GRICAD

GRENOBLE ALPES RECHERCHE
INFRASTRUCTURE DE
CALCUL INTENSIF
ET DE DONNÉES

CIMENT

Retour d'expérience iRODS

Bruno Bzeznik
Responsable du pôle calcul de GRICAD

13 decembre 2016 - Grenoble

Irods introduction

http://slides.com/irods

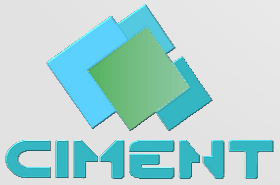

CIMENT : High Performance Computing center of the univ. Grenoble-Alpes

UNIVERSITÉ Grenoble Alpes

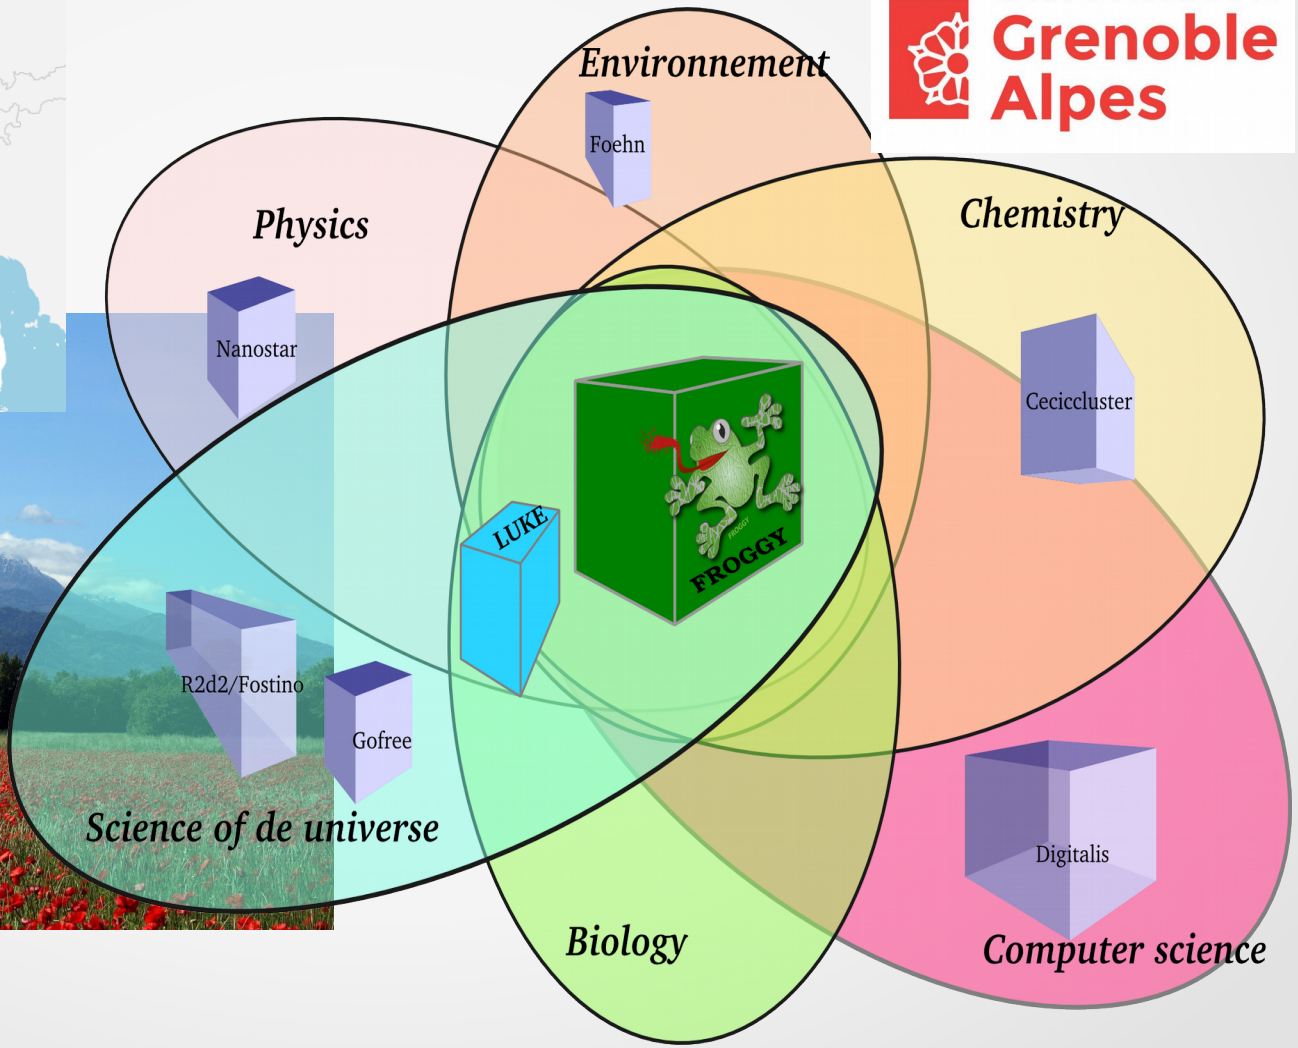

CIMENT

Platforms of the HPC center of the University of Grenoble

CiGri lightweight computing grid

OAR batch scheduler

HPC platform

Froggy

3200 Xeon E5 cores @2.6Ghz
+18 GPUS K20m

High performance distributed storage (Lustre): 90 TB

Infiniband FDR network

Remote visu nodes

OAR batch scheduler

Data processing platform

Luke

~600 cores – heterogeneous systems and continuously evolving

Local scratches on nodes
450 TB

10 Gbe network

Remote visu nodes

OAR batch scheduler

Other thematic platforms

~3000 cores heterogeneous systems federated from 10 clusters of member laboratories

