1
Dark energy parameter	0
Dark energy density	0.74
Matter density	0.26
Baryon density	0.04
Radiation density	0
ns	0.96
sigma8	0.79
h	0.72
Boxlength	648 comoving Mpc/h
npart_dm	1.07e+9
Lowest AMR level	10
Highest AMR level	16
Resolution nx (coarse grid)	1.02e+3
Mass of DM particles	1.83e+10 Msun/h
Spatial resolution	9.89 comoving kpc/h

Consortium DEUS
- scientific teams coordination
- DEUVO DB: physical objects and some raw data

In-situ data reduction

On-the-Fly
- MPI-based power spectrum
- MPI-based parallel Halos finder
- Halos properties

Shared data analysis

Services on top of the data
- Higher-order statistics for matter field and Halos
- Topological analysis
- Dynamical analysis
- Visualisation

…

Data life-cycle: persistent storage, provenance, publication

Climate and weather modelling

A continuum of time and space scales

From days to months, years, decades, and millennia

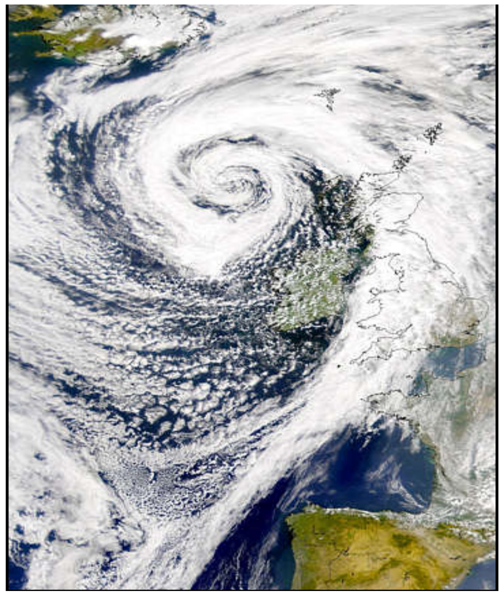

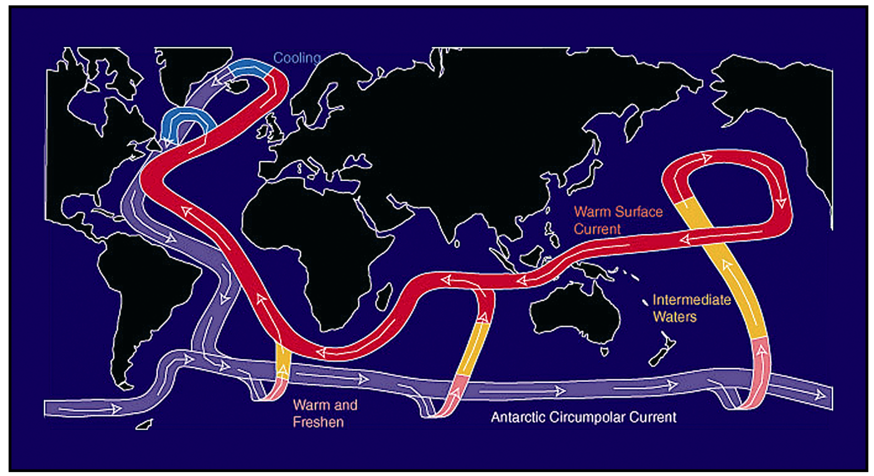

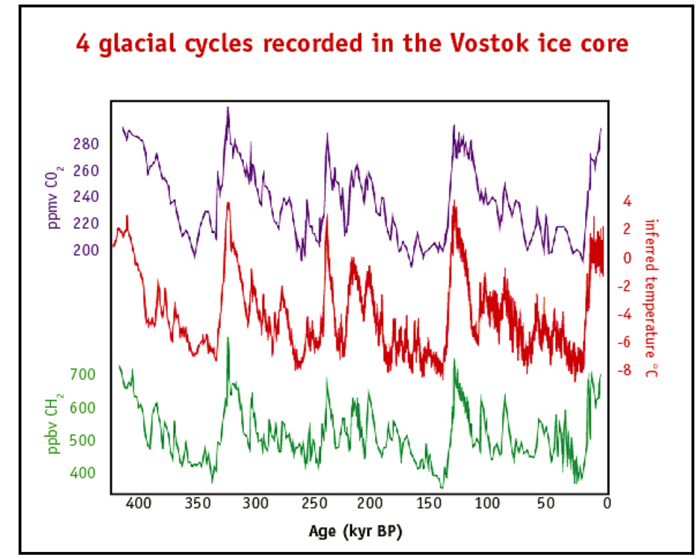

From local to regional, continental and global

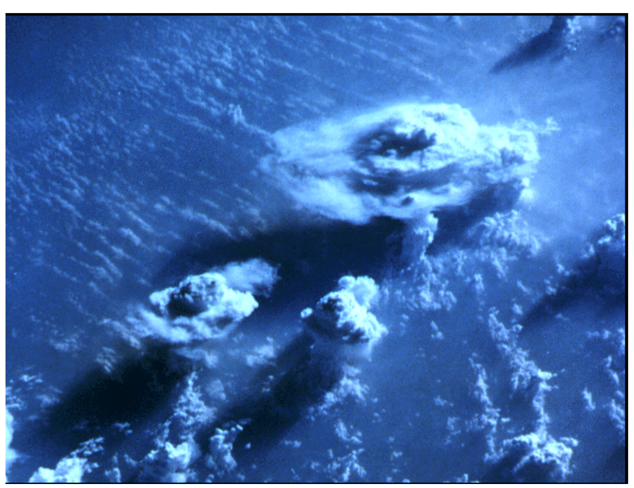

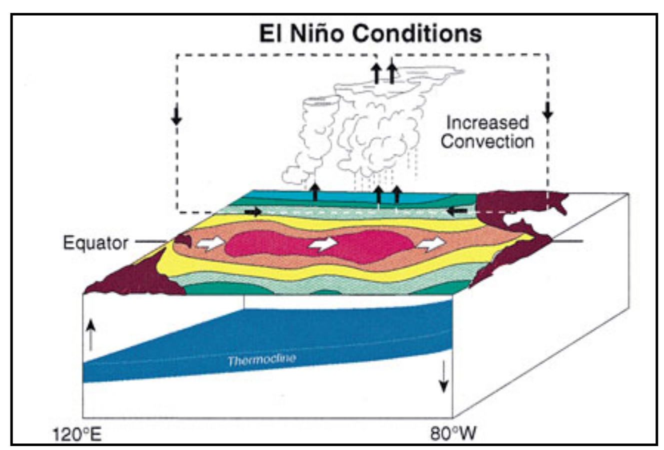

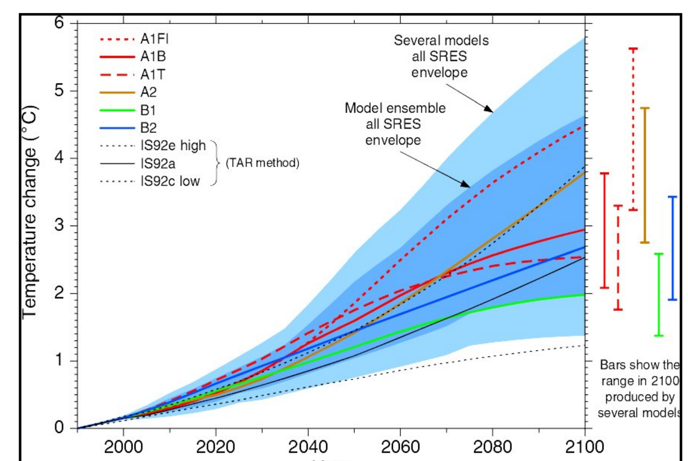

Detection, attribution and prediction of extreme events and modes of climate variability

