Commission européenne

Creating the HPC Infrastructure for the Human Brain Project

HPC 2014
Cetraro

Thomas Lippert
Director of HBP-Subproject HPC
Forschungszentrum Jülich

Vision

- Creation of a **virtual human brain model**
- Integration of all existing relevant **data**
- Simulation of the **entire** human brain model
- A **new** approach to **decoding** the brain
- Bridging Experiment and Theory in **Neuroscience**
- Supporting **therapy** of brain diseases
- Pushing **brain-inspired / Neuromorphic computing**
- **Guiding** HPC and Big Data to new horizons
- Establish an **HPC infrastructure** for Neuroscience
- Virtual brain simulator as **user facility**

Disclaimer

The most exciting questions:

- How is information stored in the brain?
- What is our communication protocol?
- How do groups of neurons and nerve fibers generate the higher brain functions?
- What is consciousness?

will not be answered immediately….

…but will be made accessible to research by a new method of investigation

Flagships Implementation Timeline

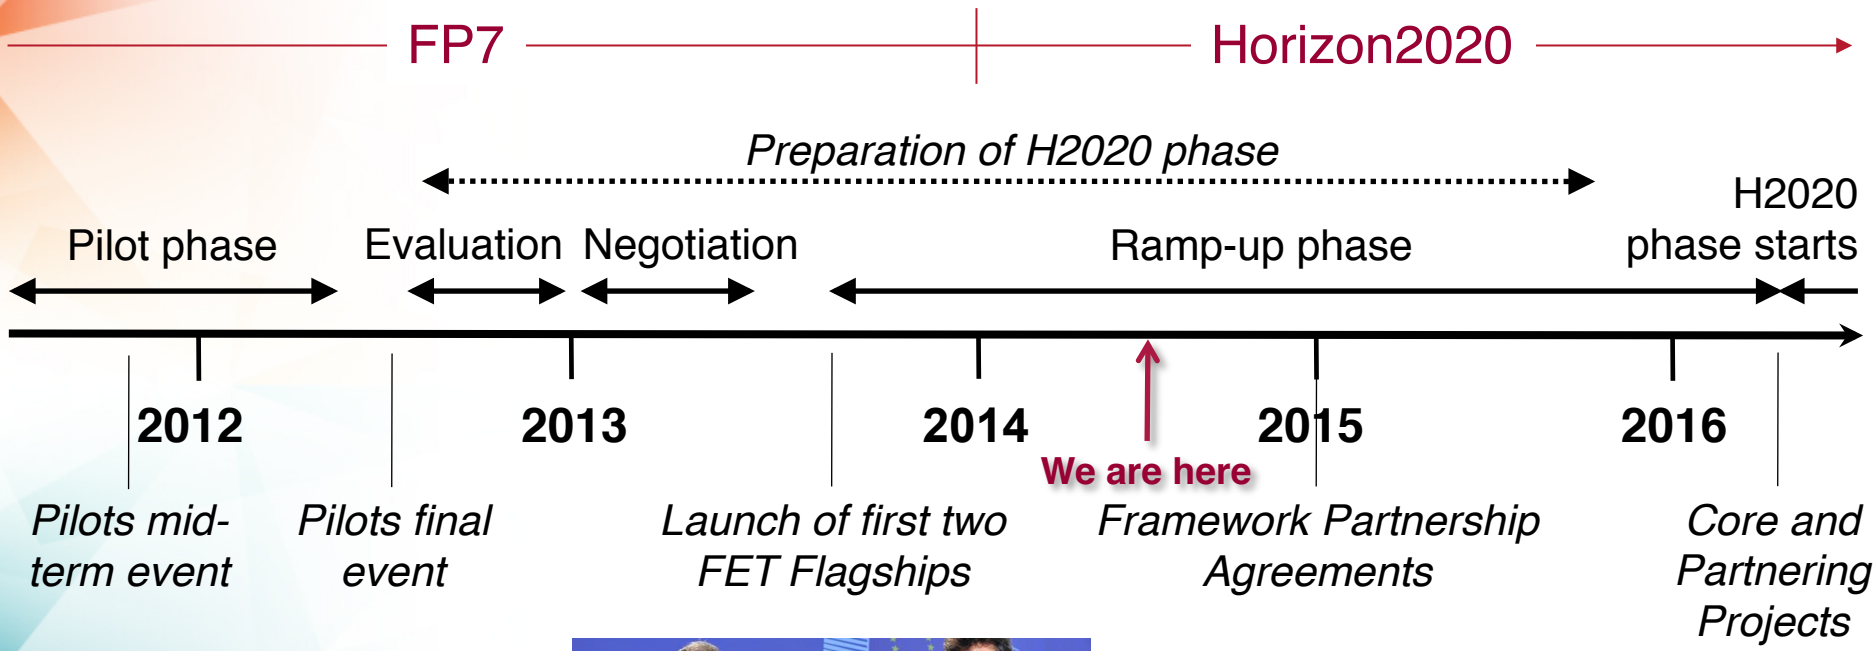

GRAPHENE FLAGSHIP

Human Brain Project

HBP Milestones

- **Jan 2013:** Announcement of HBP as one of two ICT-driven *FET Flagships* (10 years)
- **Oct 2013:** Launch of 2.5-year HBP *Ramp-Up Phase* with 80 partners from 23 countries
- **April 2014:** 32 new institutions from 13 countries join the consortium after *Competitive Call*
- **June 2014:** Proposal for *Framework Partnership Agreement (FPA)* submitted, signature end of 2014
- **April 2016 – Sept 2023:** Series of *Core Projects,* funded through *Special Grant Agreements (SGA)*
- Complementary *Partnering Projects*, funded by EC and regional, national and transnational sources

Objectives

- To build an **integrated ICT infrastructure** enabling a
- **global collaborative effort** towards understanding the human brain, and ultimately
- to emulate its **computational capabilities**

