DIR group: methods and resources

Paolo Besana, 5th August
University of Edinburgh

What is Data Intensive?

- Data volume is increasing:
  - Cost of sensors is decreasing (better CCDs, higher sampling frequency,…)
  - Ubiquitous networks allow collection and access to data
  - Better HPC allows more detailed simulations (for example, fine grained simulation of geodynamics or climate)

Example of data-intensive

- Sloan Digital Sky Survey (1992-2008)
  - 2.5 Terapixels of images from telescopes
  - 40TB of raw data, 120TB processed
  - 5TB of catalog, 35TB in the end
- LOFAR project (2012-2017)
  - Large and expensive radio telescope dishes are replaced by many cheap dipoles
    - no mechanical movements:
    - the beams are steered electronically in milliseconds
  - The theoretical output is 1.6TB/s, yielding 38PB/day raw data

Where is the problem?

Transferring and processing large amount of data is problematic

Hardware perspective:

- CPUs waste time waiting for data to process

Software perspective

- Many algorithms require in-memory processing
- Workflow languages:
  - assume flow of data to and from web services
  - based on the idea of bringing data to the computations, assuming CPUs are the scarce resource

View on data intensive problems

Requires the interactions between experts

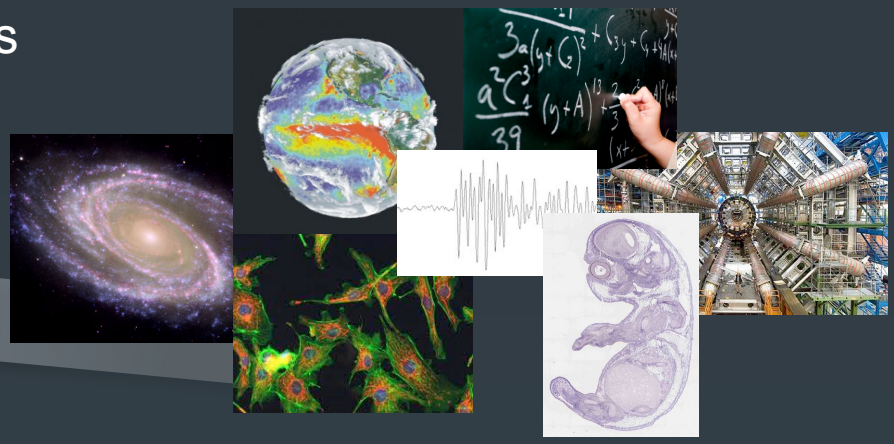

Domain experts

Data-Analysis experts

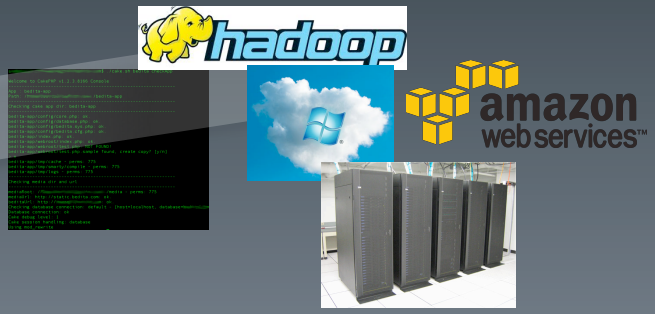

Data-intensive engineers

DIR Group

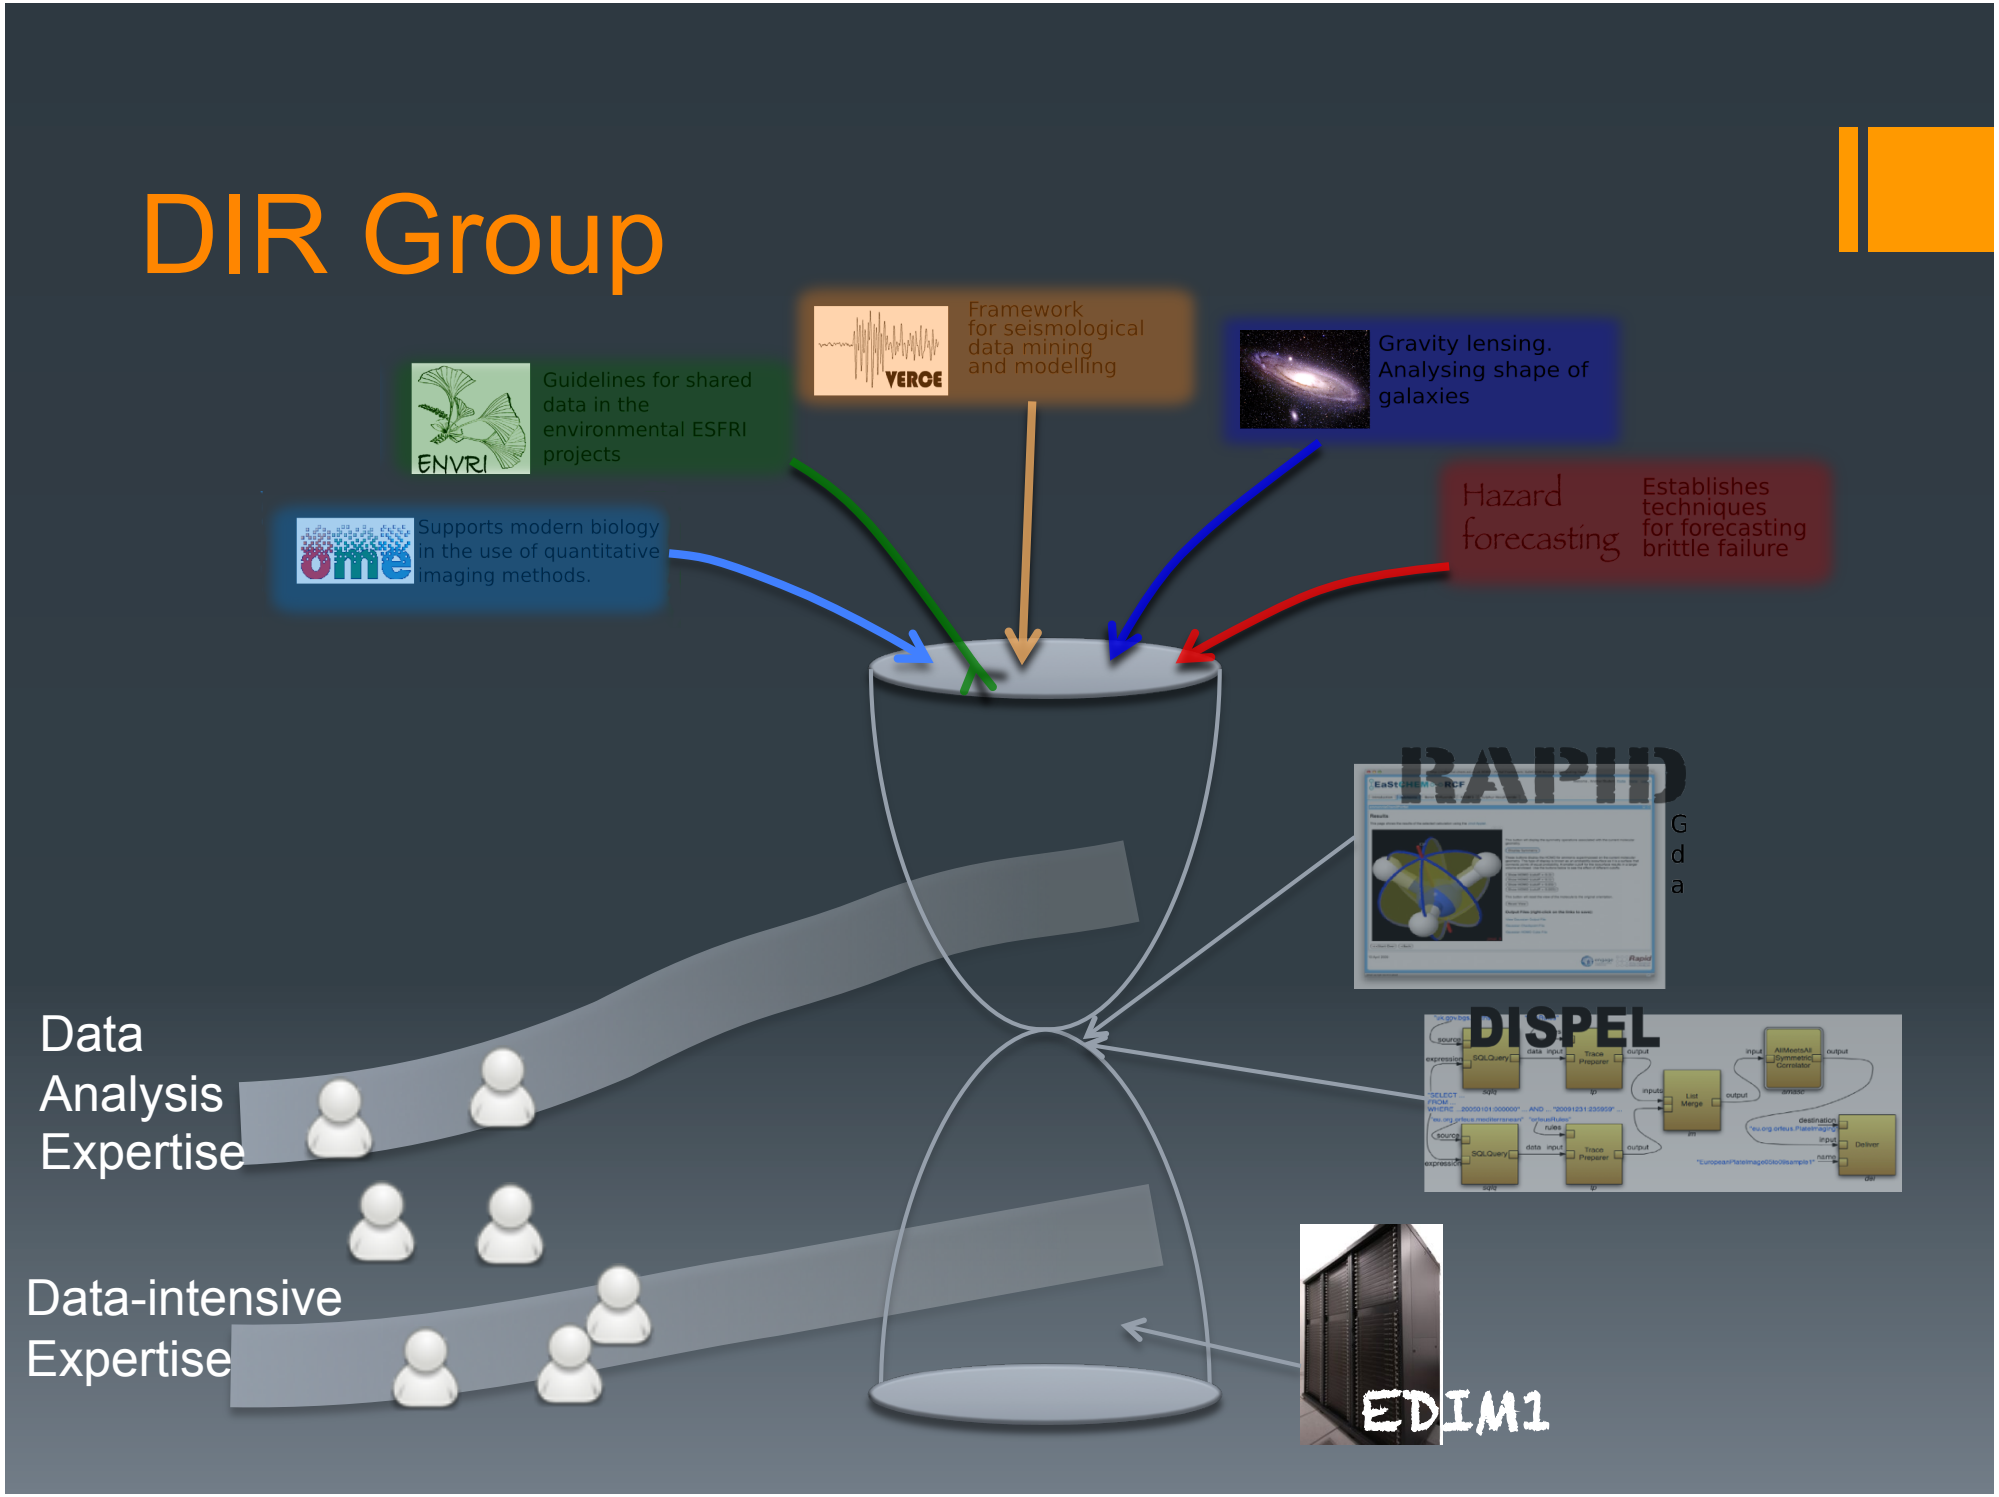
ENVRI
Guidelines for shared data in the environmental ESFRI projects
VERCE
Framework for seismological data mining and modelling
Gravity lensing. Analysing shape of galaxies
ome
Supports modern biology in the use of quantitative imaging methods.
Hazard forecasting
Establishes techniques for forecasting brittle failure
RAPID
DISPEL
EDIM1
Data Analysis Expertise
Data-intensive Expertise