Climate simulations and observations

The volume of worldwide climate data is expanding creating **challenges for both physical archiving and sharing**, as well as for **ease of access and finding** what is needed particularly if you are not a climate scientist.
Overpack *et al.*, Science (2011)

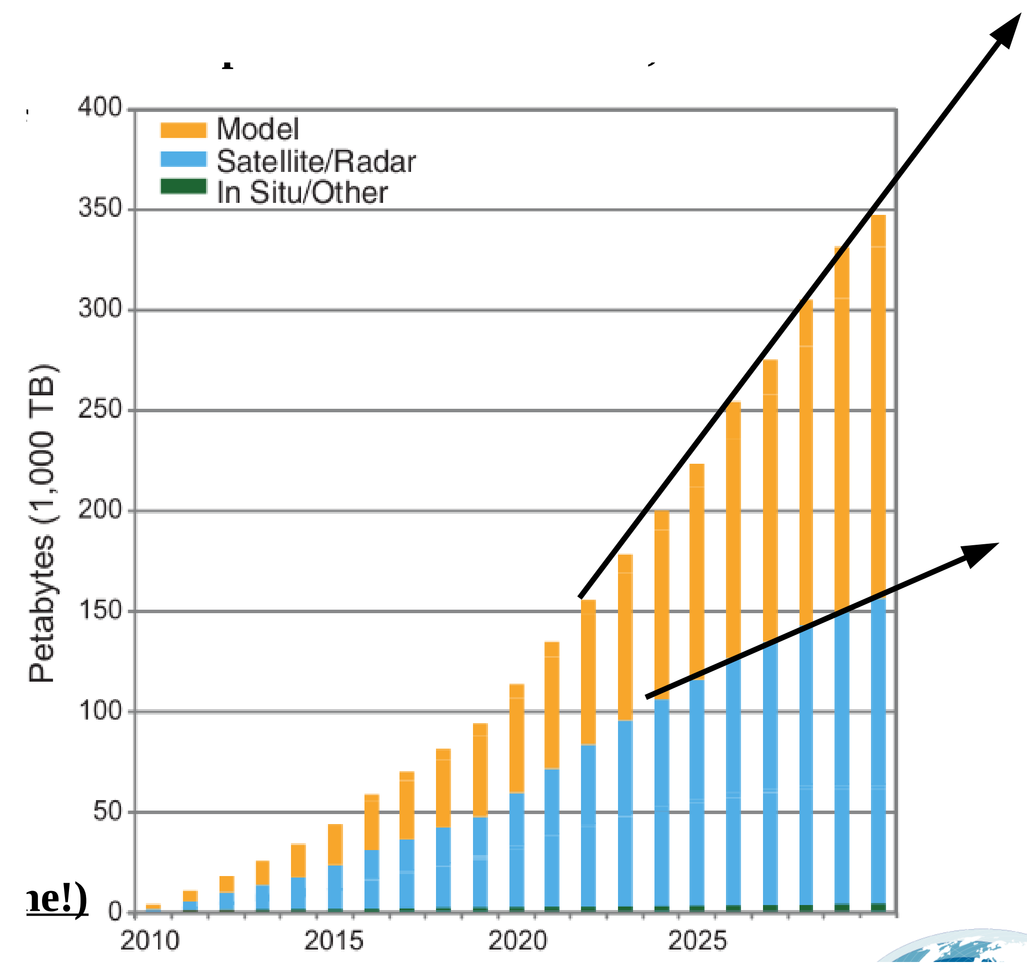

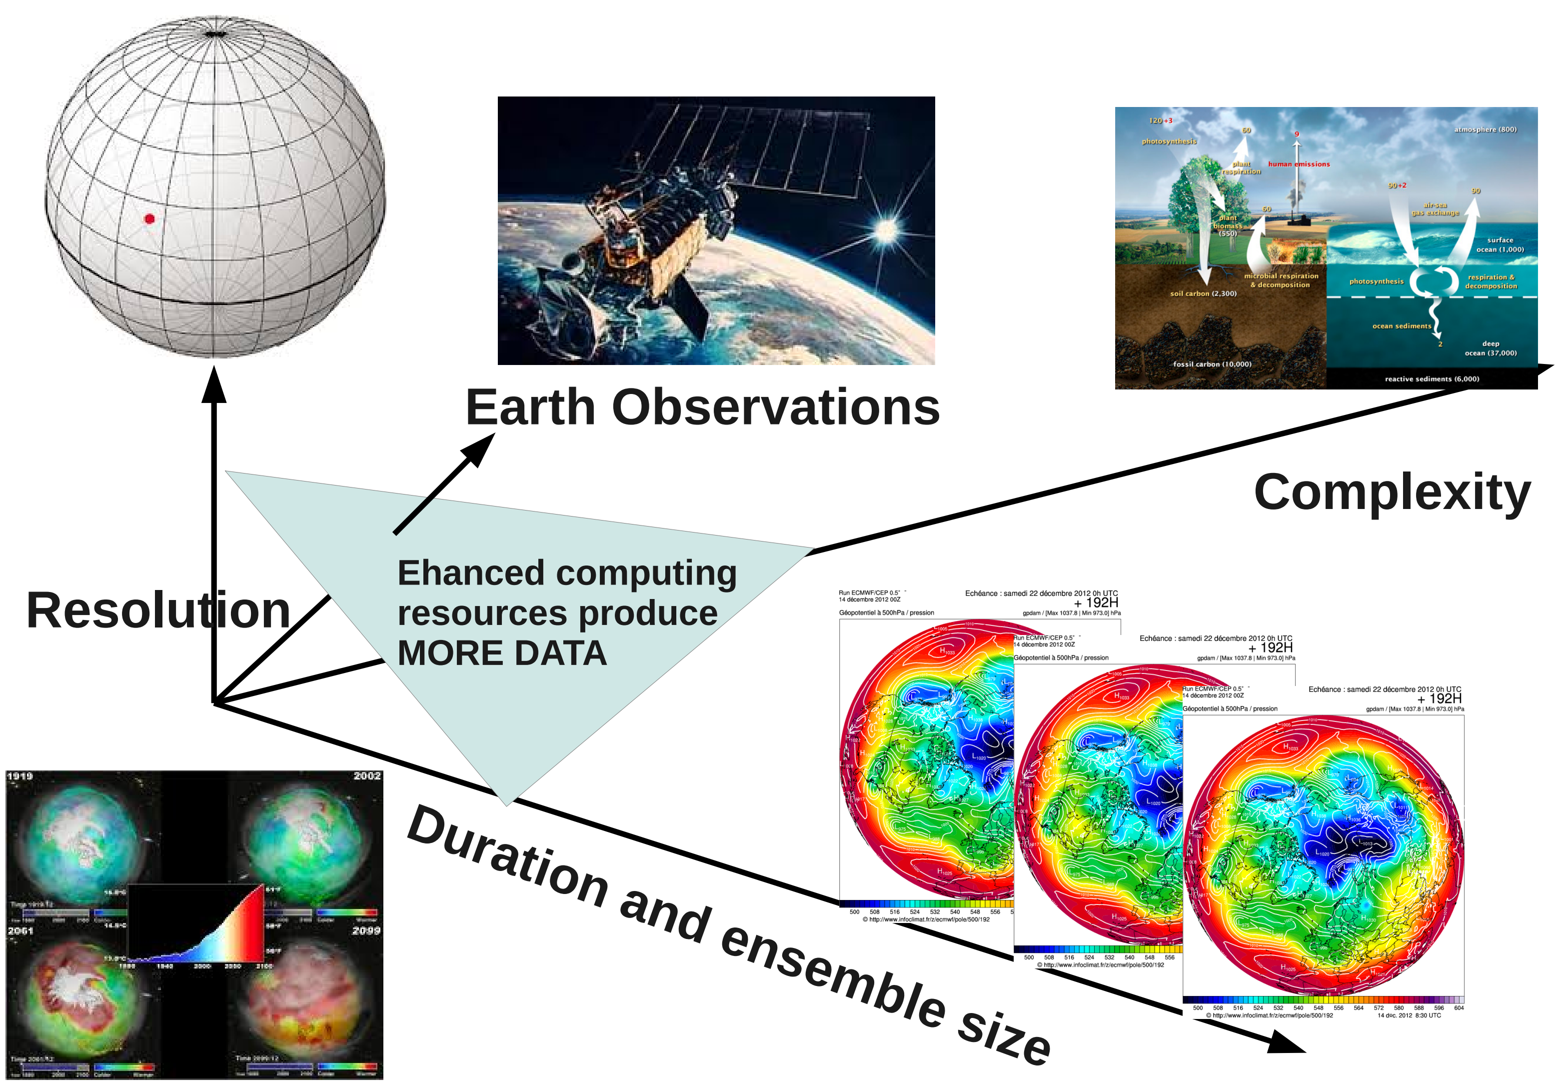

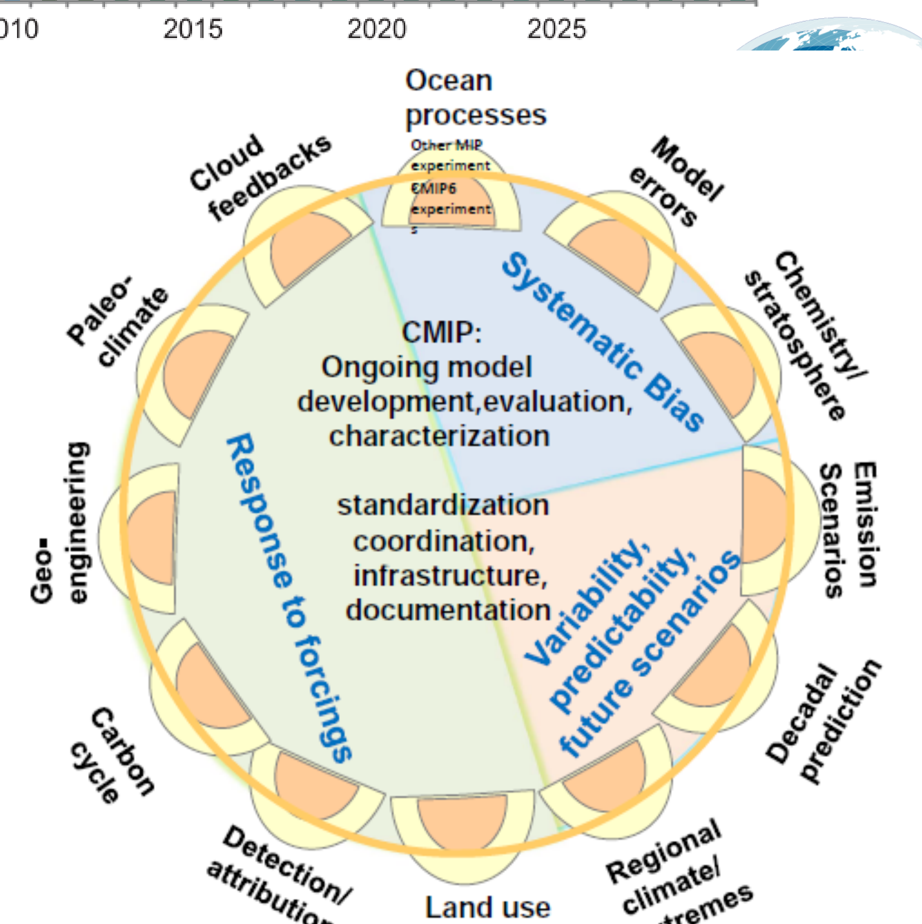