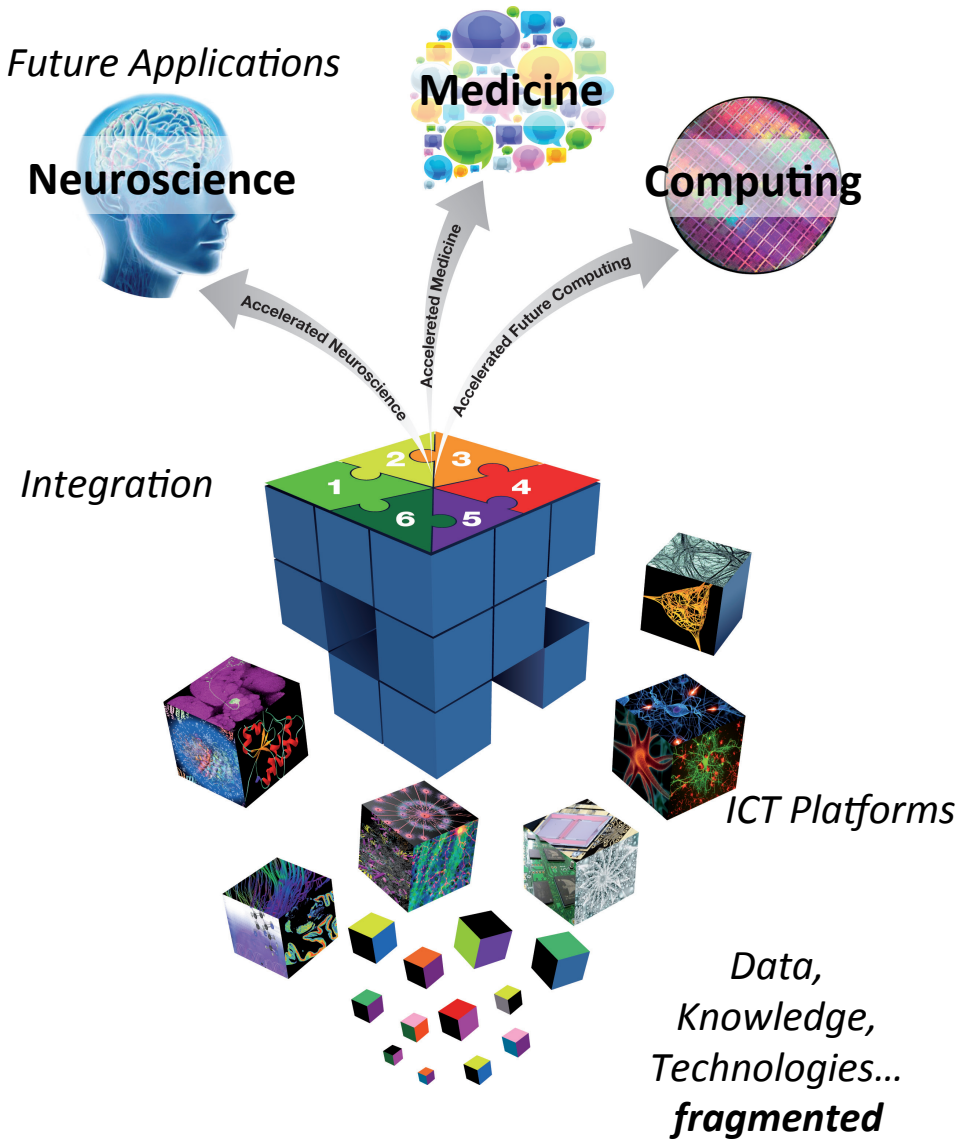

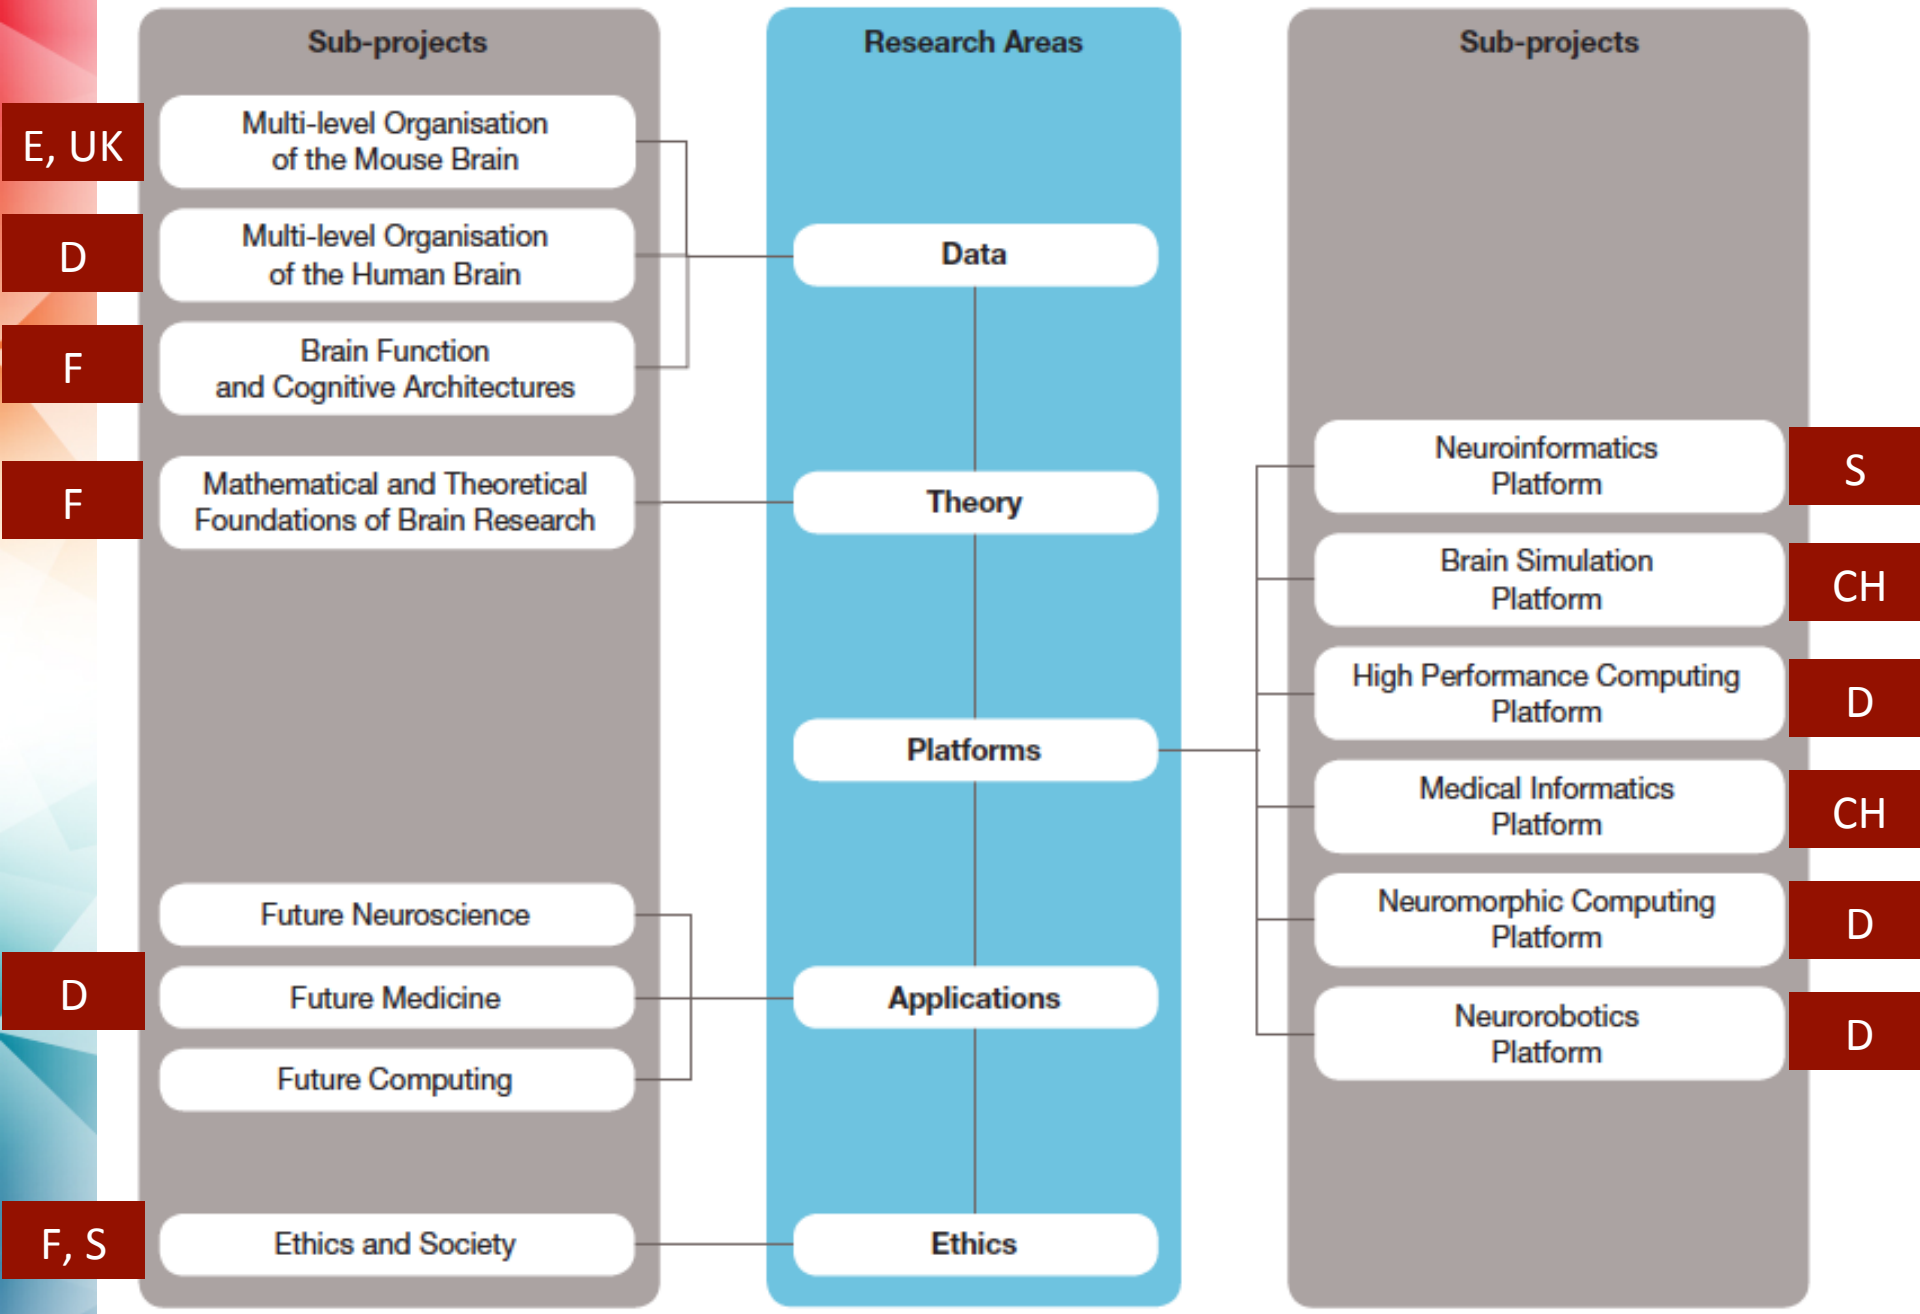
Sub-projects
Research Areas
Sub-projects
E, UK
Multi-level Organisation of the Mouse Brain
D
Multi-level Organisation of the Human Brain
F
Brain Function and Cognitive Architectures
F
Mathematical and Theoretical Foundations of Brain Research
Data
Theory
Platforms
Applications
Ethics
Neuroinformatics Platform
S
Brain Simulation Platform
CH
High Performance Computing Platform
D
Medical Informatics Platform
CH
Neuromorphic Computing Platform
D
Neurorobotics Platform
D
Future Neuroscience
D
Future Medicine
Future Computing
F, S
Ethics and Society