SW architectures

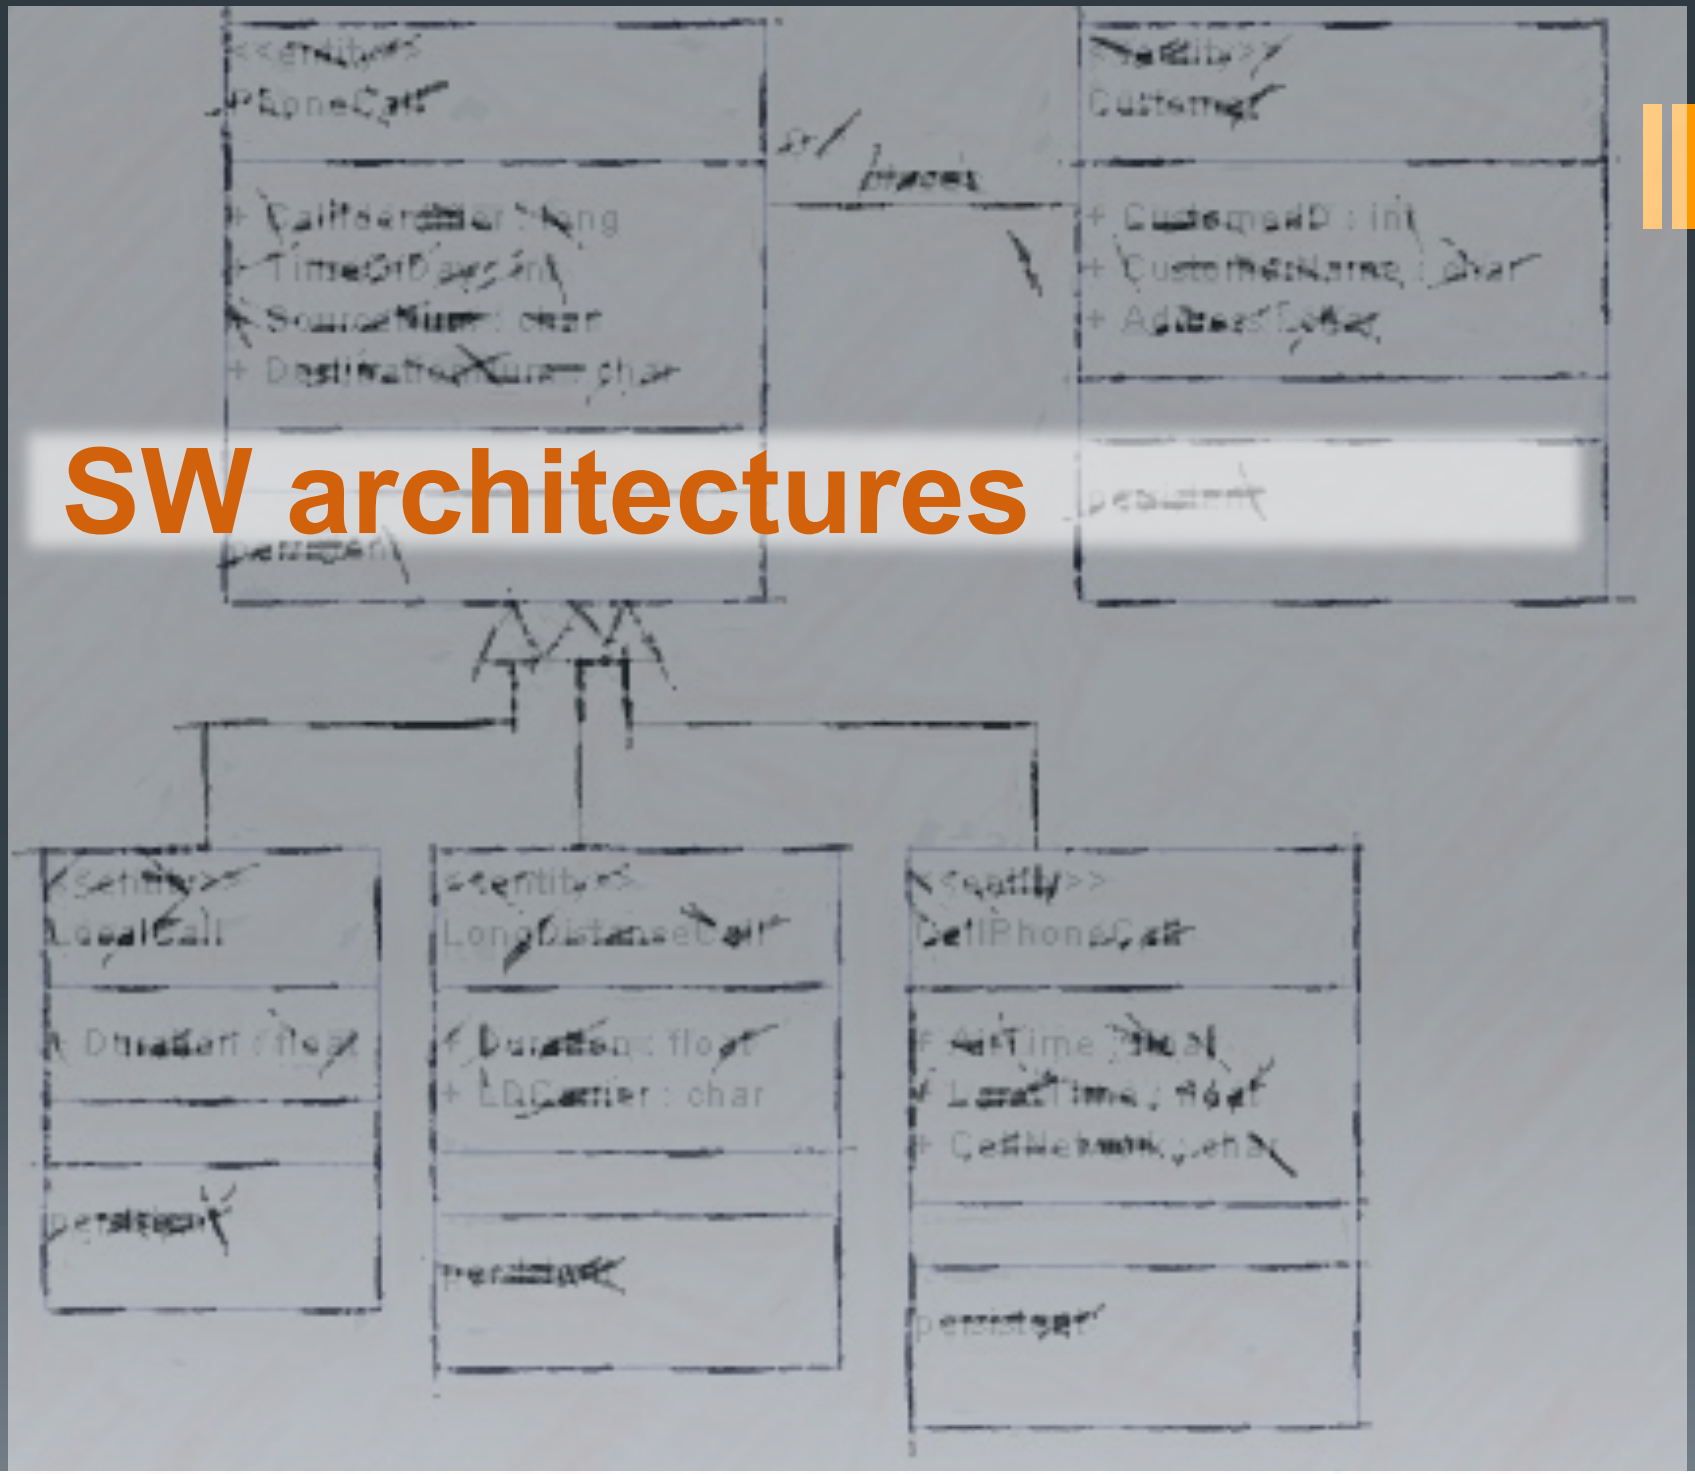

DISPEL

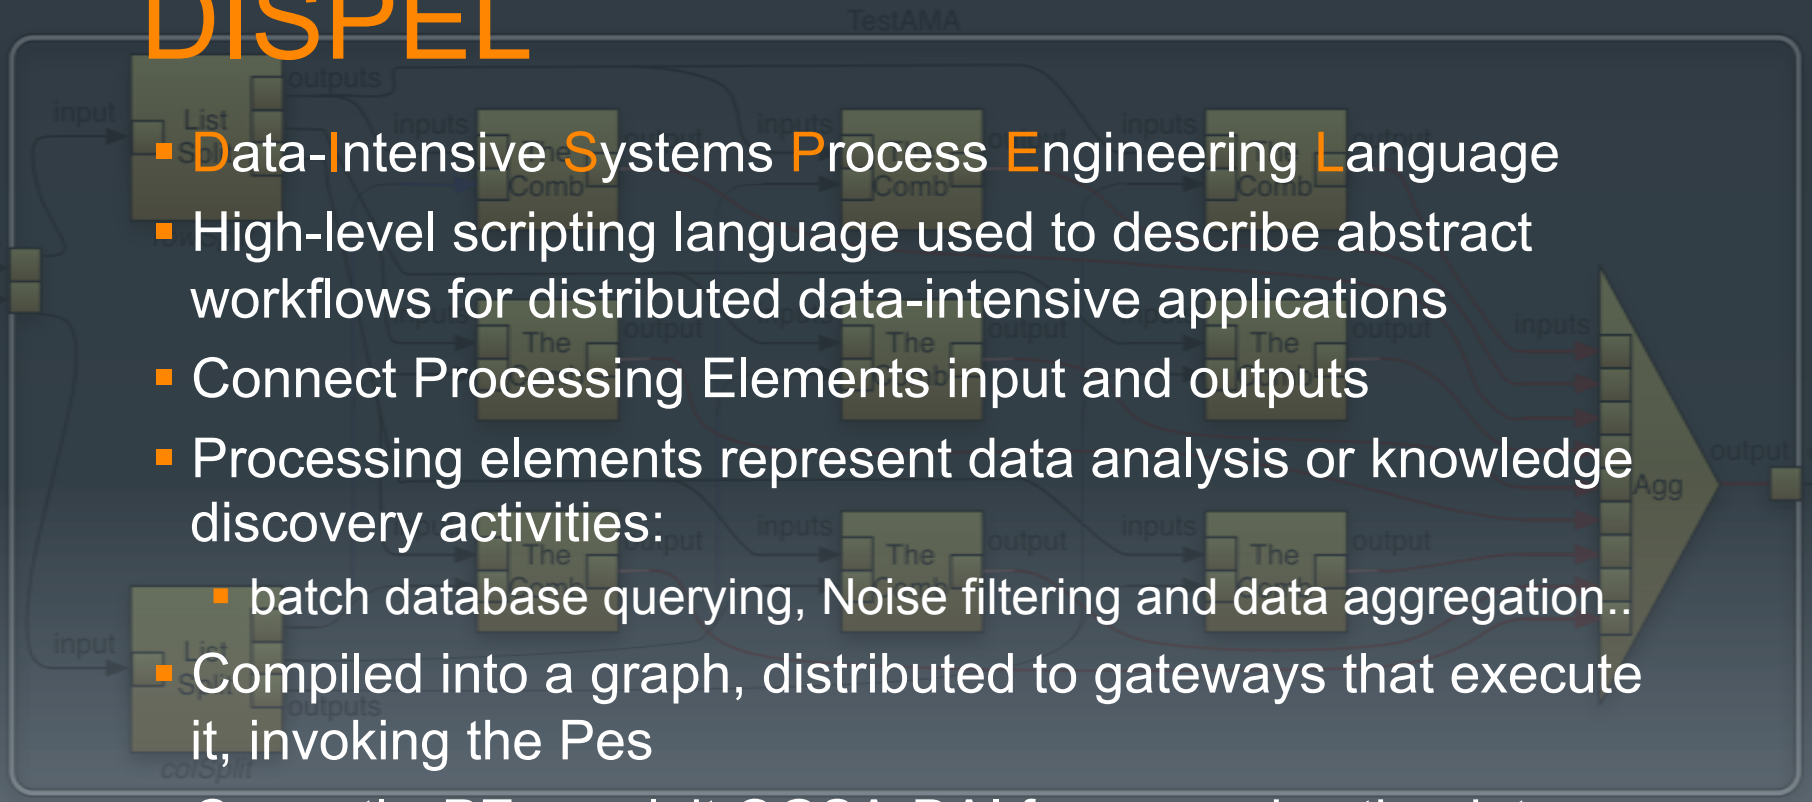

- Data-Intensive Systems Process Engineering Language
- High-level scripting language used to describe abstract workflows for distributed data-intensive applications
- Connect Processing Elements input and outputs
- Processing elements represent data analysis or knowledge discovery activities:
  - batch database querying, Noise filtering and data aggregation..
- Compiled into a graph, distributed to gateways that execute it, invoking the Pes
- Currently, PEs exploit OGSA-DAI for accessing the data