The IPSL-CM application

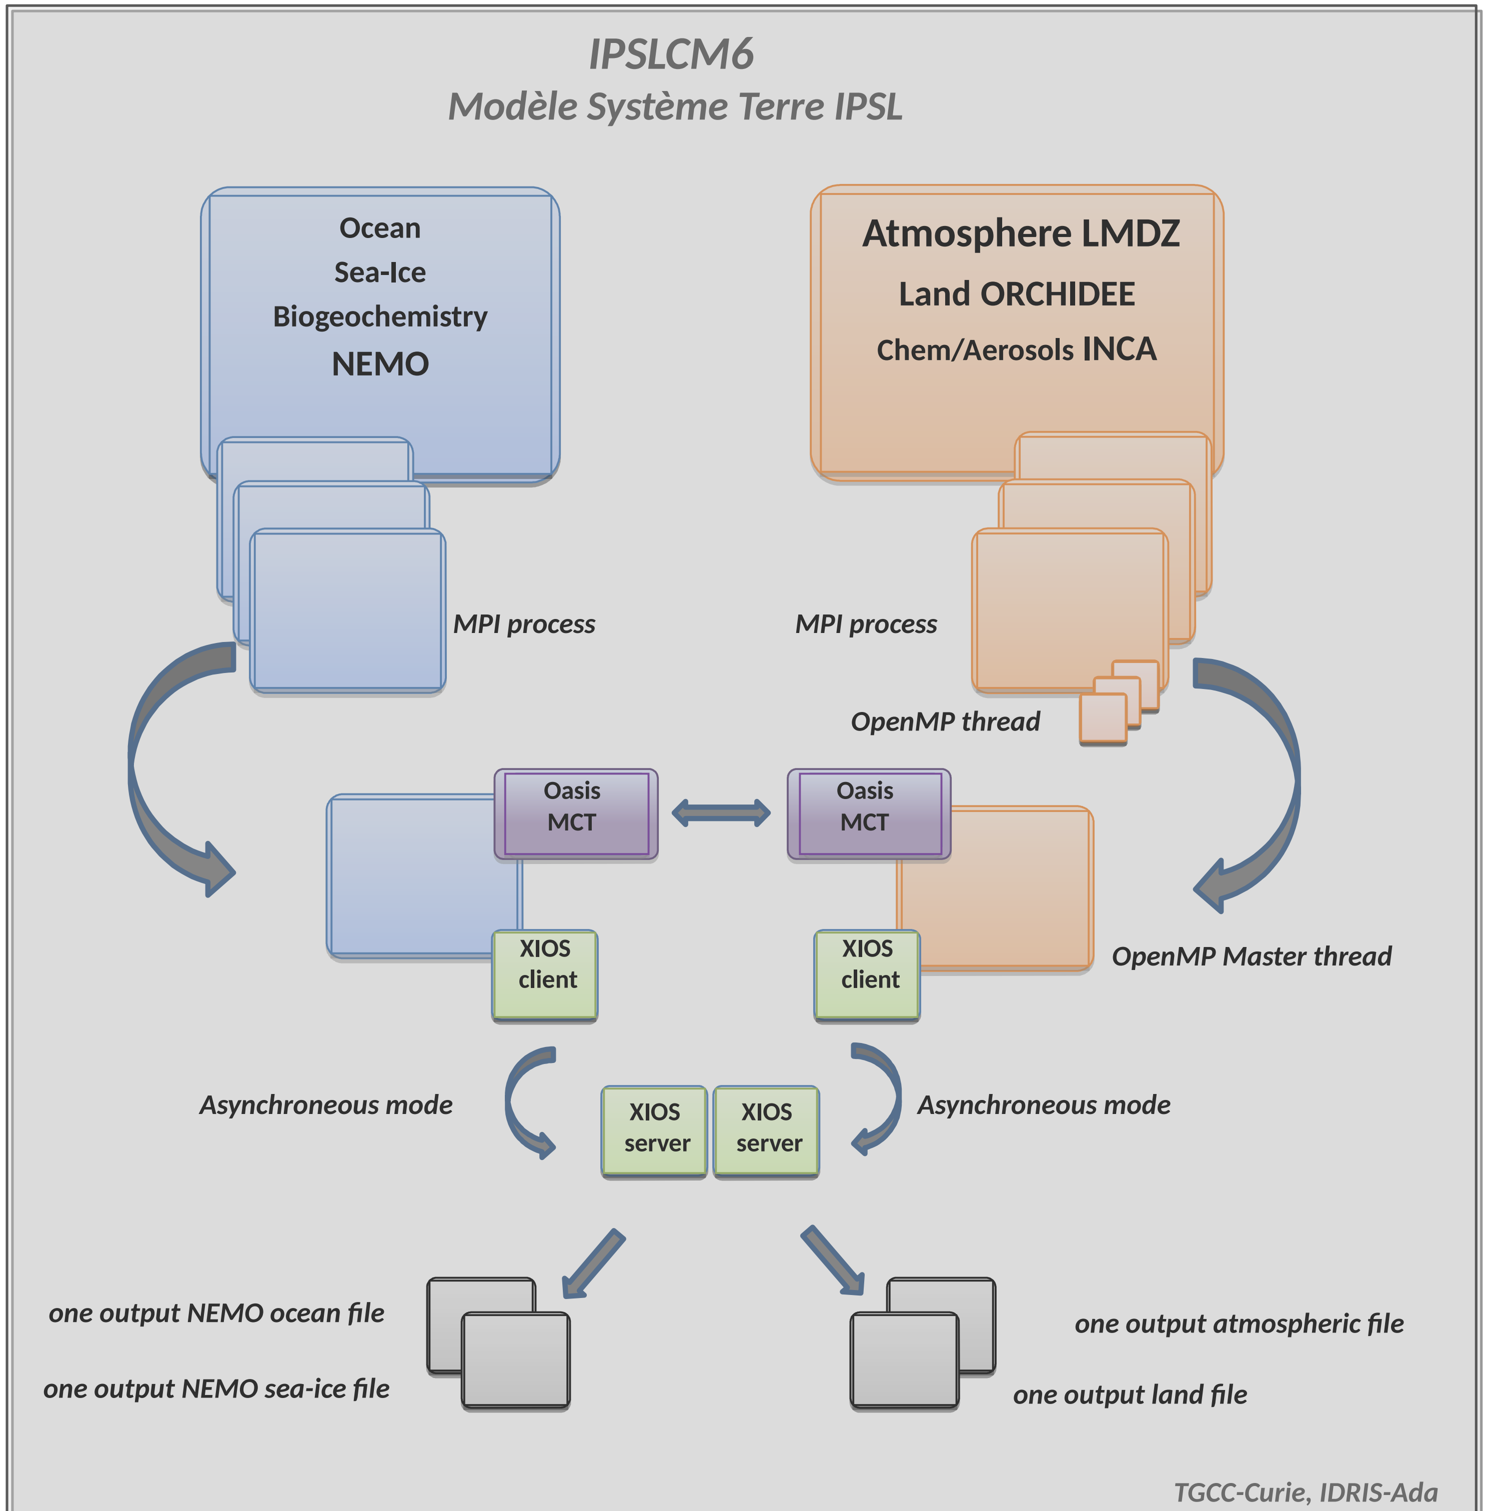

Large number of models with
- a number of **configurations**
- a number of **experiences**
- an ensemble of **realisations**