ICT Platforms of the HBP

- Neuroinformatics (Sten Grillner - Sweden)
  - Tools to manage, navigate and annotate brain atlases
- Brain simulation (Henry Markram - Switzerland)
  - Simulate unifying brain models integrating all available data
- Medical informatics (Richard Frackowiak - Switzerland)
  - Data mining on a large volume of clinical data
- Neuromorphic computing (Karlheinz Meier – Germany)
  - Develop and provide access to neuromorphic devices
- Neurorobotics (Alois Knoll - Germany)
  - Interface a detailed brain model to a simulated body
- High Performance Computing (Thomas Lippert – Germany)
  - Exascale Capability / Big Data / Future Computing (Hybrid)

Contribute to the Future of Computing

What they could provide – a few selected items

High Performance Computing

- Interactive, visual. Exascale supercomputing
- Massive distributed volumes of heterogeneous data
- Convergence with neuromorphic technology

Neuromorphic Computing

- First generic, large-scale neuromorphic systems
- Beyond Turing - no algorithmic operation
- Beyond von Neumann - immerse computation in memory
- Technology integration (3D, non-CMOS backends)

Neurorobotics

- Virtual robots with two-way, closed loop interfaces
- Link to brain models and neuromorphic systems
- Physical prototypes and applications

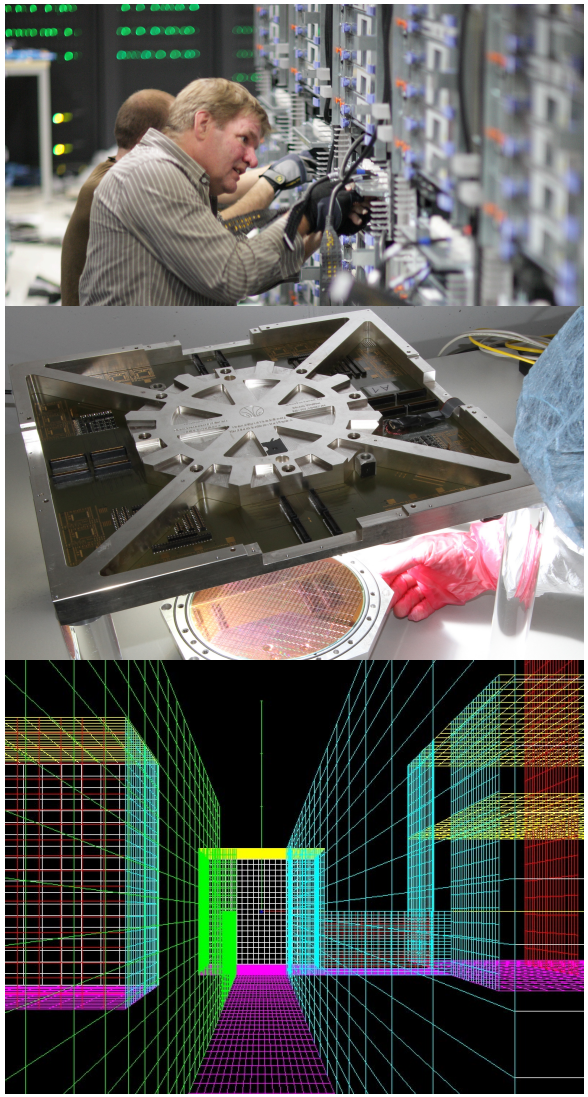

The Multi-Scale Simulation Challenge

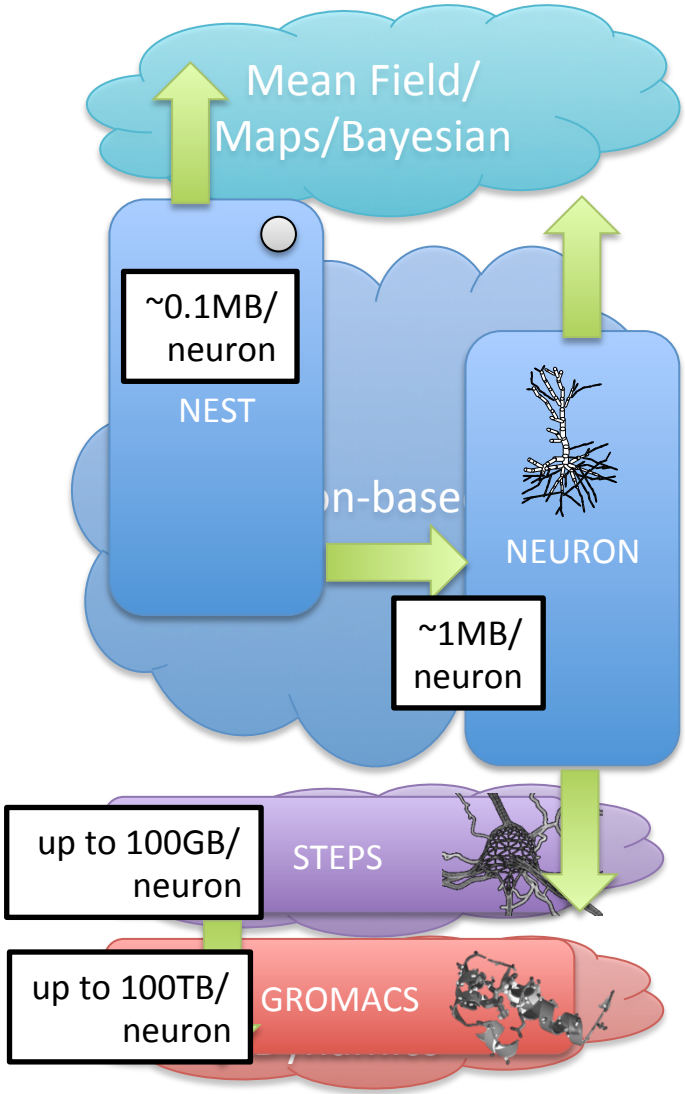

NEST scaling

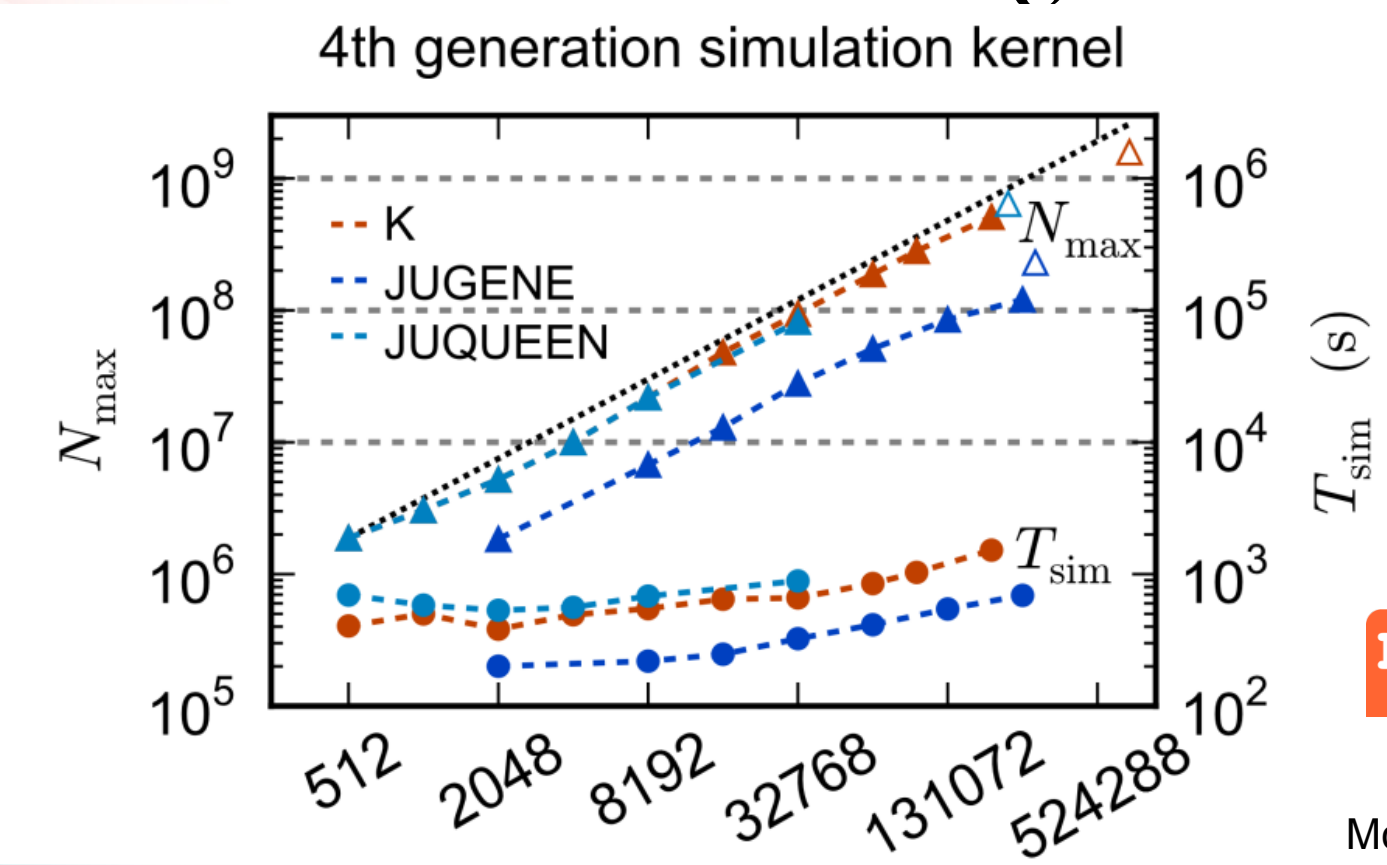

nest:: initiative

Diesmann, Morrison et al.

- aim: full-scale models at cellular and synaptic (contacts) resolution
- maximum-filling benchmarks
- 1% of HB on JUQUEEN (10^9 neurons, 10^{13} synapses

The HBP Data Challenge

... The tremendous complexity of the structure of the grey matter is so intricate, that it withstands the tenacious curiosity of researchers und will continue to withstand for many hundred years to come.

