inventors for the digital world

Modeling, Numerical Simulation, HPC & Cloud Towards HPC-Big Data Convergence

© Inria / picture

Inria: the French National Institute
For Computer Science and Applied Mathematics

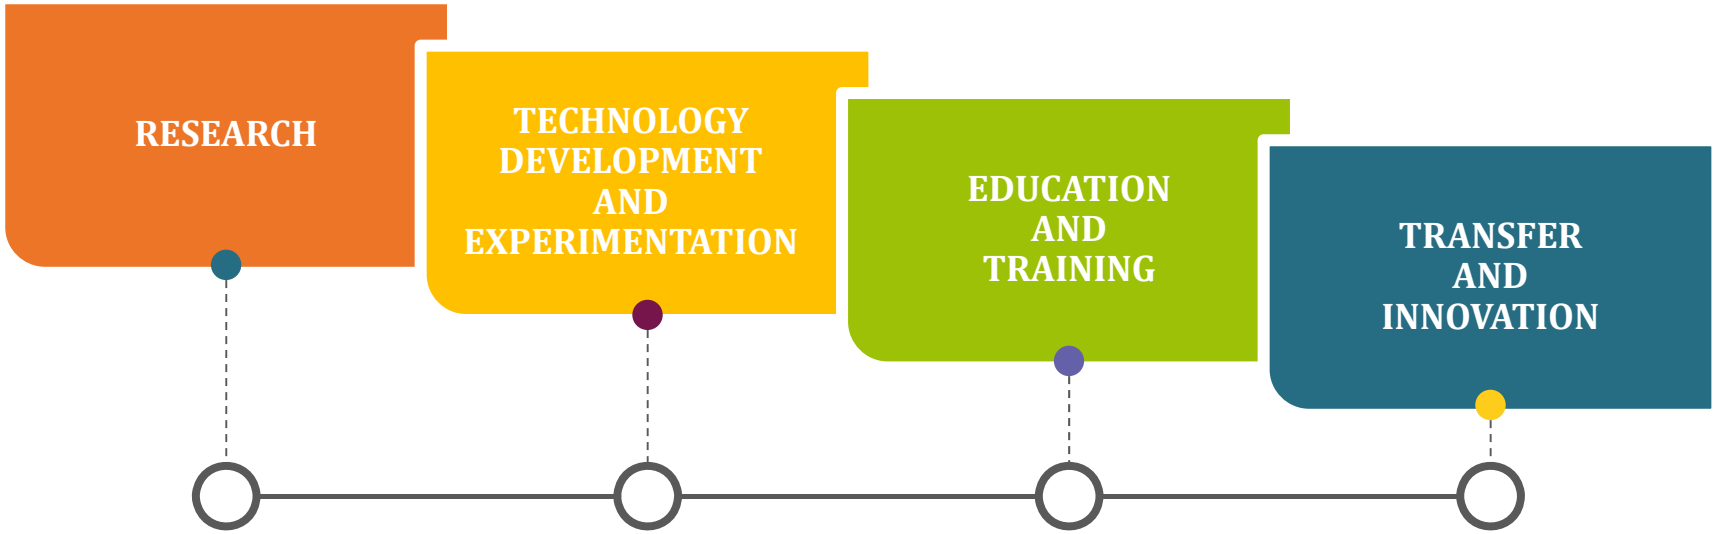

A scientific and technological public institution under the dual authority of the Ministry of Research and the Ministry of Industry