Large number of variables, and files

Large volume of secondary data

~ 10 PBs scale

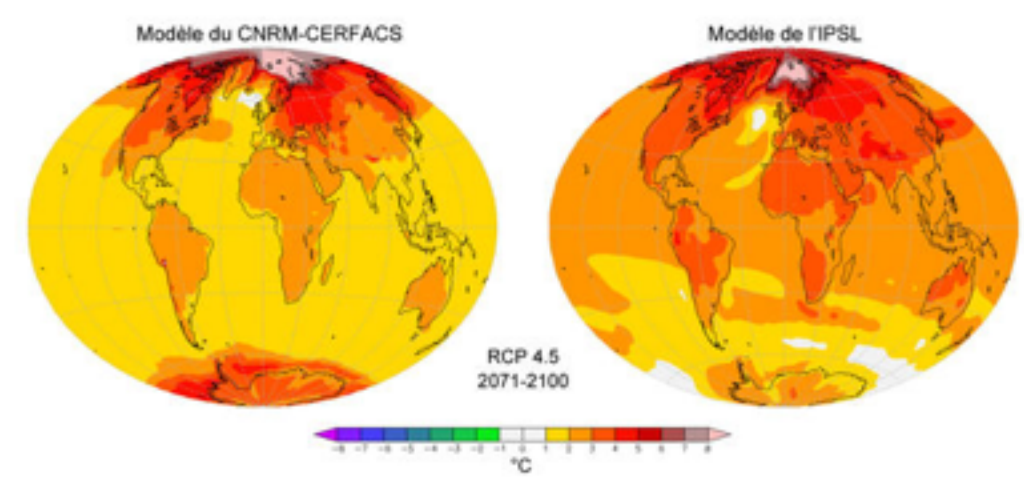

A resilient and flexible runtime environment

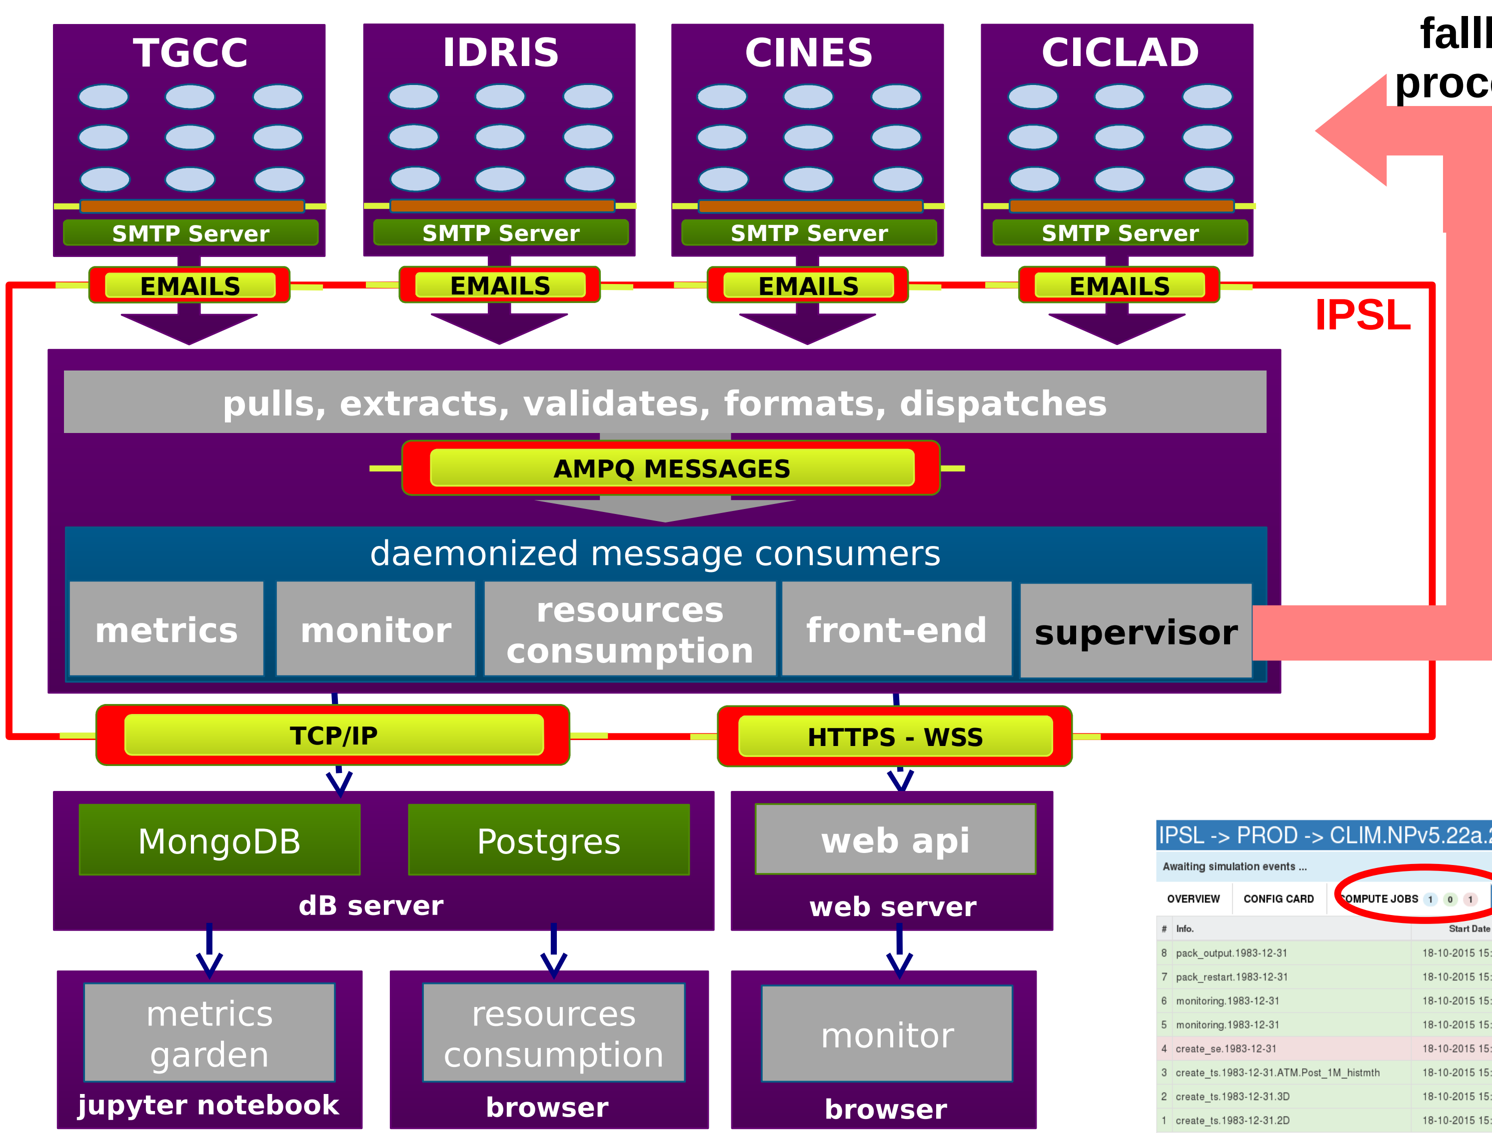

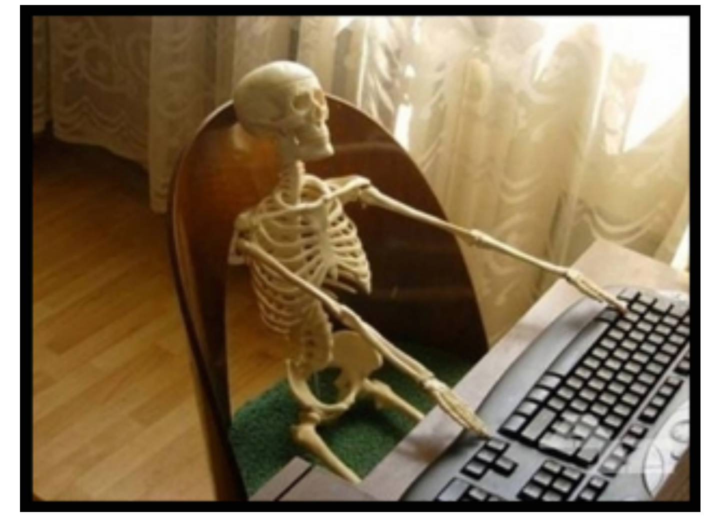

Aujourd'hui :

- 100 utilisateurs
- qqs 10 job/jour

A flexible, resilient and reactive provenance-driven system

- **Providing Run-Time feedback** on the process with **tuneable metadata** and **provenance-driven controlled data movement**
- **Avoiding** useless waits **for long and unfruitful run**
- **Fostering Dynamic Steering, Diagnostics, saving computing cycles, storage and energy ($$)!**

Numerical laboratory: Earth System Grid Federation

~ 10 PBs scale

Web processing service (WPS)