Santiago Ramon Cajal

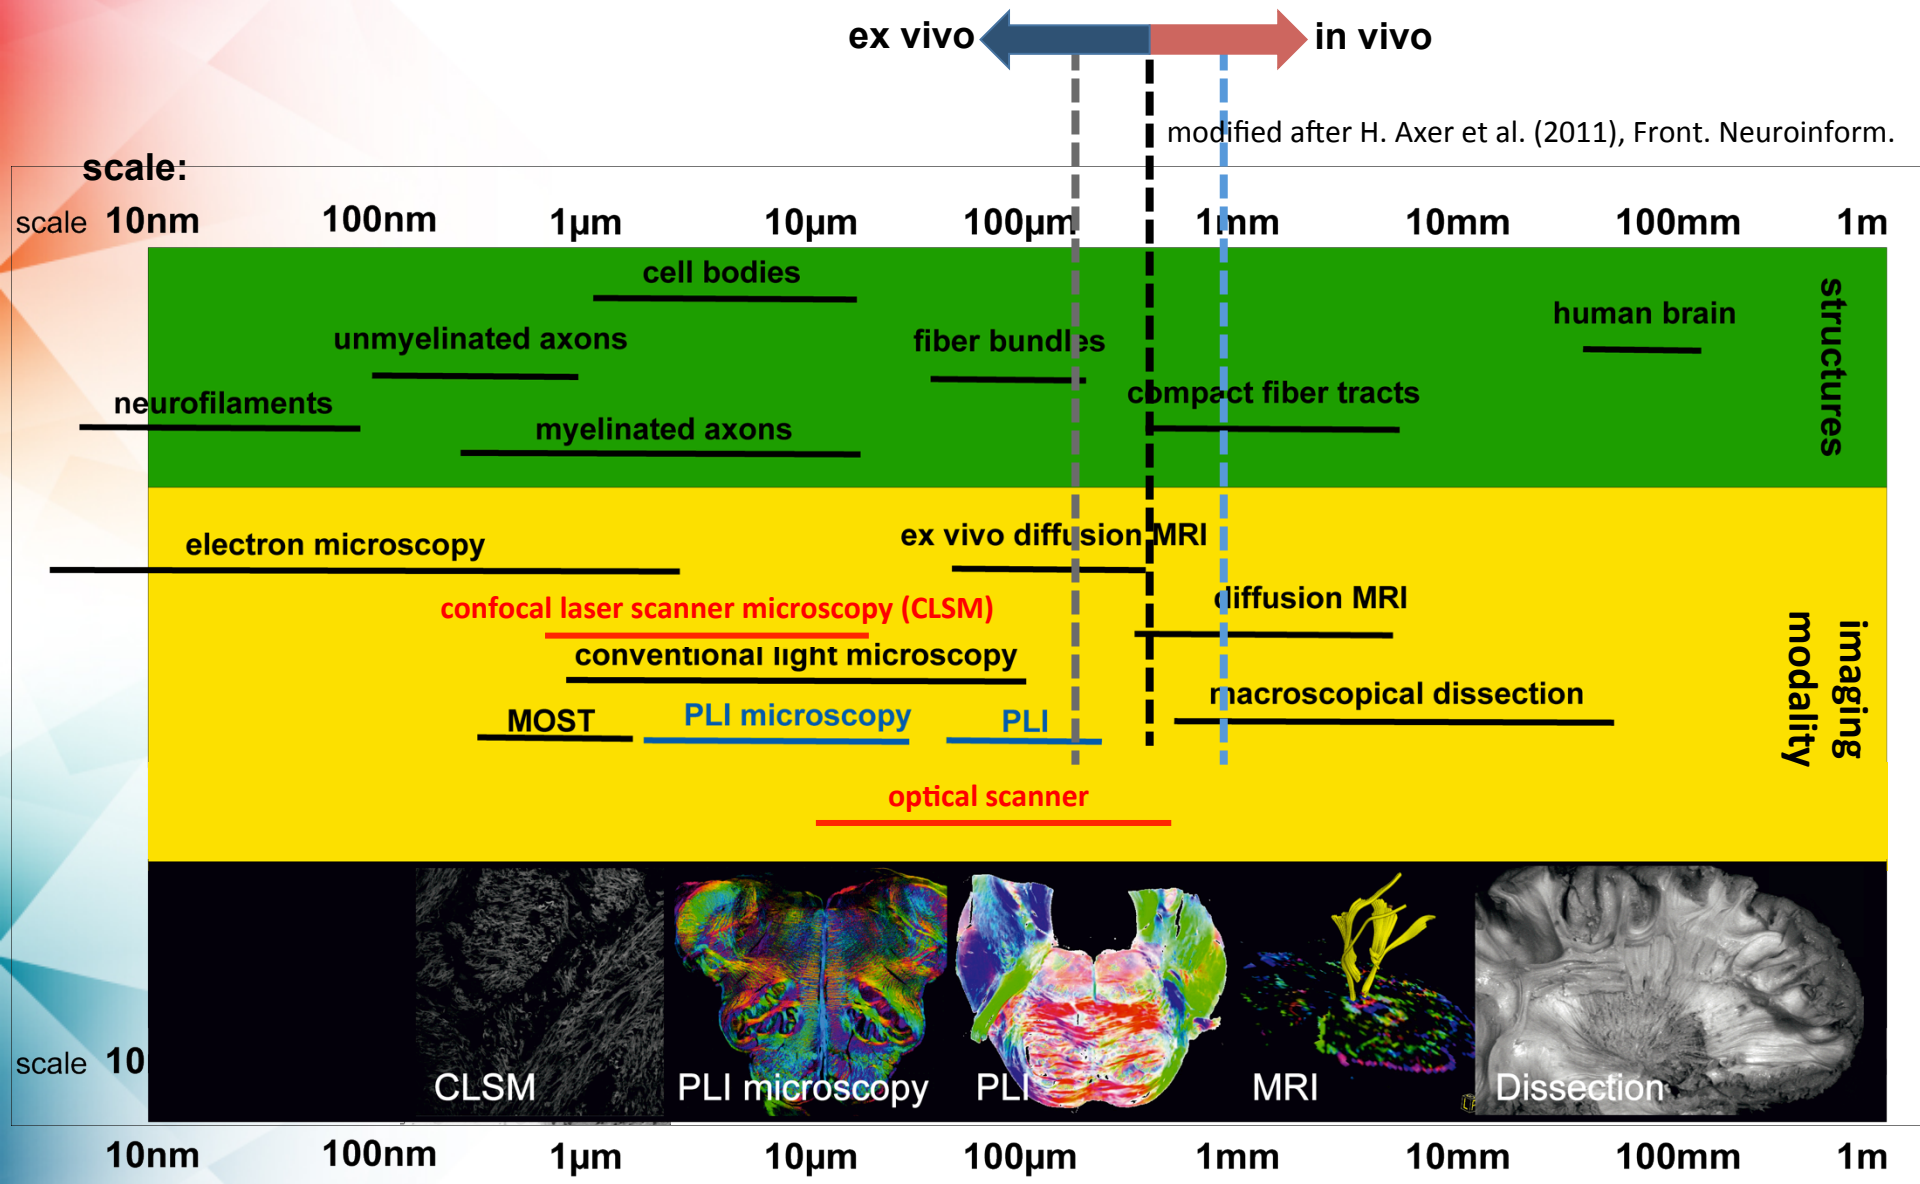
ex vivo
in vivo
modified after H. Axer et al. (2011), Front. Neuroinform.
scale:
scale 10nm 100nm 1µm 10µm 100µm 1mm 10mm 100mm 1m
cell bodies
unmyelinated axons
fiber bundles
human brain
neurofilaments
myelinated axons
compact fiber tracts
structures
electron microscopy
ex vivo diffusion MRI
confocal laser scanner microscopy (CLSM)
diffusion MRI
conventional light microscopy
macroscopical dissection
MOST
PLI microscopy
PLI
optical scanner
imaging modality
scale 10
CLSM
PLI microscopy
PLI
MRI
Dissection
10nm 100nm 1µm 10µm 100µm 1mm 10mm 100mm 1m