NFS filesystems:
a few TB per cluster

Infiniband QDR network

Common distributed storage (IRODS) 1Po

The IRODS infrastructure setup in Ciment

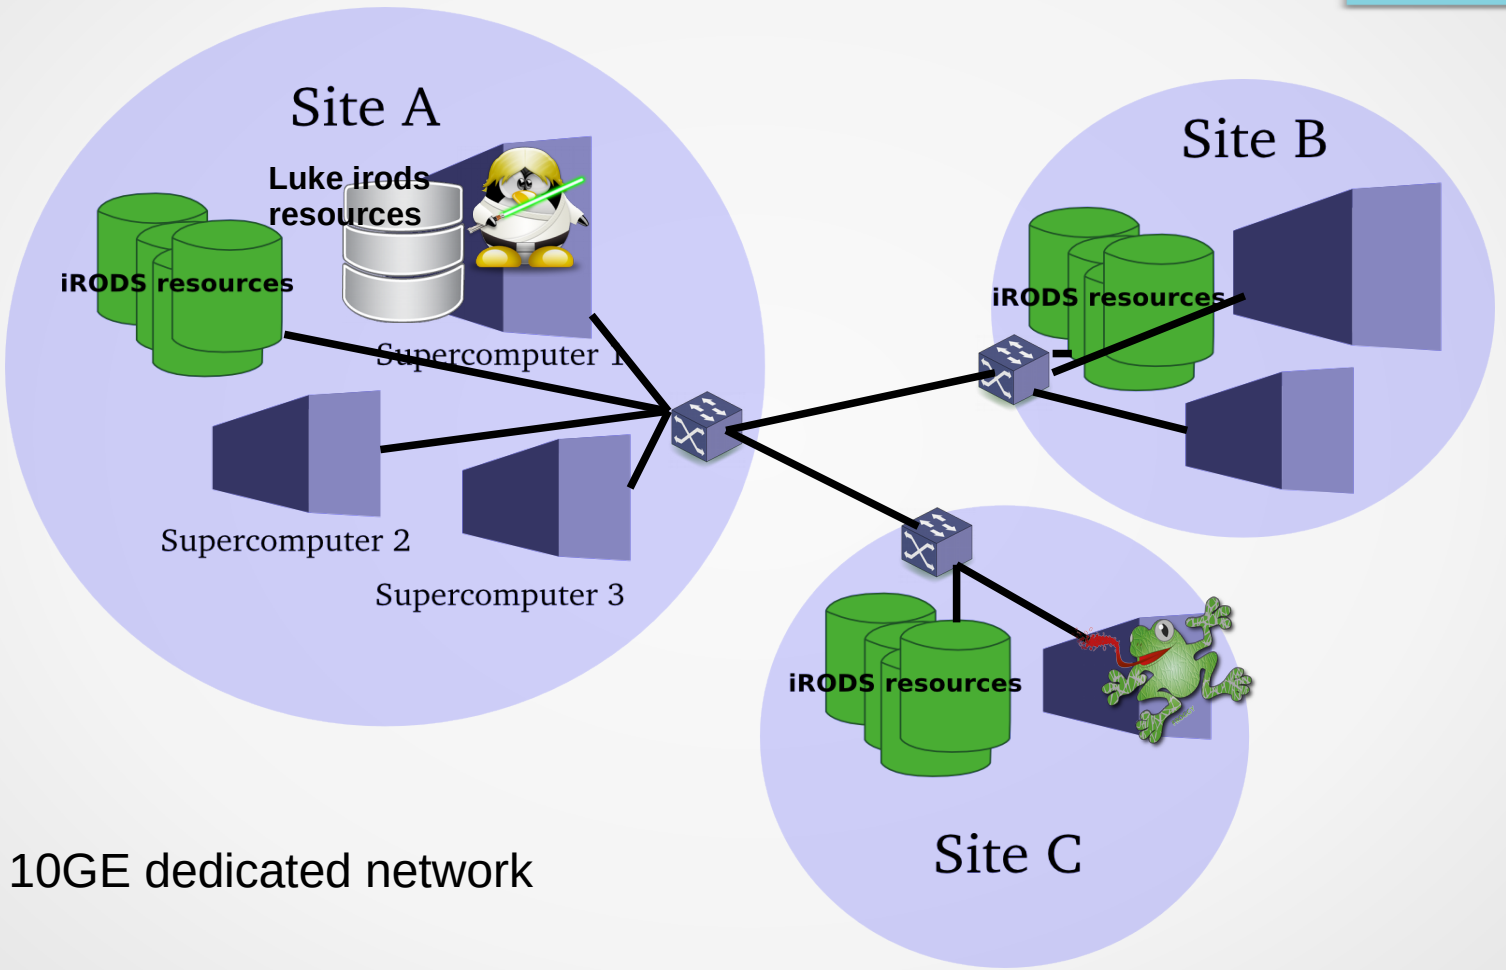

CIGRI and IRODS

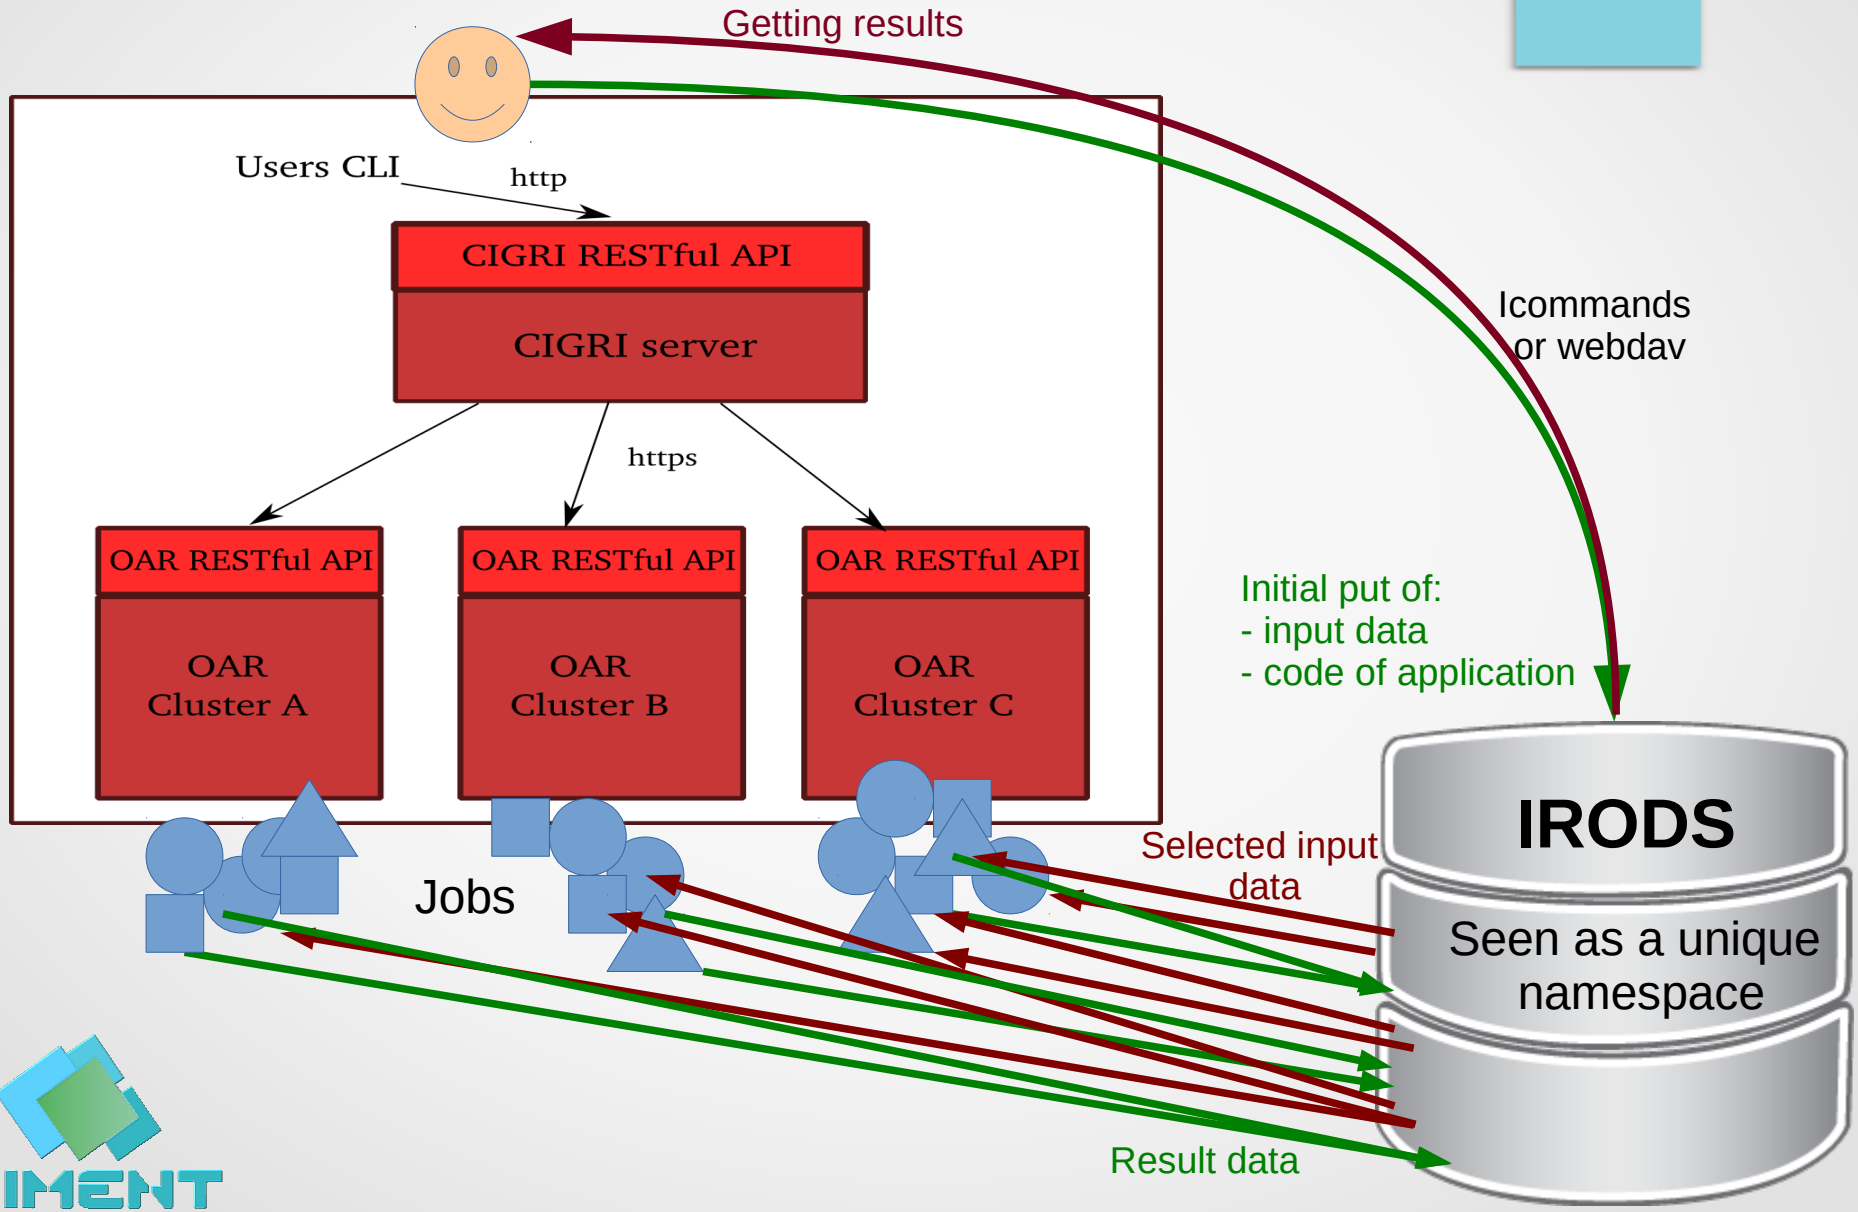

iRODS history at CIMENT

- 2010: joint project with LPSC. “Institut des grilles” funding.
  - iRODS v2
  - 288 TB (Dell R510 + MD1200)
- 2012:
  - iRODS v3
  - Labs contributions → 450TB
- Today, decembre 2016
  - IRODS v4.2
  - Continuous labs contributions → > 1Po
  - IRODS not only used for CiGri, but also for sharing scientific data, long term scratching,...

CIMENT

Going in and out from the iRODS cloud

- “Cargo” resource:
  - Scp, http, ftp,… + registration (ireg)
- Webdav interface
  - https://github.com/UtrechtUniversity/davrods
- iRods cloud browser (web interface)
- iRods zones

CIMENT

CIMENT IRODS in numbers

- December 2016 status:
  - 188.590.693 files
  - 5.470.142 collections (directories)
  - 342 TB used / 1PB total
  - 29 resources (20-80 TB RAID arrays)
  - 15 iServers (+ 5 on Luke)
  - Continuously increasing
- 6 typical months on cigri:
  - Number of Cigri jobs: 2,7 millions
  - Number of IRODS transactions: 6,6 millions
    → average ~ 43000 transactions / day

CIMENT

CIMENT IRODS in numbers

Near 200M files counting

```
$ time iquest "select count(DATA_ID)"
DATA_ID = 188593149
real   0m36.816s
```

- A ‘du’ on a directory containing 3M files? No problemo

```
$ time iquest "select count(DATA_ID) where COLL_NAME like
'/cigri/home/isterre%'"
DATA_ID = 3325211

real   0m2.168s

$ time iquest "select sum(DATA_SIZE) where COLL_NAME like
'/cigri/home/isterre%'"
DATA_SIZE = 24584072995889

real   0m16.252s
```

CIMENT

CIMENT contributions to IRODS

- Participation to the iRods users group (once a year at Chapel Hill)
- Bash completion script
- Real life performance tests with big collections in an HPC environment.
- Using “connection control”
- NIX packaging (irods 4.2)
- Will suggest an incremental retry option for each icommand until actually coded (or suggest a patch...)

CIMENT

Scientific use cases

CIMENT

Seismology : Whisper, European seismological project

Project: Detect sligth changes of properties in the solid Earth

Data : Noise Continuously recorded by seismic stations worldwilde.
The computations produce even more data
More than 200 TB managed at the same time

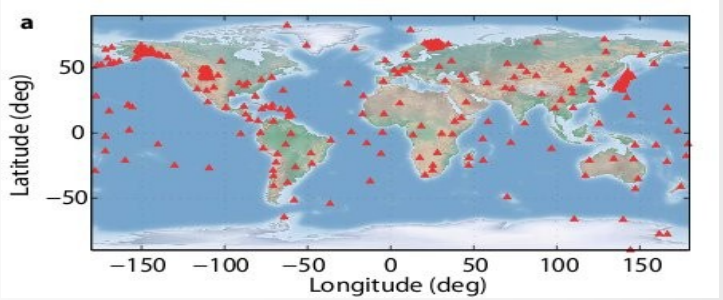

Data Intensive processing :
Use intensively the data grid environement.
Specific python library is devellopped
Lot of feedback on cigri and irods

Scientific :
Many papers, posdocs and students use the data grid environment for whisper

About varation of velocity change
of the tohoku earthquake in Japan (Science)