RAPID

- Rapid is an efficient way of designing and delivering portal interfaces for tasks and applications requiring access to remote compute resources.
- The expert specifies information about the task (transfer, options to display to the user, security, etc) in an XML file.
- Rapid translates the XML file into a job submission and a web-based portlet for user interaction.

HW architectures

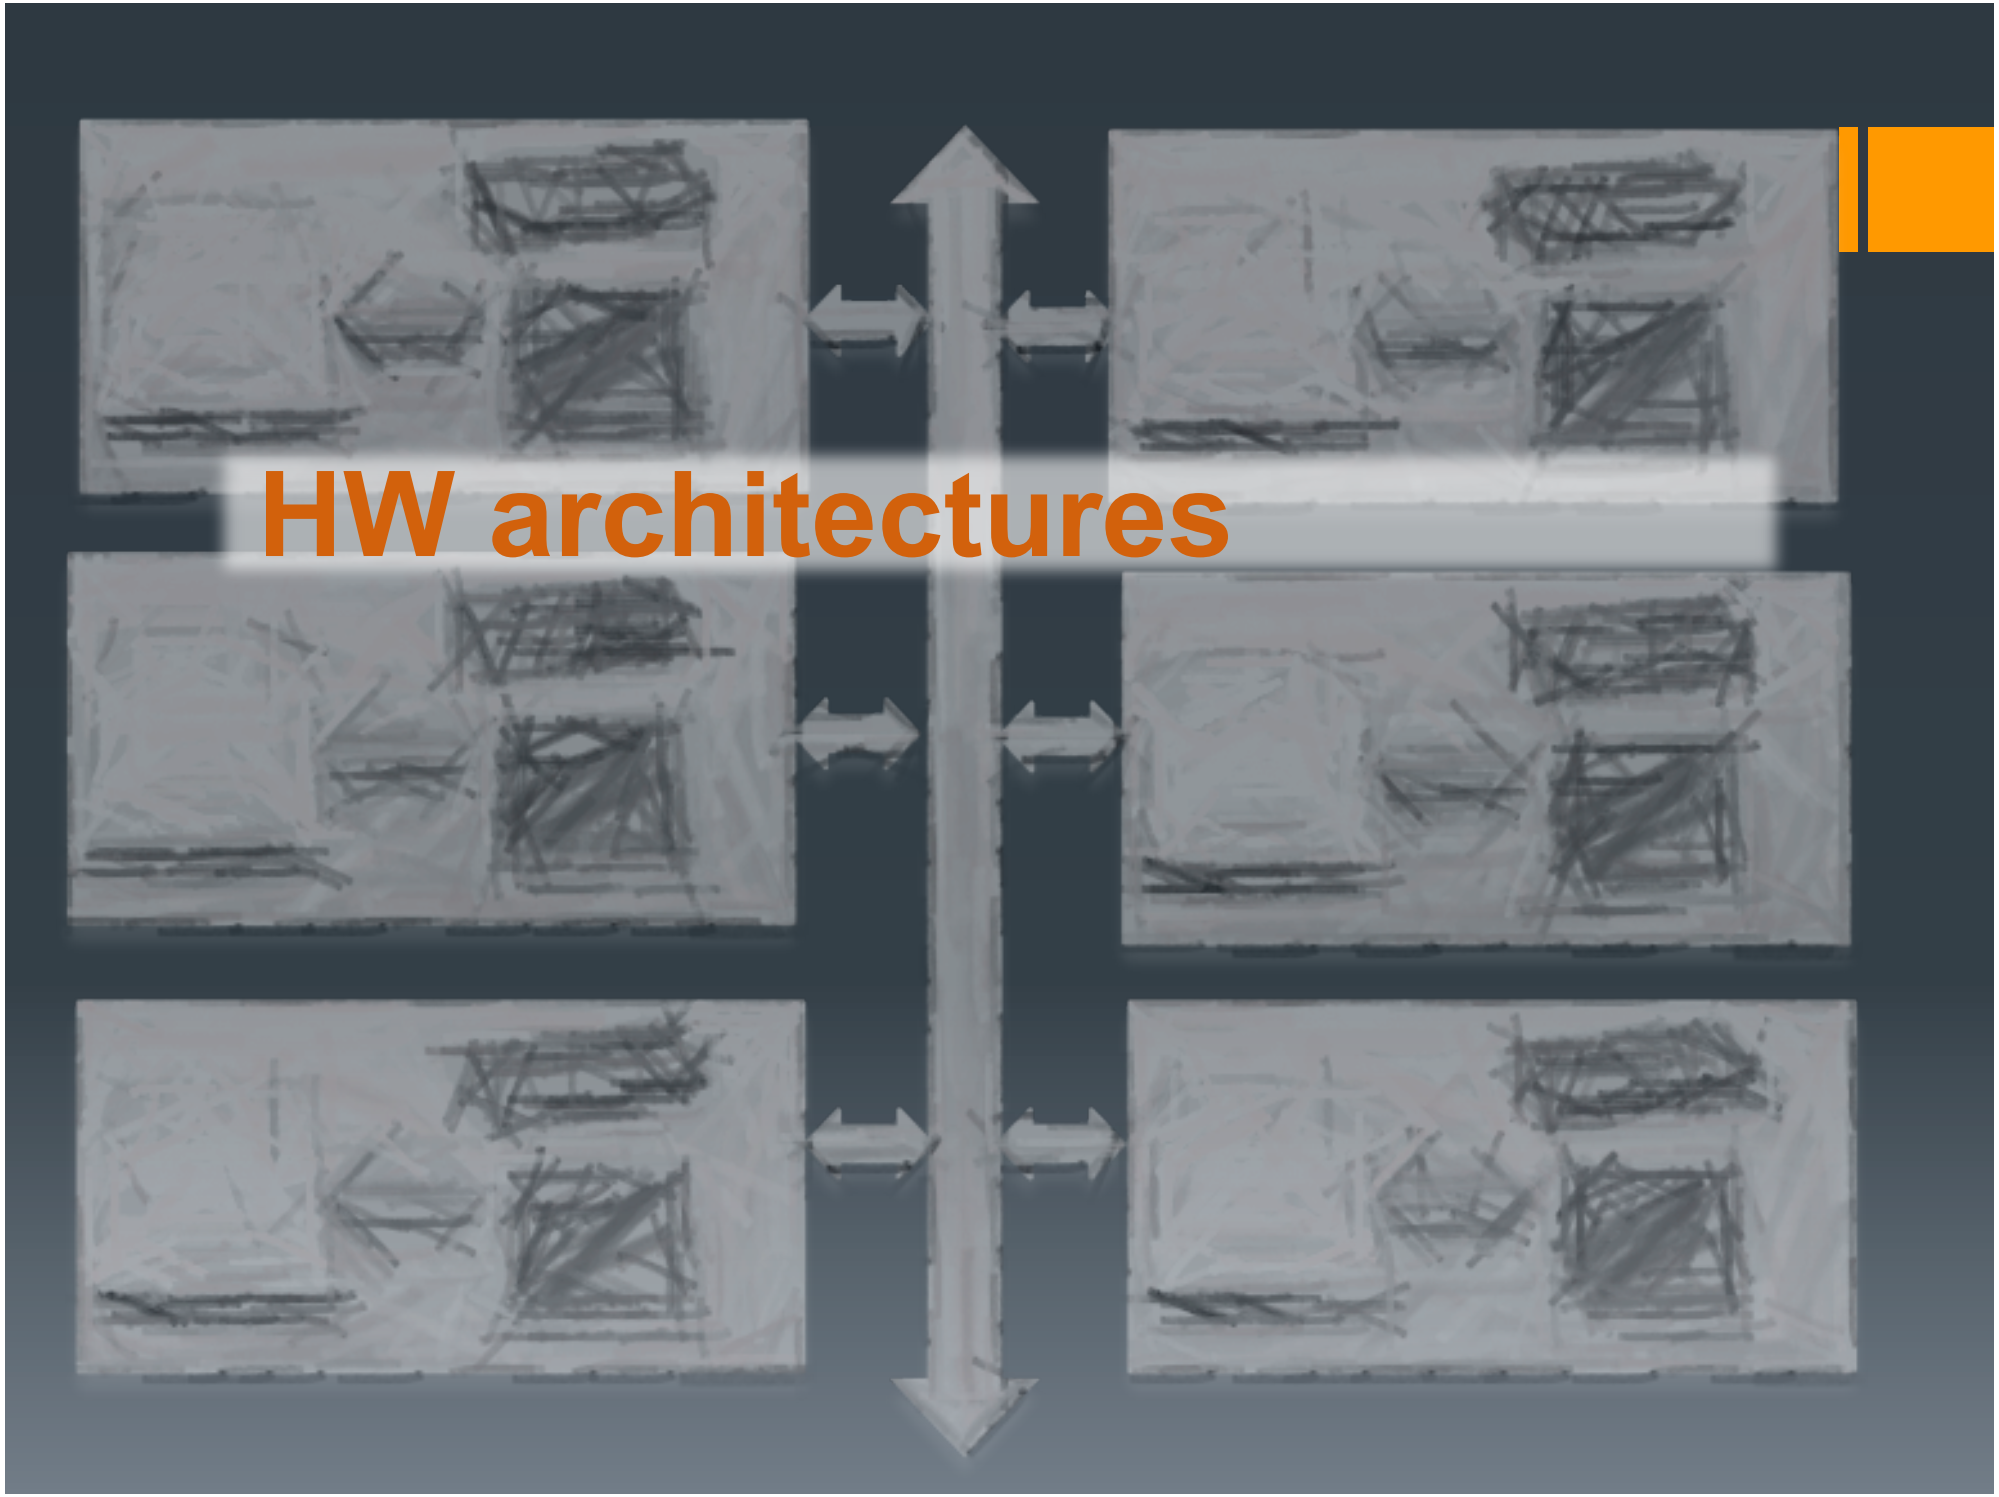

Classical approach

- Traditional model of processing separates processing from storage:
  - it may not be efficient (transfer can be a bottleneck)
  - It's often very expensive