Cellular architecture at different spatial resolutions

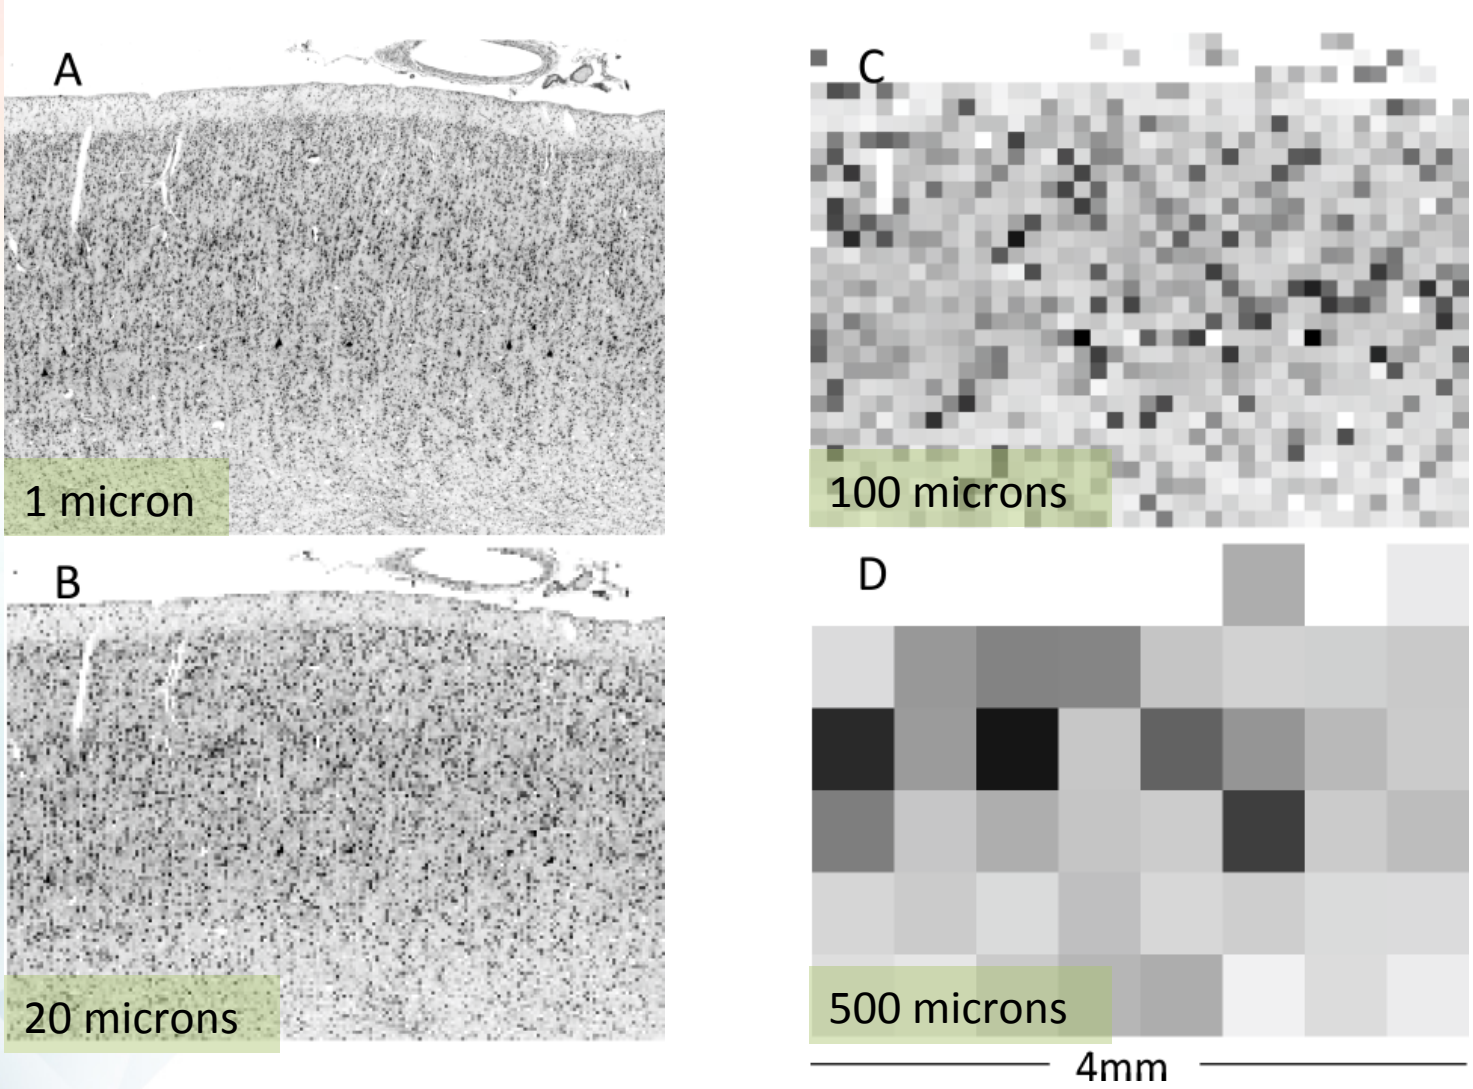

Storage Requirements

- Cellular strucuture:
  - 300 TB / brain
  - 3 PB / sample
- Polarizing light imaging microscope:
  - 4 PB / brain
  - 40 PB / sample (10 brains)
- Simulation:
  - NEST: 10 PB / brain (in memory)
  - NEURON: 100 PB / brain (in memory)

HPC Platform Structure

- Technology Evaluation (D. Pleiter)
- Mathematical Methods, Programming Models and Tools (J. Labarta)
- Interactive Visualization, Analysis and Control (T. Kuhlen)
- Exascale Data Management (A. Ailamaki)
- Integration and Operations (Th. Schulthess)
- User Support and Community Building, Scientific Coordination (Th. Lippert /B. Orth)

Vision for operating HBP HPC system

- Heterogeneous workflows:
  - Users running multiple jobs concurrently within a single session
  - Dynamic change of job composition
- Job examples
  - Large-scale simulation jobs (spiking neuronal network simulators)
  - Data analysis jobs
  - Visualization pipelines
- Job features
  - Multiple MPI processes
  - Multi-threaded processes
- Dynamic change of session resources

Pre-Commercial Procurement (PCP)

- R&D required to realize this vision
  - Available roadmaps do not indicate appropriate solutions being developed
- A PCP is designed to procure R&D services
  - Exemption from public procurement directives
  - Respects principles of public procurement: fair, open and transparent
  - Competitive process:
    - Open tender for bids to framework contract
    - Number of contractors reduced IN each phase of PCP:
      Phase I (Solution design): 5 competitors
      Phase II (Prototype development): 3 competitors
      Phase III (Pilot system development): 2 competitors

HBP PCP

– Expected outcomes

- R&D of HPC system components that allow for data-intensive interactive simulations, analysis and visualization, integrated into a future HPC architecture capable of providing a peak performance up to 50 PFlop/s
- Smaller-scale pilot systems, demonstrating the readiness of the developed technologies

– Main technical goals

- Significant enlargement of memory capacity based on power-efficient, dense memory technologies
- Scalable visualisation capabilities tightly integrated into a scalable HPC architecture with minimum data movement
- Dynamic resource management as a basic requirement to realize supercomputers featuring interactivity

Tools for the HBP

- Goals
  - Prepare tool sets for extreme-scale systems and HBP interactive workflows
  - Support HBP code teams in analyzing and tuning their simulation codes (NEST, Neuron, …)
- Partners
  - Jülich Supercomputing Centre (JSC)
  - Barcelona Supercomputing Center (BSC)
- Tool sets
  - Score-P, Scalasca, CUBE (JSC)
  - Extrae, Paraver, Dimemas (BSC)

Tool Scalability Verified

- Scalasca 1.4.3
- BlueGene/Q
- NASA NAS benchmark BT-MZ
- 1,048,704 cores
- Largest successful trace measurement and analysis ever done!