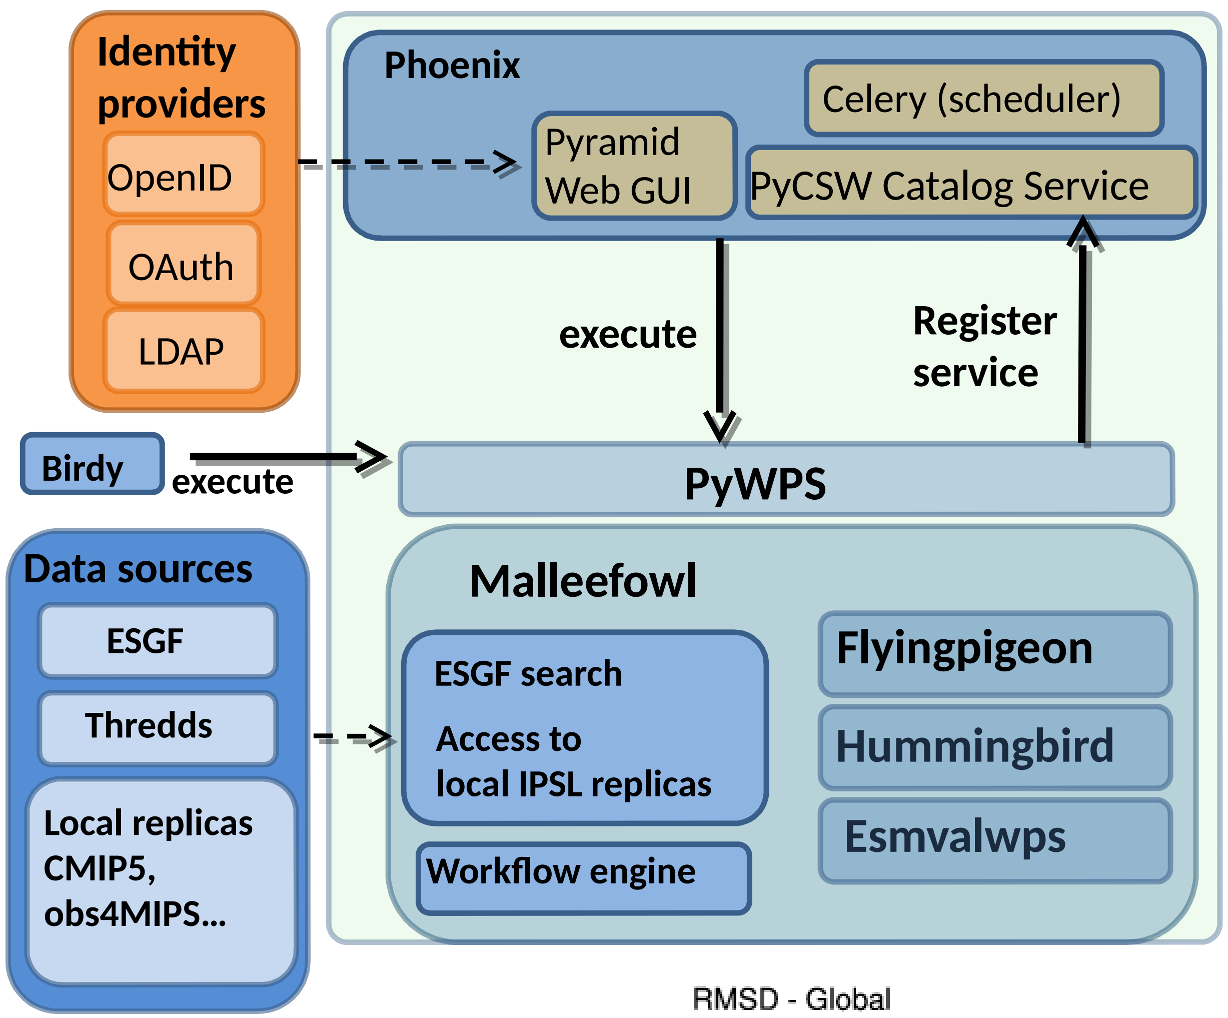

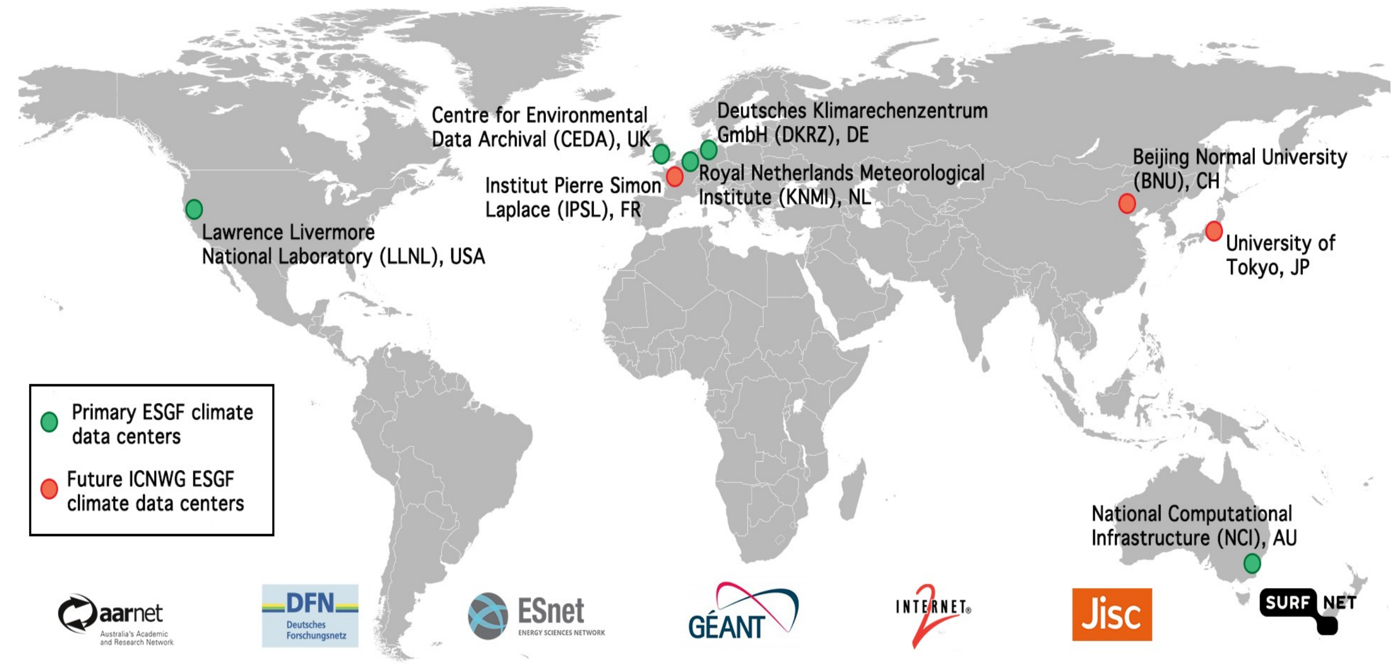

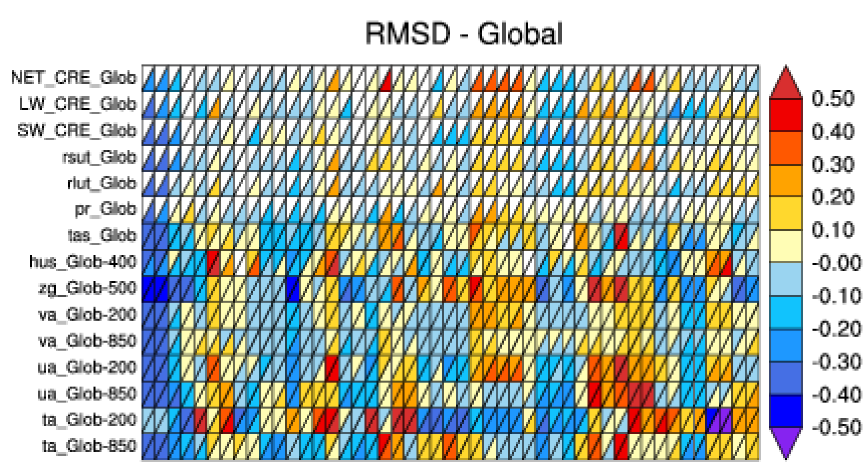

Climate Model Assessment Framework (CLiMAF)

- **Exploration and analysis** of climate simulations
- **Share data processing and analytic** methods and tools
- **Advanced management** of simulations and analysis
- **Induction of a broad community** of researchers and users
- **Accelerate the full path of data use** from capture to delivery of information
- **Web services on top of data analysis platforms**
- **Pervasive provenance system**

IPSL, courtesy of S. DENVIL

From HPC simulations to data-intensive platforms

Largest simulations at the petabytes scale

- From regional to global scales (climate, seismology, magnetohydrodynamics, etc.)
- From supernovae to turbulence
- Need for community access/reuse of the best and the latest secondary data through numerical laboratories with pervasive provenance system

Create new challenges:

- How to **move/output data during simulations** (vertical re-use, I/Os, parallel storage)
- How to **reduce data** through **in-situ analytics**
- How to **stag in and stag out data** (high-speed transfer protocol, access policies)
- How to **explore/visualise data** (render on top of the data, immersive analysis)
- How to **analyse/instruments data** (data analytics, immersive analysis, value added services …)