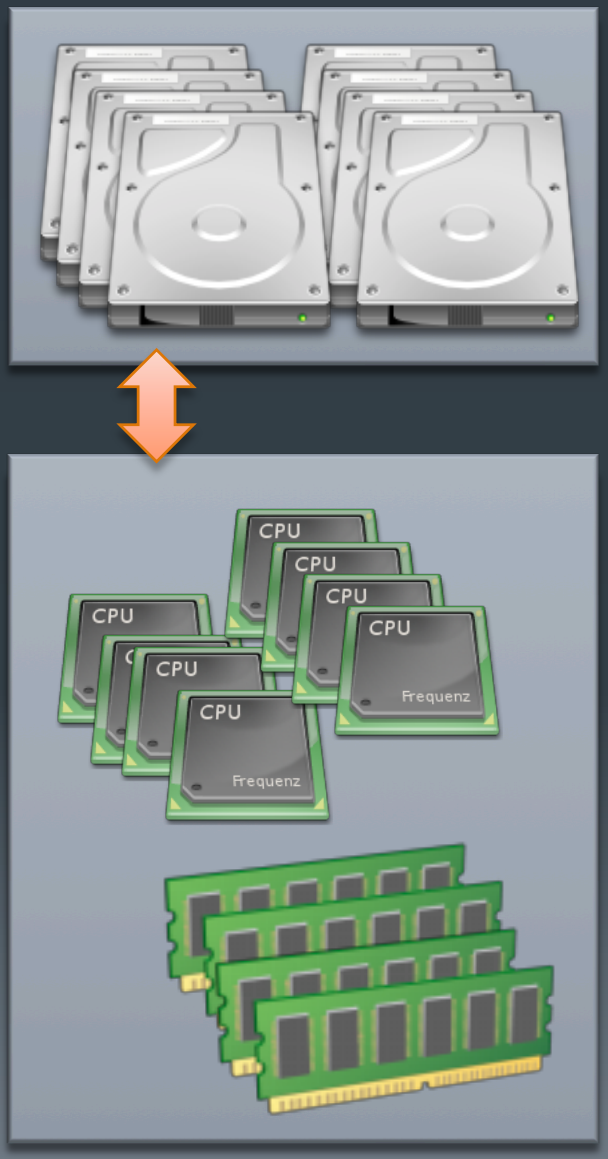

An alternative

- Ideas from Jim Gray (Microsoft) and Alex Szalay (JHU, The Sloan Digital Sky Survey)
- Create a network of "*data-bricks*":
  - low consumption node, with large storage capability
- Aims at processing data locally, reducing the need to transfer data to processing nodes

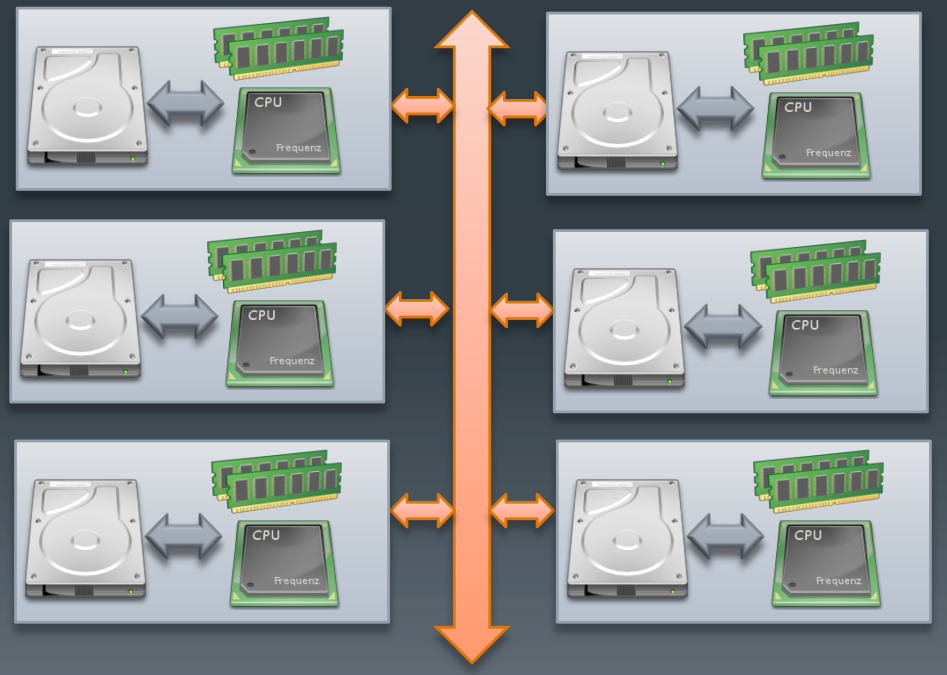

EDIM1

- Experimental hardware architecture used to:
  - calibrate the hardware architecture
  - identify efficient software architectures for data intensive applications
- Worked with scientists to identify data intensive applications as test beds