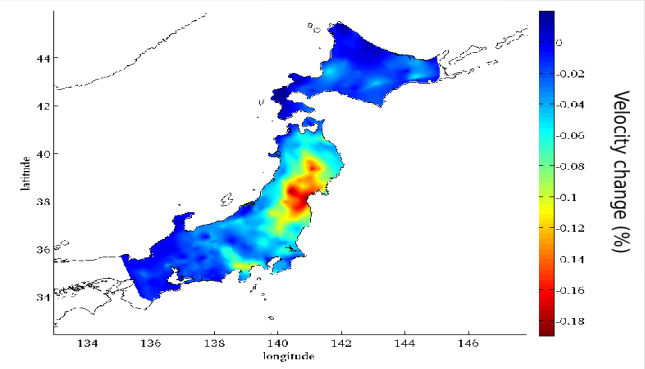

CIMENT

Rosetta / CONSERT

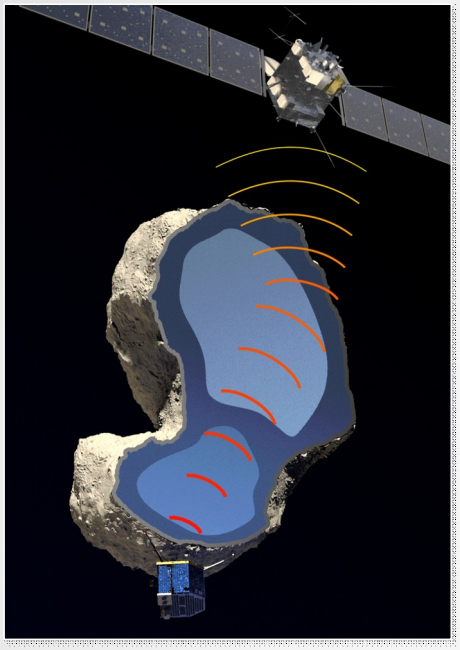

COmet Nucleus Sounding by Radiowave Transmission

An experiment on-board **Rosetta** of the European Space Agency

Performing radar tomography of the comet nucleus
of 67P/Churyumov-Gerasimenko

CONSERT radio wavefront propagating through nucleus towards Rosetta

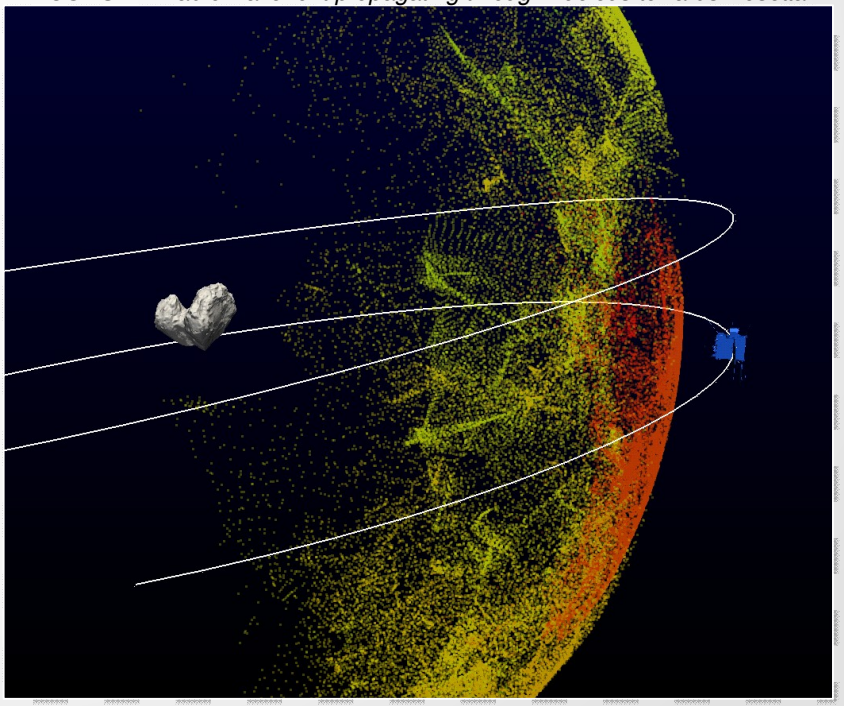

CIMENT with iRods were used for:

- preparation of space operations, and especially for Philae landing (12 Nov. 2014),
- inversion of dielectric properties, deriving better knowledge on composition and structure.