Research-driven

Huge variations in Data lifecycle and commitments

- **On-the-fly (in-situ) analysis and visualisation** (immediate, do not keep)
- **Collaborative reuse and analysis** secondary data (short/mid term, local)
- **Community services and analytic tools** (mid /long term, community commitment)
- **Archival, curation, provenance, trust** of secondary data (long term, community commitment)

Different from today supercomputer usage and access policies

A variety of data and computing resource access patterns

! Dump data into large HPC providers, move data out to analysis platforms

Compute and Data-analysis federated infrastructures

a research-driven strategy

A research-driven variety of infrastructures

Ashby's Law of Requisite Variety

Only variety absorbs variety

Data-intensive analysis platform and HPC

- Caches and persistent caching storage close to data-intensive analyse platform
- Data-intensive computing architectures and HPC simulation architectures
- Render on top of the data together with value added services, data analytics
- Induction to a broad research and user community (access and security)

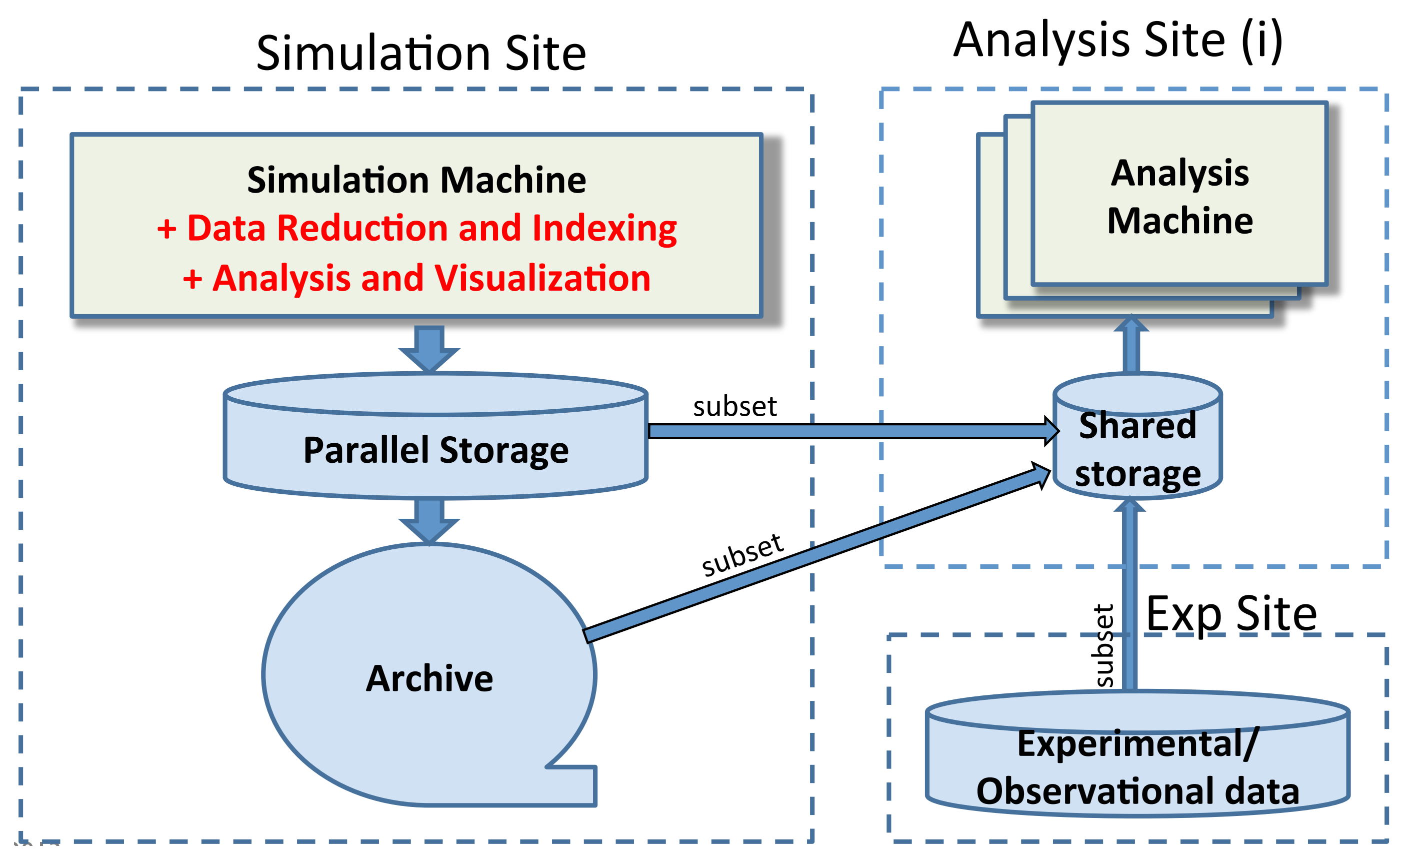

Bridges and Gateways

System-building: technology-based and research-driven services

Technology transfer across domains and locations: variations of original design and emergence of competing systems

Gateways consolidation:

research-driven technical solution with social choice integrated within research communities of practice

federation of dissimilar autonomous systems into research-driven networks

(adapted from Edwards et al., 2007)