Physical layout

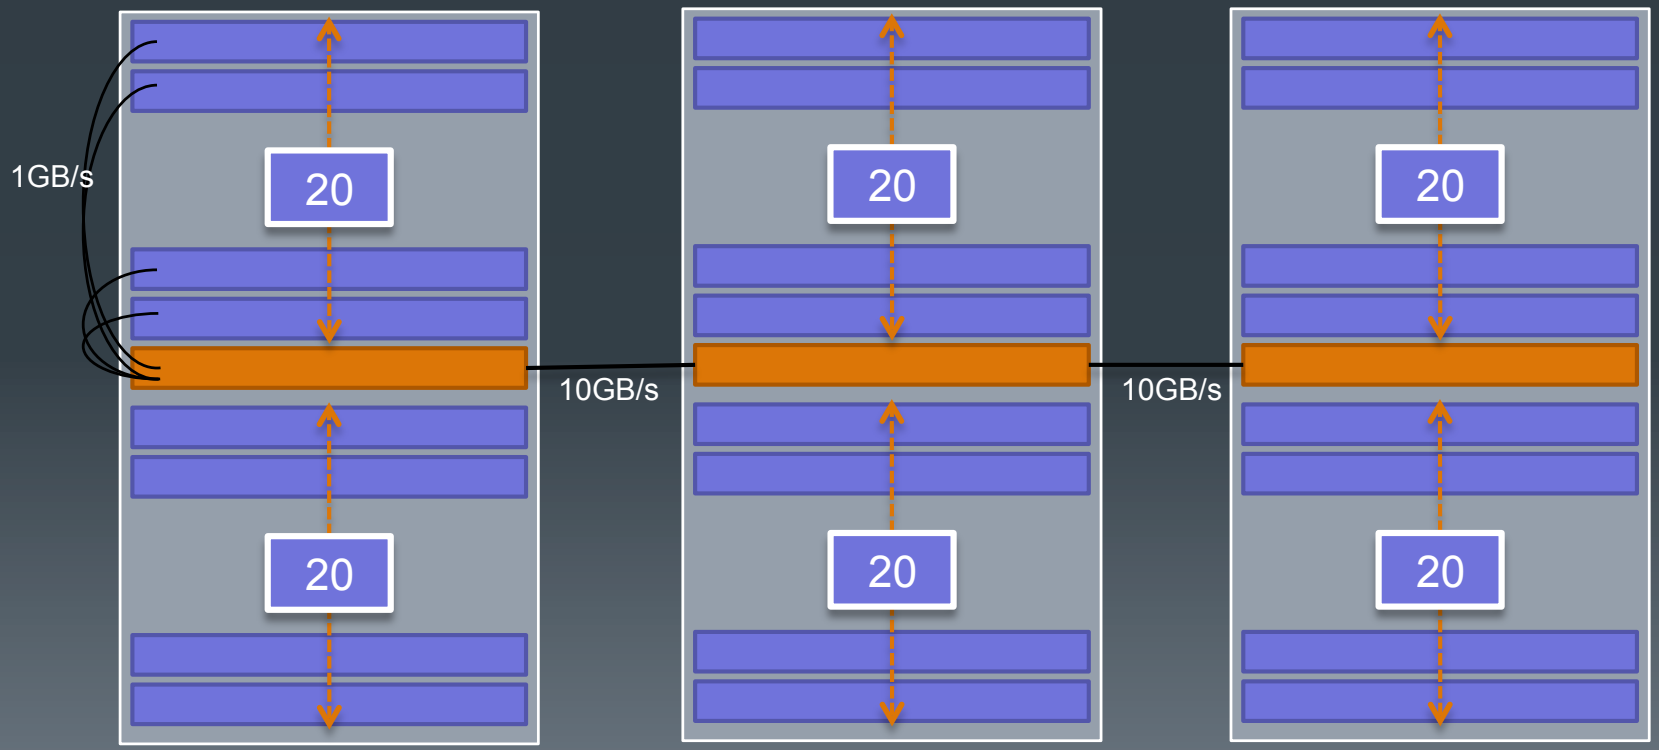

A node on EDIM 1

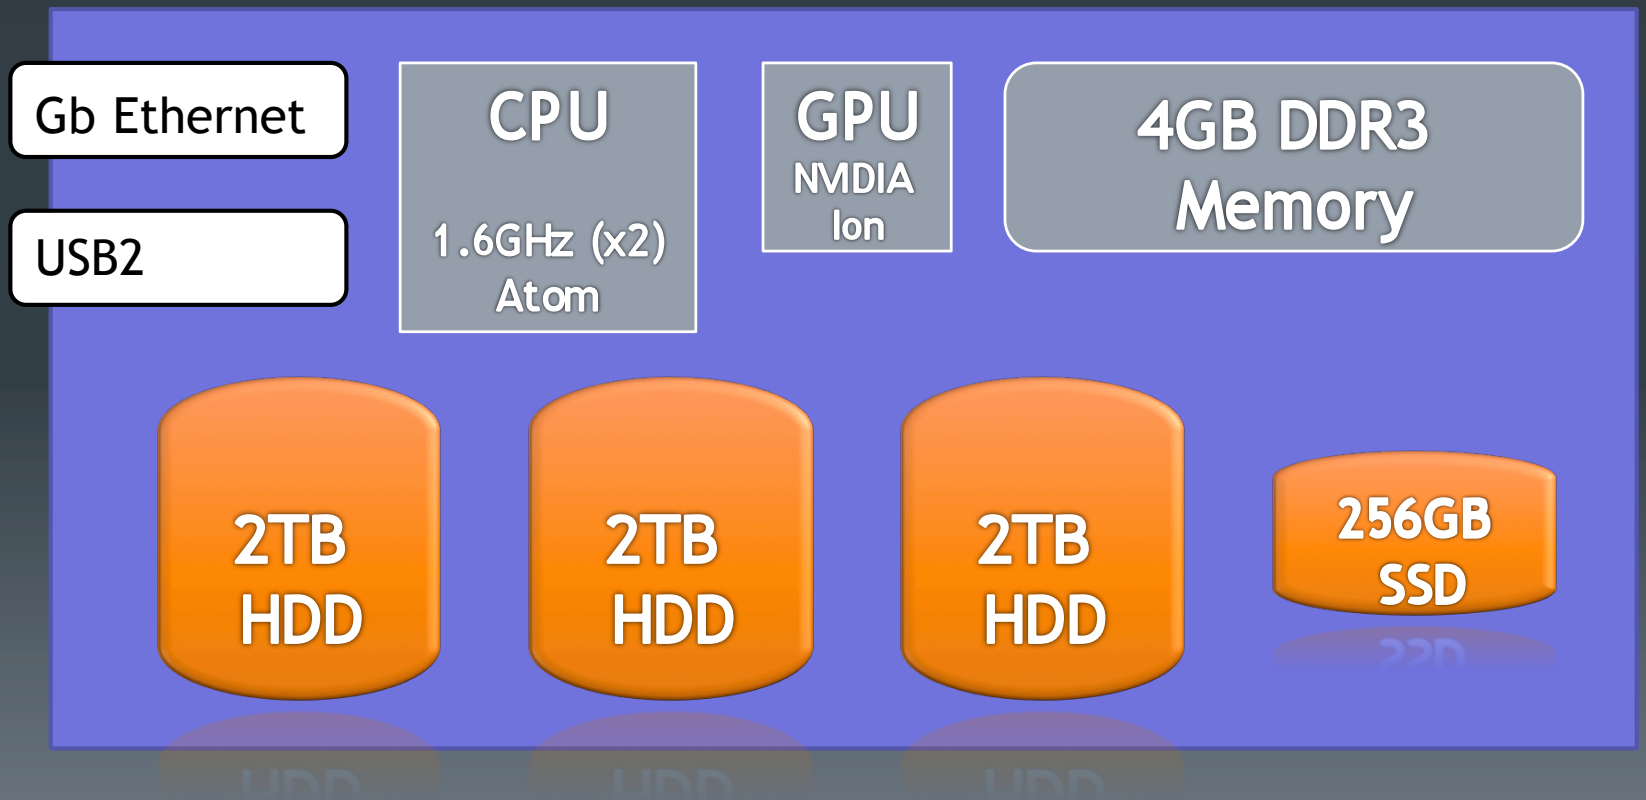

Access policy

- It is not a job submission system
- Nodes are reserved to experiments (by negotiation)
- Full control over the nodes: the software stack to install on the nodes must be specified
- We use ROCKS to manage
  - the creation of the software stack from the specifications
  - the installation
  - The naming of the nodes

Examples of applications

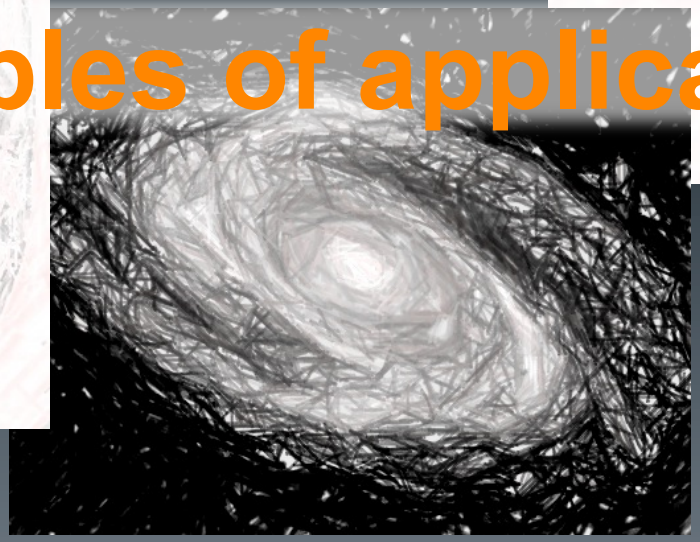

Mainly from ADMIRE project

Gene interaction in development

What genes interact during mouse embryo development?