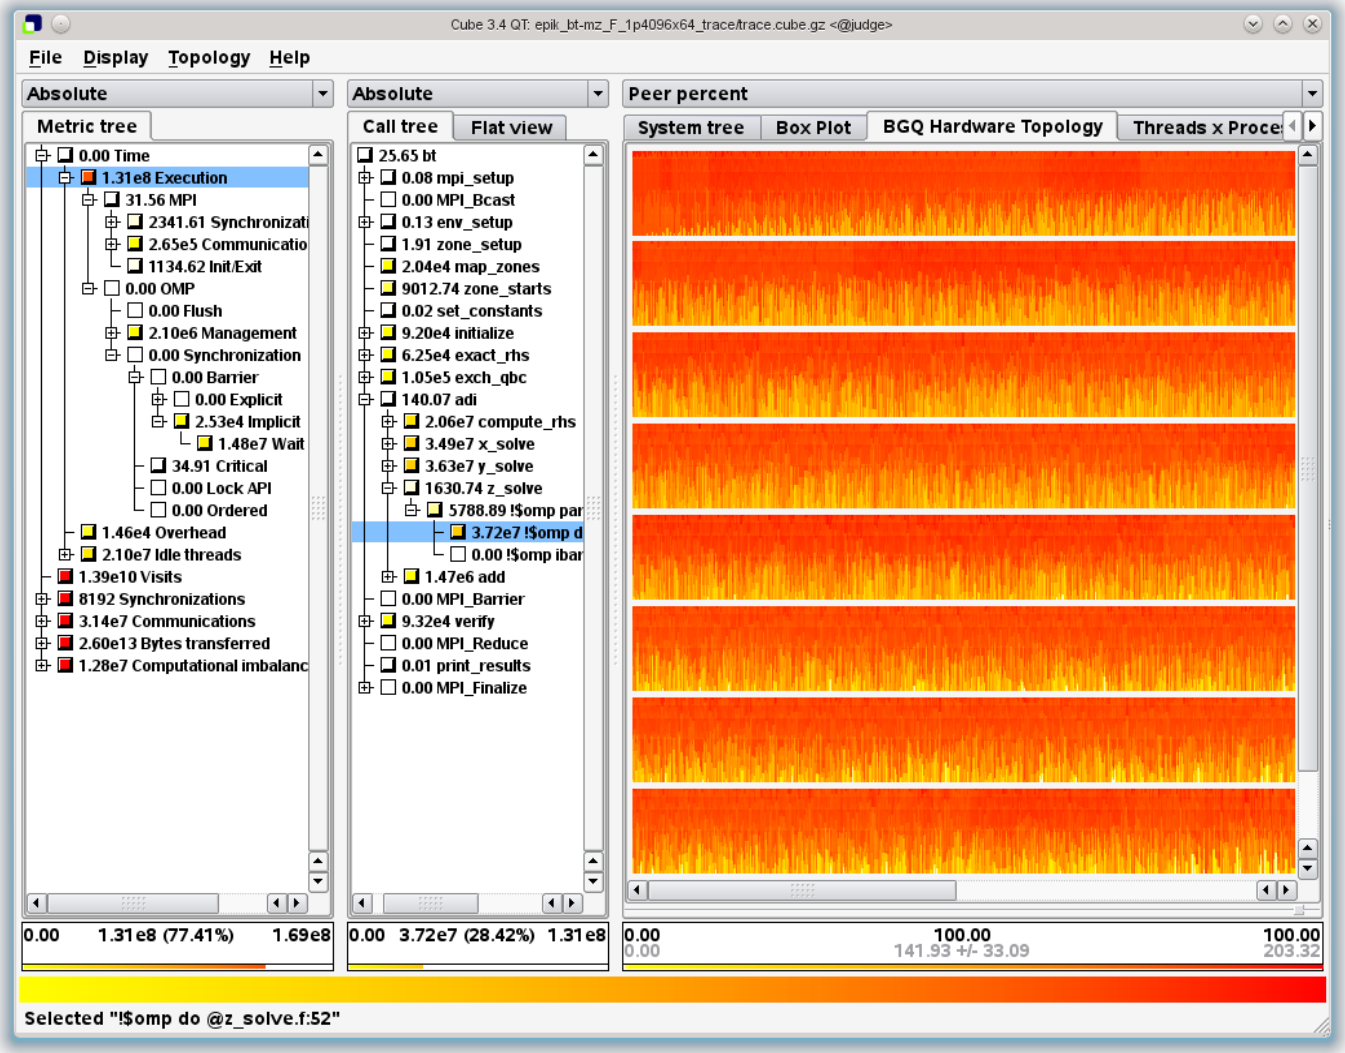

NEST Analysis

- Scaling analysis of NEST on BlueGene/Q
  - Measure on varying core numbers (e.g. 1, 2, 4, 8 racks)
  - Automatically determine and classify scaling behavior of each program region