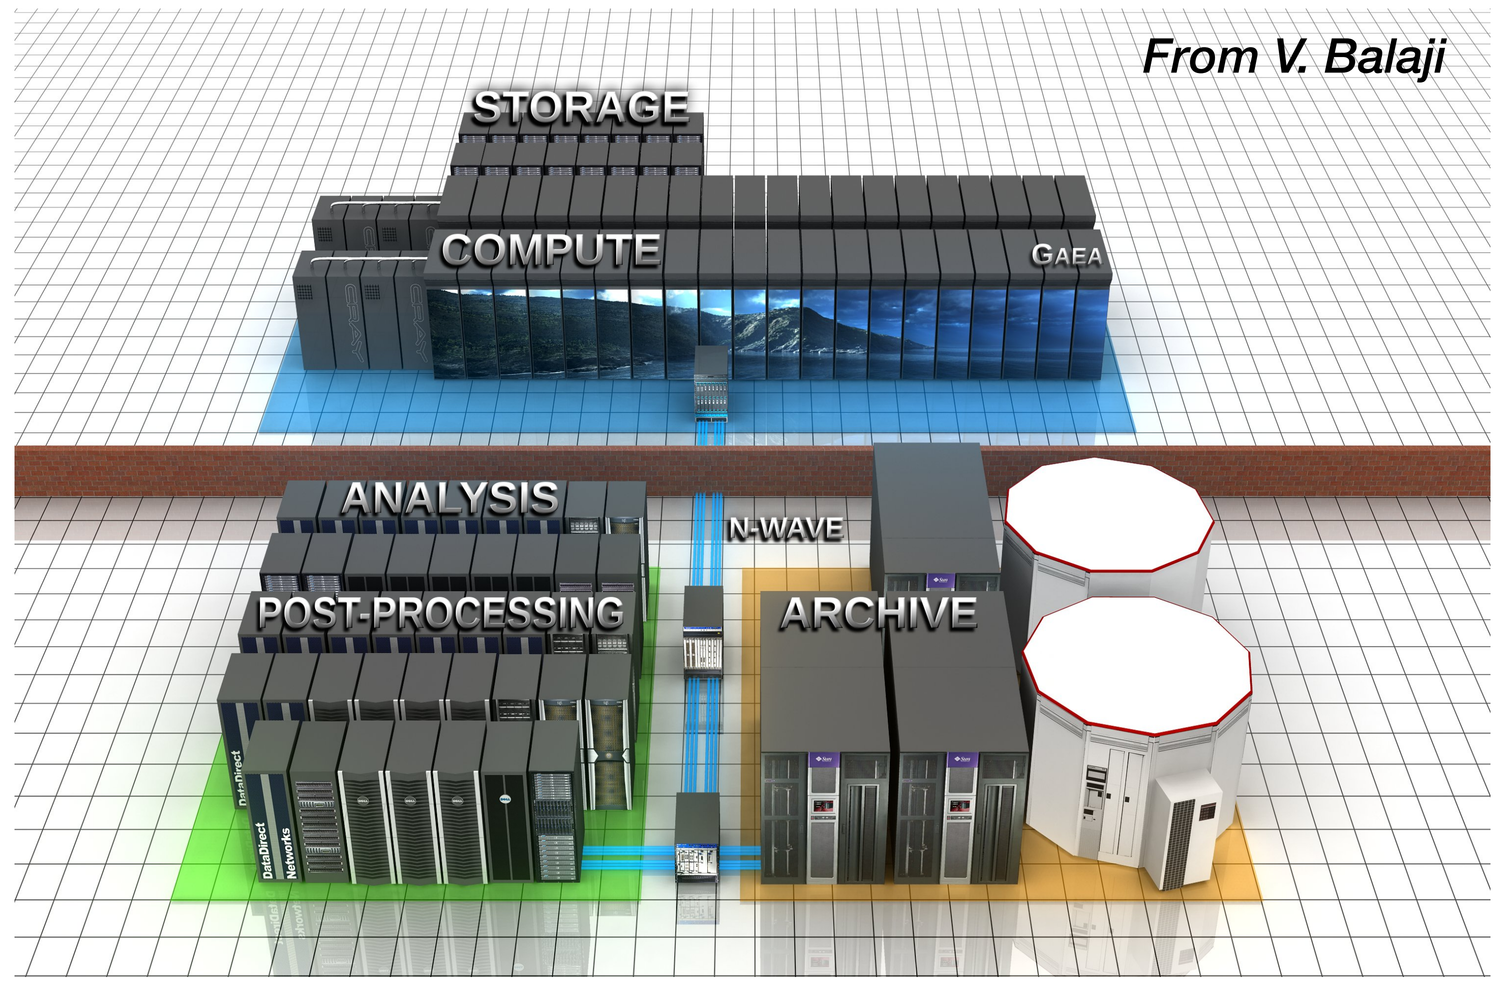

E-Infrastructure challenges and strategy

System and infrastructure Big data analysis

Where should the caches and persistent storage be ?

- Caches and *persistent* caching storage: sharing large chunk of observations and simulated data, optimise costs of moving data
- Not directly at the supercomputer (too expansive storage)
- Analysis computations and visualisations on top of the data
- High-speed transfer protocols from/to data sources (HPC, large instruments, data archives)

Complex data movements scheduling

- Data and metadata bases (scalability)
- Provenance-driven triggering and management
- Extended file management systems and model
- Augmented services with added-value to large community

Data organisation

- Most of these data are not hard to partition (scale-out)
- Provenance management system and lineage metadata
- Fined grained data streaming flows
- Tier of large memory systems (random access)

Challenges and strategies

✳ **Difficulty getting things to run in multiple providers contexts**

- explore new virtualisation technology and Sand boxed environment
  - Linux Containers, docker, Google Kubernetes
- support software developments and maintenance
  - prepare the new generation of HPC architectures (exascale challenge, in-situ data analytics)

✳ **Difficulty to provide uniform access, trust and security model**

- leveraging existing identification systems across infrastructures

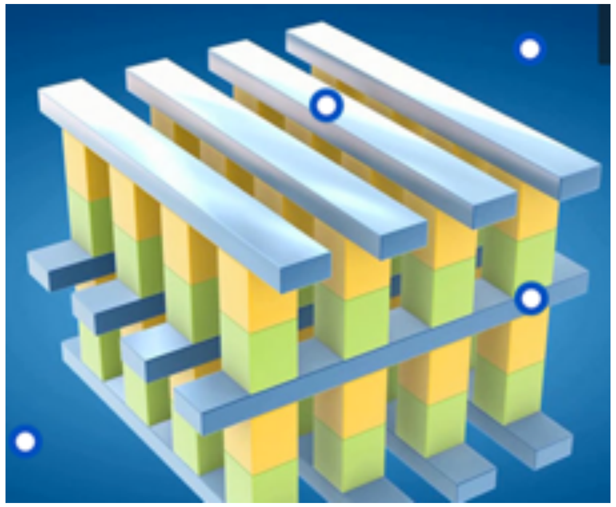

✳ **Difficulty handling data and computation strategy**

✤ reduce computational costs

- well-matched architectures to each stages

✤ reduce data movement costs

- in-situ analytics, persistant storage, caching strategy, compression

✤ re-use of calculations and data

- effective metadata and provenance system information

✳ **Align methods with research infrastructures**

- balanced and aligned investment for the full path of data use
  - maximise overall value of generating, collecting, preserving, curating data

✳ **HPC and Data Infrastructures tailored by scientific use cases**

- a variety of access and usage patterns requirements

✳ **Interdisciplinary task forces**

- share mutual understanding of methods and technologies (Astrophysics, Climate, …)
- Interdisciplinary task forces
- Computer scientists must meet flexible federation challenges

Data-intensive analysis platforms

A scientific e-science environment capable of "observing" (explore, analyse and model) massive and complex data generated by large-scale instruments, observation and monitoring systems, and numerical simulations in the sciences of Universe.

- Innovative methods, software, ICTs for large scale data-intensive computation and massive data statistical analysis that ultimately **induce a broad base of researchers to new research practices**
- Emergence of **cross-disciplinary expertise** in data-intensive computing and data analytics across scientific domains, research informatics, HPC and Data system engineers
- Accelerates full data use path: **valorisation of massive data generated by large-scale instruments, observation and monitoring systems**
- **Training and `intellectual ramps"** to engage a **new generation of researchers** to harvest data capabilities in their research practices to address new research challenges
- Community building around **simulation and data analytics shared application-software** together with **provenance and services** for open research and application science
- Consider **full path of data use and data life cycles** -> federation of HPC and data-intensive analysis platforms

A flexible and scalable federation of autonomous infrastructure providers/organisations

Data resources - Data-intensive analysis platforms - HPC infrastructures

E-infrastructure and data Management CRA

BELMONT FORUM

5/1/2015

A Place to Stand: e-Infrastructure and Data Management for Global Change Research

Belmont Forum e-Infrastructures & Data Management Community Strategy and Implementation Plan

"Give me a place to stand, and I will move the world"

- Archimedes

E-Infrastructures and Data Management Steering Committee

Data-intensive e-Infrastructure Action Theme 3

- Identify and fund interdisciplinary use-cases for federated data- and e-infrastructures in environmental and global change challenges.
- Identify and fund large-scale Data and Model Inter-comparison Projects (DMIP) that are relevant to global change research.
- Through the above outcomes, inform data- and e-infrastructure policy with case-proven best practices that respond to concrete issues.

Milestones

29-31 August 2016, Paris, France: 2 ***scoping workshops in Paris***: cross- and trans-disciplinary data-intensive use cases analysis; data and model inter-comparison (DMI) use cases

10-16 September 2016, Denver, Colorado: ***International Data Week***: 2017 Belmont Call finalisation for data-intensive cross and trans disciplinary use cases and DMI projects

http://www.bfe-inf.org

Contact: Mark Asch (ANR, France),mark.asch@agencerecherche.fr