Ecology : *The geography of evolutionary convergences*

Principle

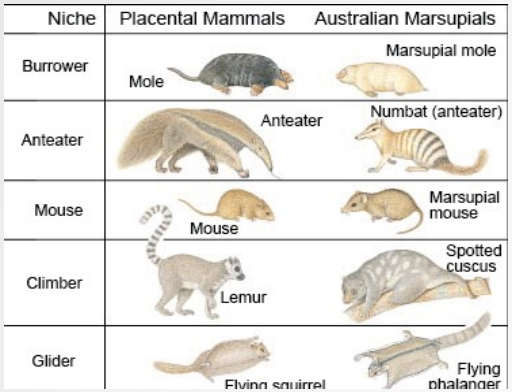

Computations

- Measure of morphological and species similarities between sites
  ⇒ 6 000 000 values

- Detect assemblages that morphologically ressemble each other but contain very different species

Data : 3600 pixels / 5000 mammals

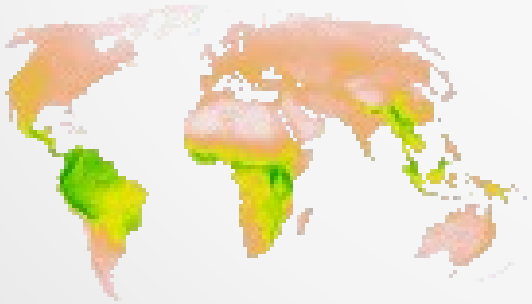

Results

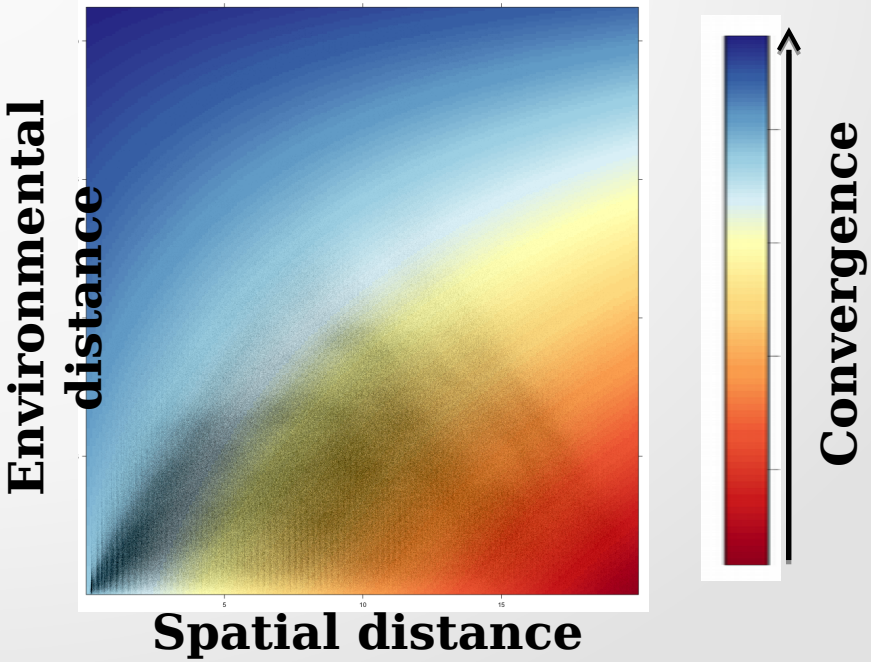

Funding

Particle Physics - LHC

First stable proton collisions at 13 TeV
June 3rd 2015

ATLAS EXPERIMENT

Computing model in Particle Physics with Colliders

- Event by event computation → grid computing is ideal
- The LHC experiments use a grid of ~ 160 computing centres around the world (WLCG)
- CIGRI+iRODS : used as a local farm for ATLAS analyses lead in Grenoble (LPSC/CNRS)

An new area has just began, an un-preceded high energy Physics goal: hunt for exotic particles

Analysis on CIGRI for ATLAS

- Search for extra dimensions in di-photon final states
- Event cross section computation
- “CIGRI is an asset”

- Already used for the earlier phase of the LHC (Run 1)


New Journal of Physics
The open access journal at the forefront of physics

This is to certify that the article

Search for extra dimensions in diphoton events from proton–proton collisions at $\sqrt{s}$ = 7 TeV in the ATLAS detector at the LHC
by The ATLAS Collaboration

has been selected by the editors of *New Journal of Physics* for inclusion in the exclusive 'Highlights of 2013' collection. Papers are chosen on the basis of referee endorsement, novelty, scientific impact and broadness of appeal.

Professor Eberhard Bodenschatz
Editor-in-Chief
New Journal of Physics
www.njp.org

DPG | IOP Institute of Physics

NJP 15, 043007 (2013)

At the Tevatron (US)

Phys. Rev. Lett. 108, 151804