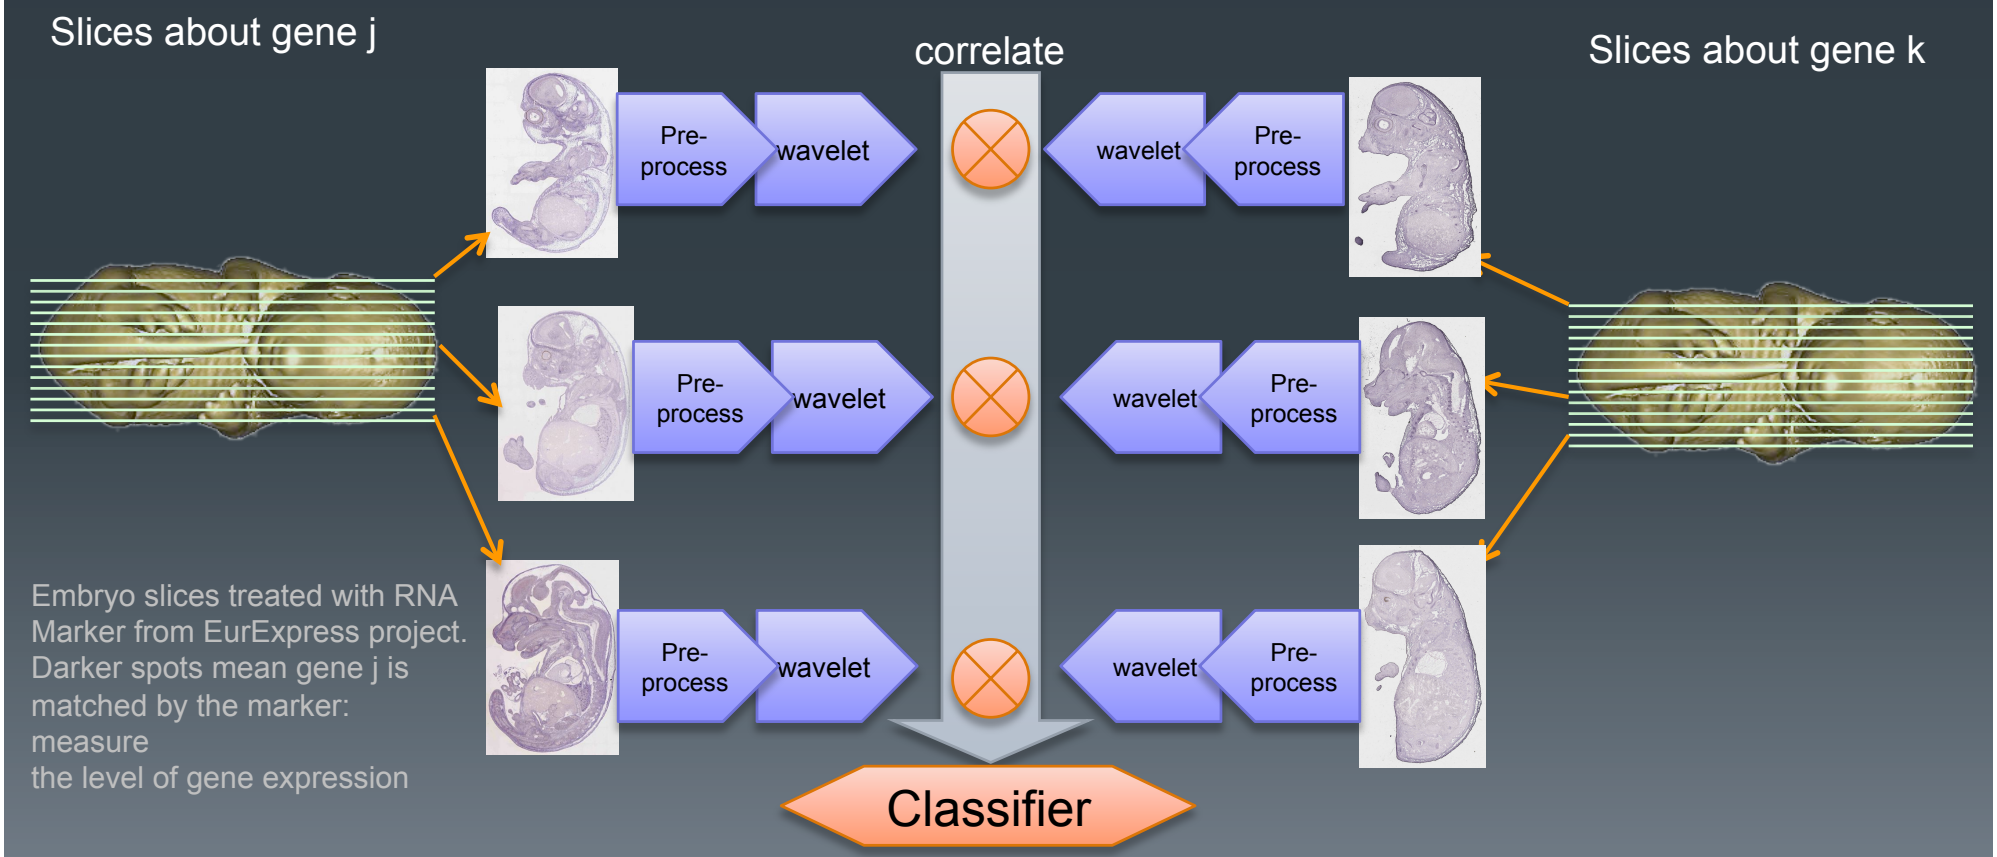

Embryo slices treated with RNA Marker from EurExpress project. Darker spots mean gene j is matched by the marker: measure the level of gene expression

Gravitational lensing

What is the distribution of dark matter in the Universe?