- First result
  - Significant improvement of neural network buildup phase

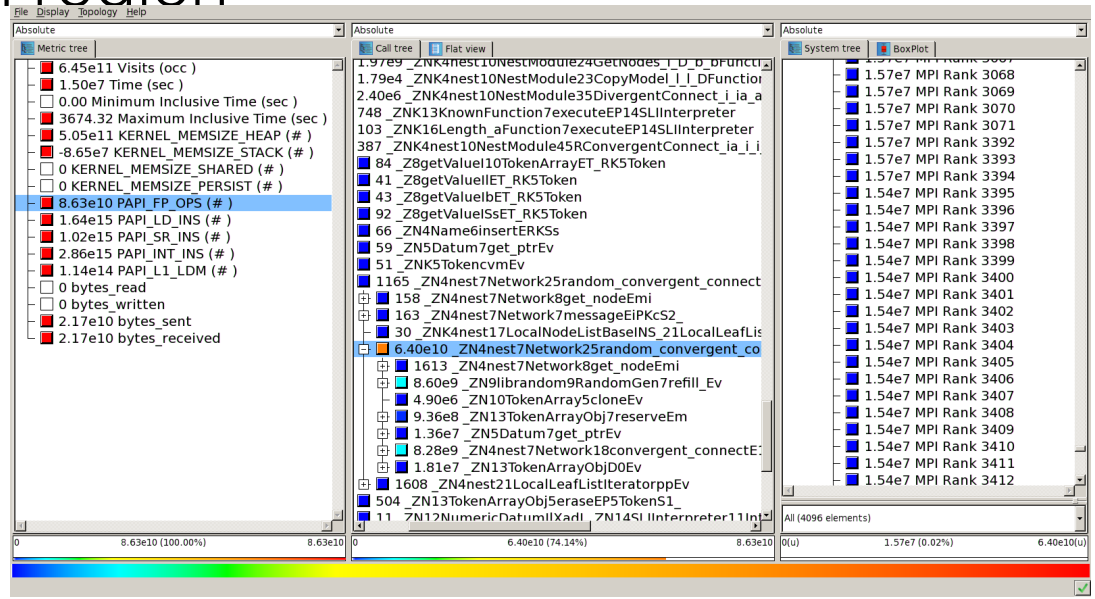

(*) www.vi-hps.org/projects/catwalk/

HPC Platform Architecture

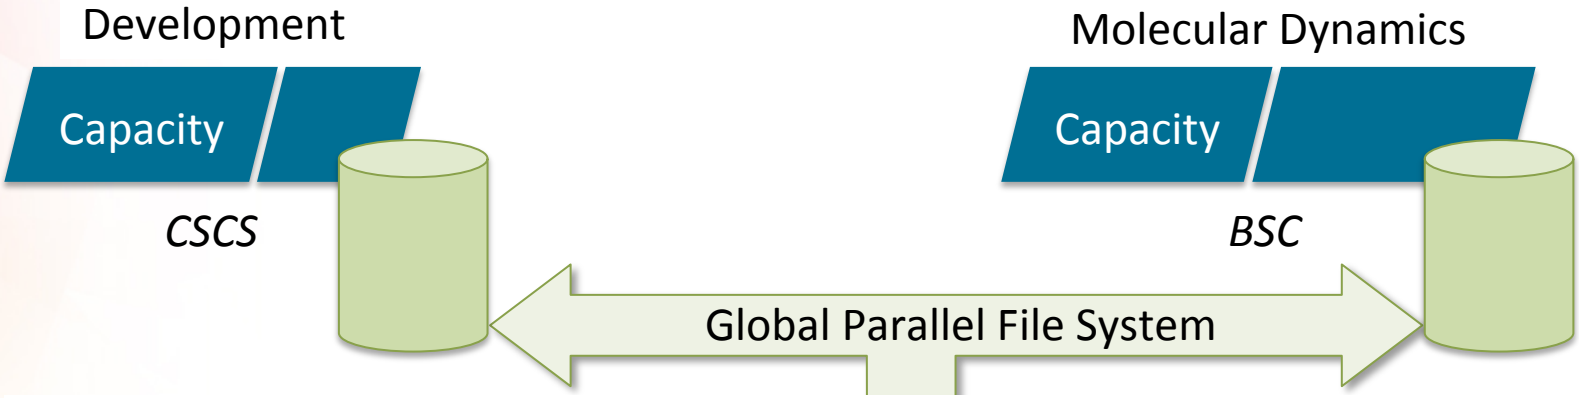

Partner project programme in H2020

Access to PRACE (based on peer review)

- CEA (Theory Simulations)
- LRZ (Robotics Simulation)
- others to join...

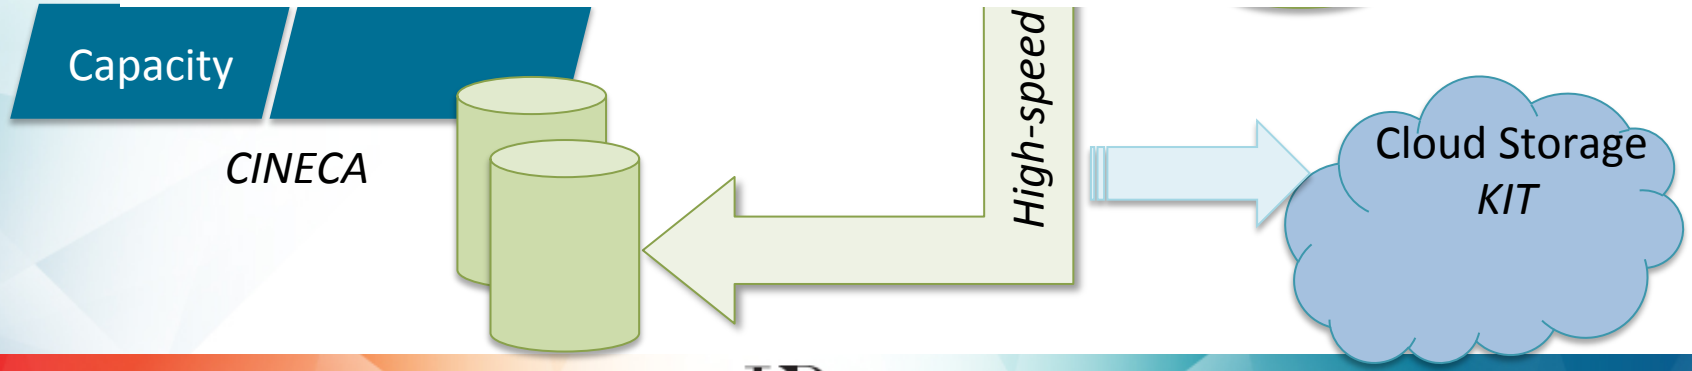

Roadmap for Brain Facility

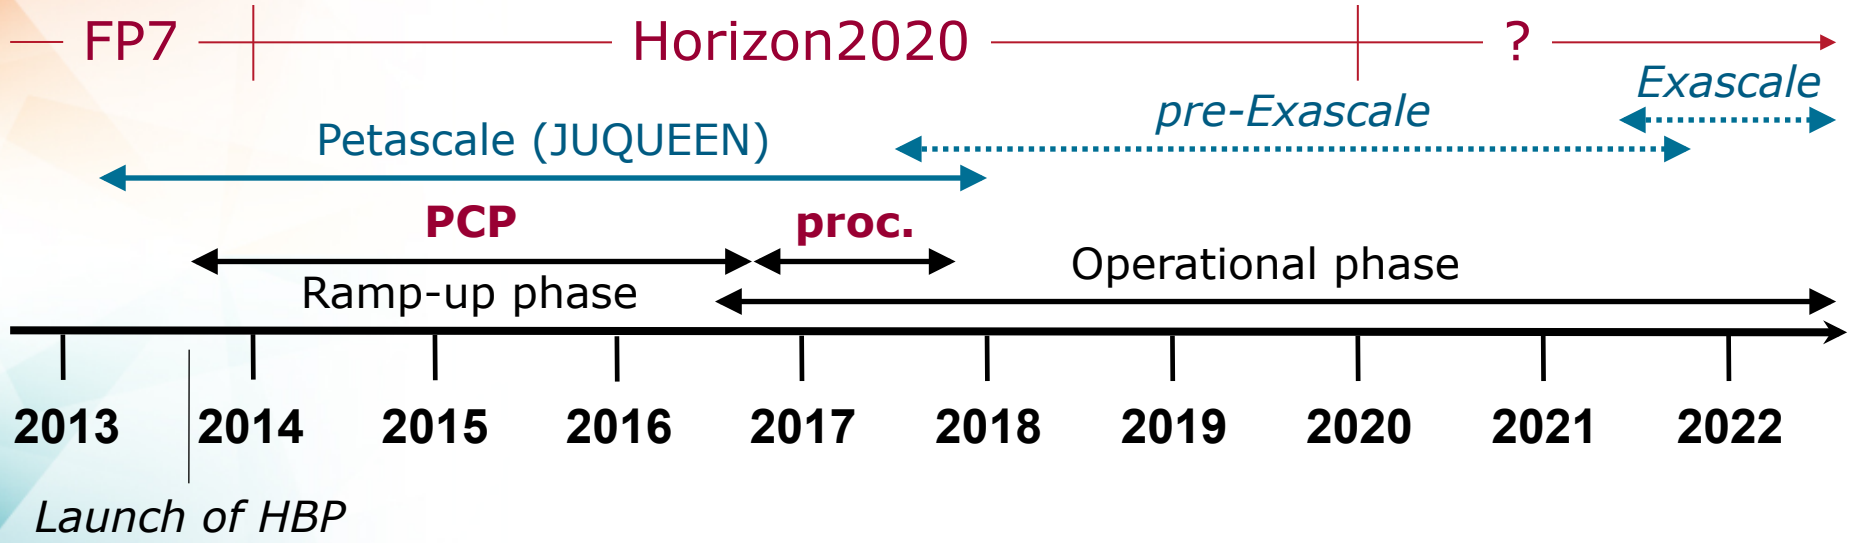

Contacts

The Human Brain Project Consortium

http://www.humanbrainproject.eu

Contacts:
Prof. Henry Markram
Director Blue Brain Project
Coordinator Human Brain Project
Email: henry.markram@epfl.ch

Prof. Thomas Lippert
Director of HPC Subproject
Email: th.lippert@fz-juelich.de

Human Brain Project

JÜLICH
FORSCHUNGSZENTRUM