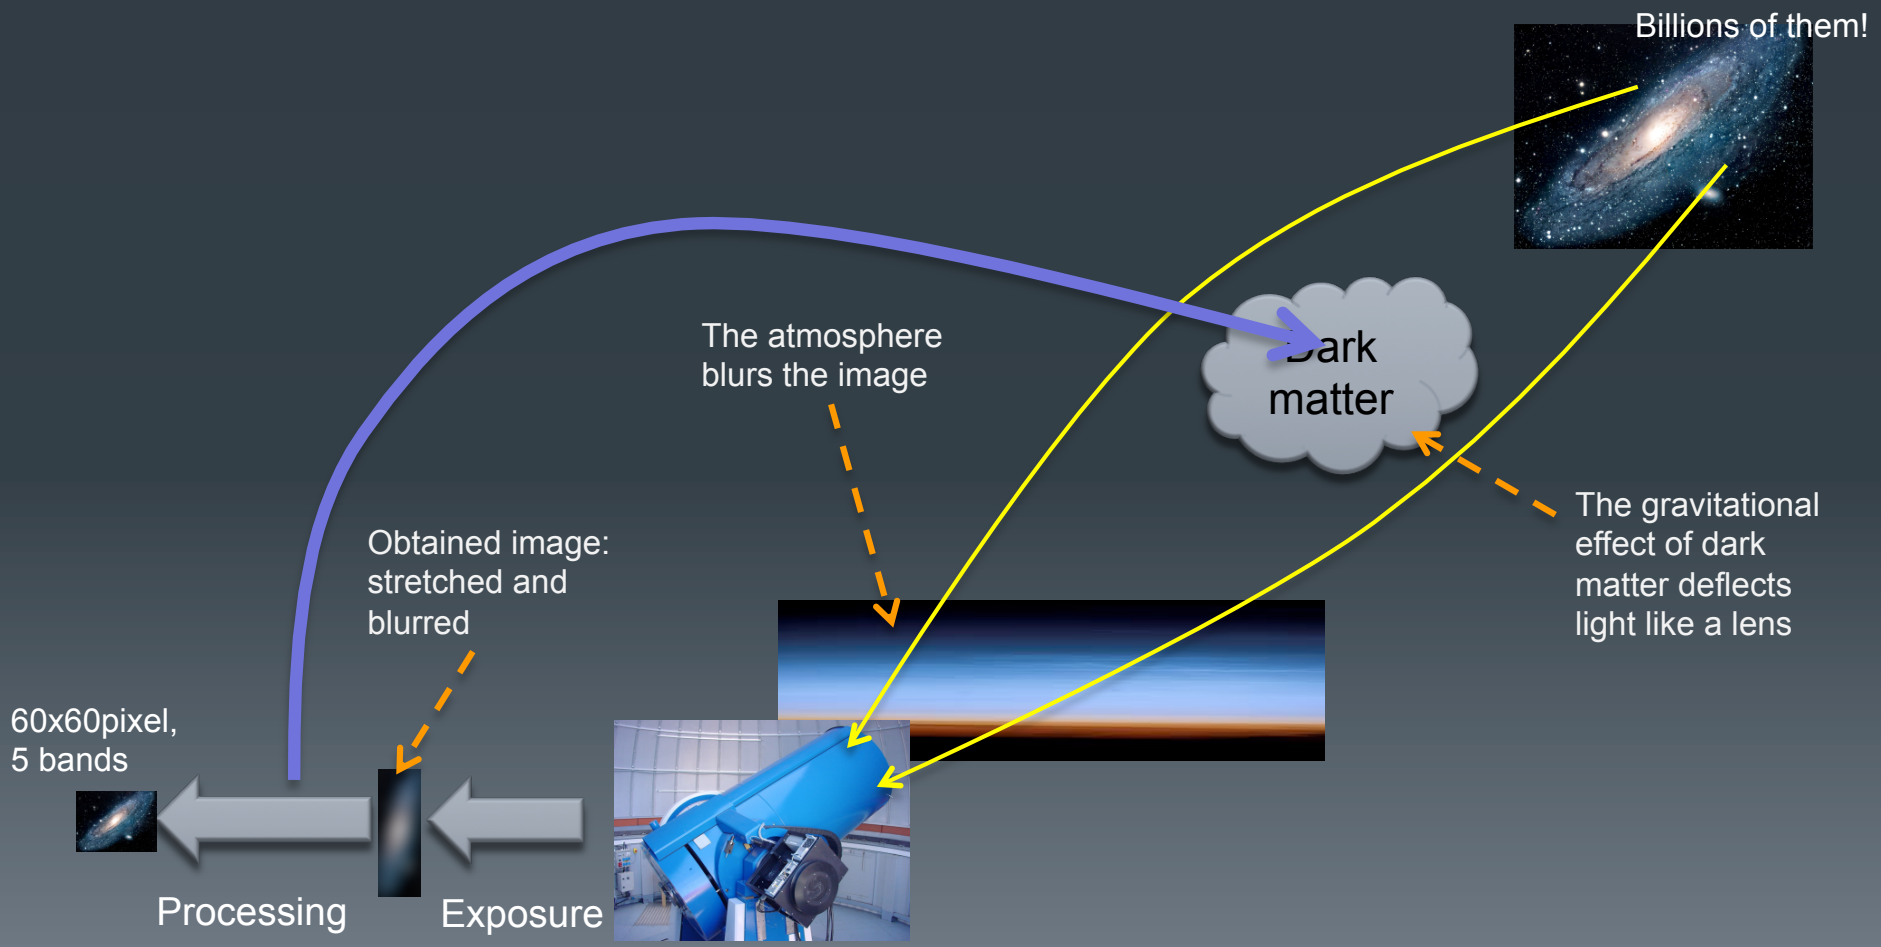

Cross correlation of waves

What is the structure of the earth?

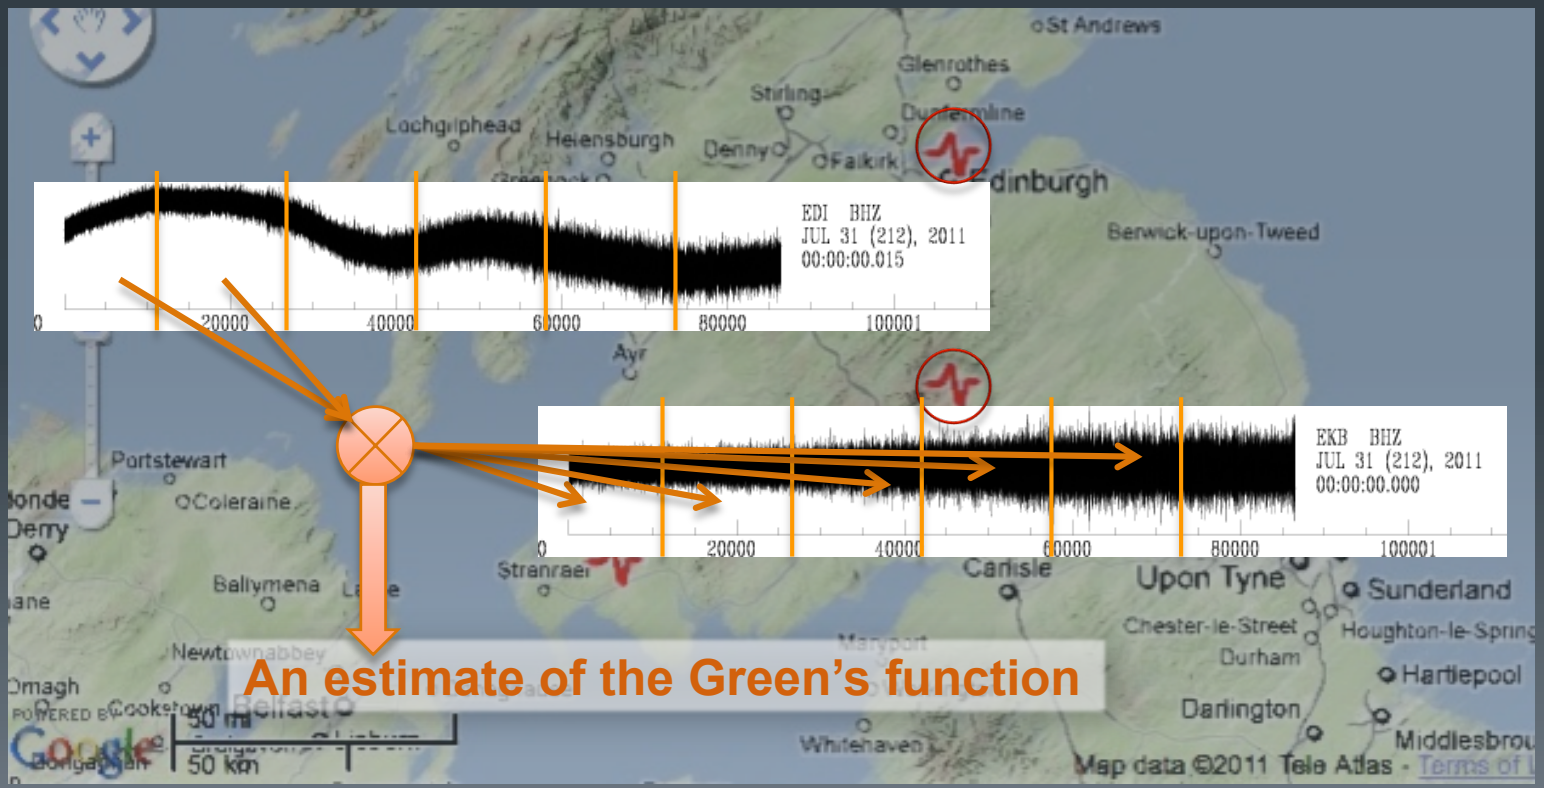

Conclusion

- Amount of data produced and available is increasing
- Bringing data to computation is not efficient
- New approaches are required
- DIR group has expertise and resources to tackle the problem
  - Software resources such as DISPEL and RAPID
  - Hardware resources such as EDIM1
  - Experience in large data-intensive applications. The ADMIRE project used as test beds:
    - Biology (gene interaction)
    - Cosmology (gravity Lensing)
    - Seismology (cross correlation of noise)

Thanks!
Any Question?