Dedicated to scientific excellence for technology transfer and society

Inria
inventors for the digital world

Science at Inria

ALGORITHMS & PROGRAMMING

DATA SCIENCE & KNOWLEDGE ENGINEERING

MODELING & SIMULATION

OPTIMIZATION & CONTROL

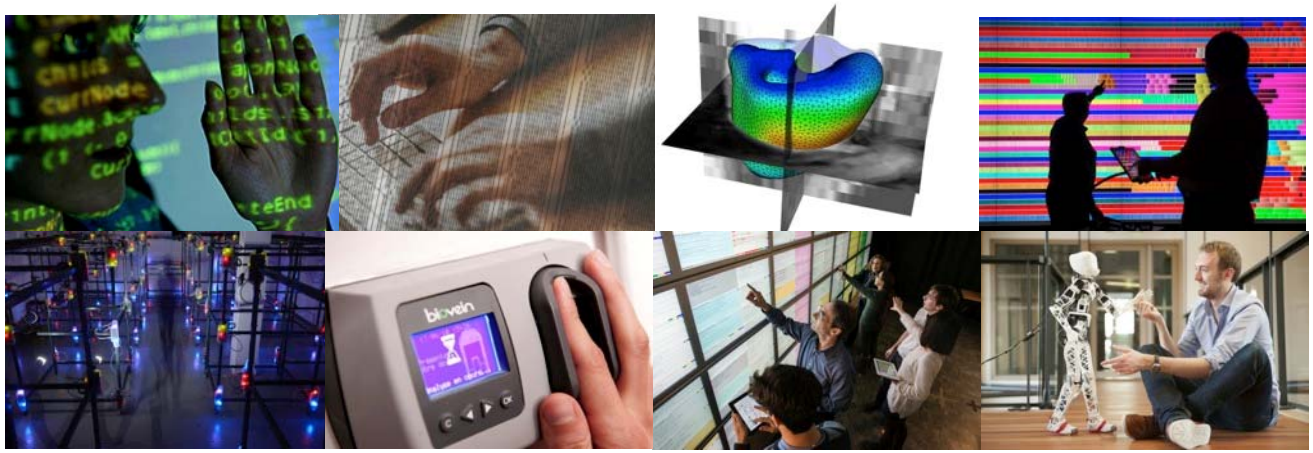

ARCHITECTURE, SYSTEMS & NETWORKS

SECURITY & CONFIDENTIALITY

INTERACTION & MULTIMEDIA

ARTIFICIAL INTELLIGENCE & AUTONOMOUS SYSTEMS

Socio-economic Areas of Application

HEALTH

ENERGY

SECURITY & RESILIENCE

ENVIRONMENT

CLIMATE

TRANSPORT

CULTURE & ENTERTAINMENT

ECONOMY

FINANCE

FOOD & AGRICULTURE

Technology Transfer Mechanisms

TRANSFER OF TECHNOLOGIES AND EXPERTISE

+

STARTUP CREATION

=

150
startups

including **75%** in activity or bought out

3 000
Jobs created

Joint laboratories
(joint labs, innovation labs, labcoms)

R&D partnerships
(collaborative projects)

Technology transfers
(software and patents)

Transfer of knowledge / know-how
(expertise, mobility)

Providing them with structural assistance (IT-Translation)

Facilitating funding support

Working in partnership and through networks (regional incubators)

Numerical Simulation, HPC & Cloud @Inria

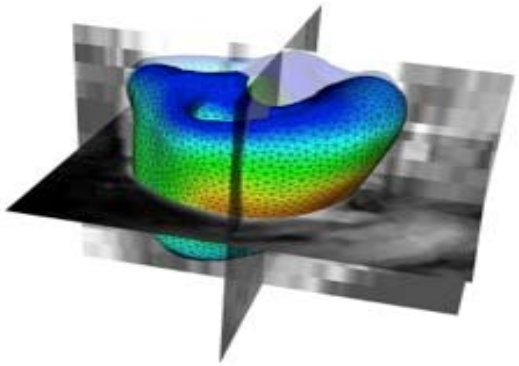

- Methodological work : modelling, numerical models, uncertainties, data analytics)
- Computer science for HPC : models and tools for programming, numerical libraries, performances, energy use, visualization
- Architecture and compilation
- And a lot work using HPC & Cloud: energy, environment, life sciences, …

Forces in Numerical Simulation and HPC

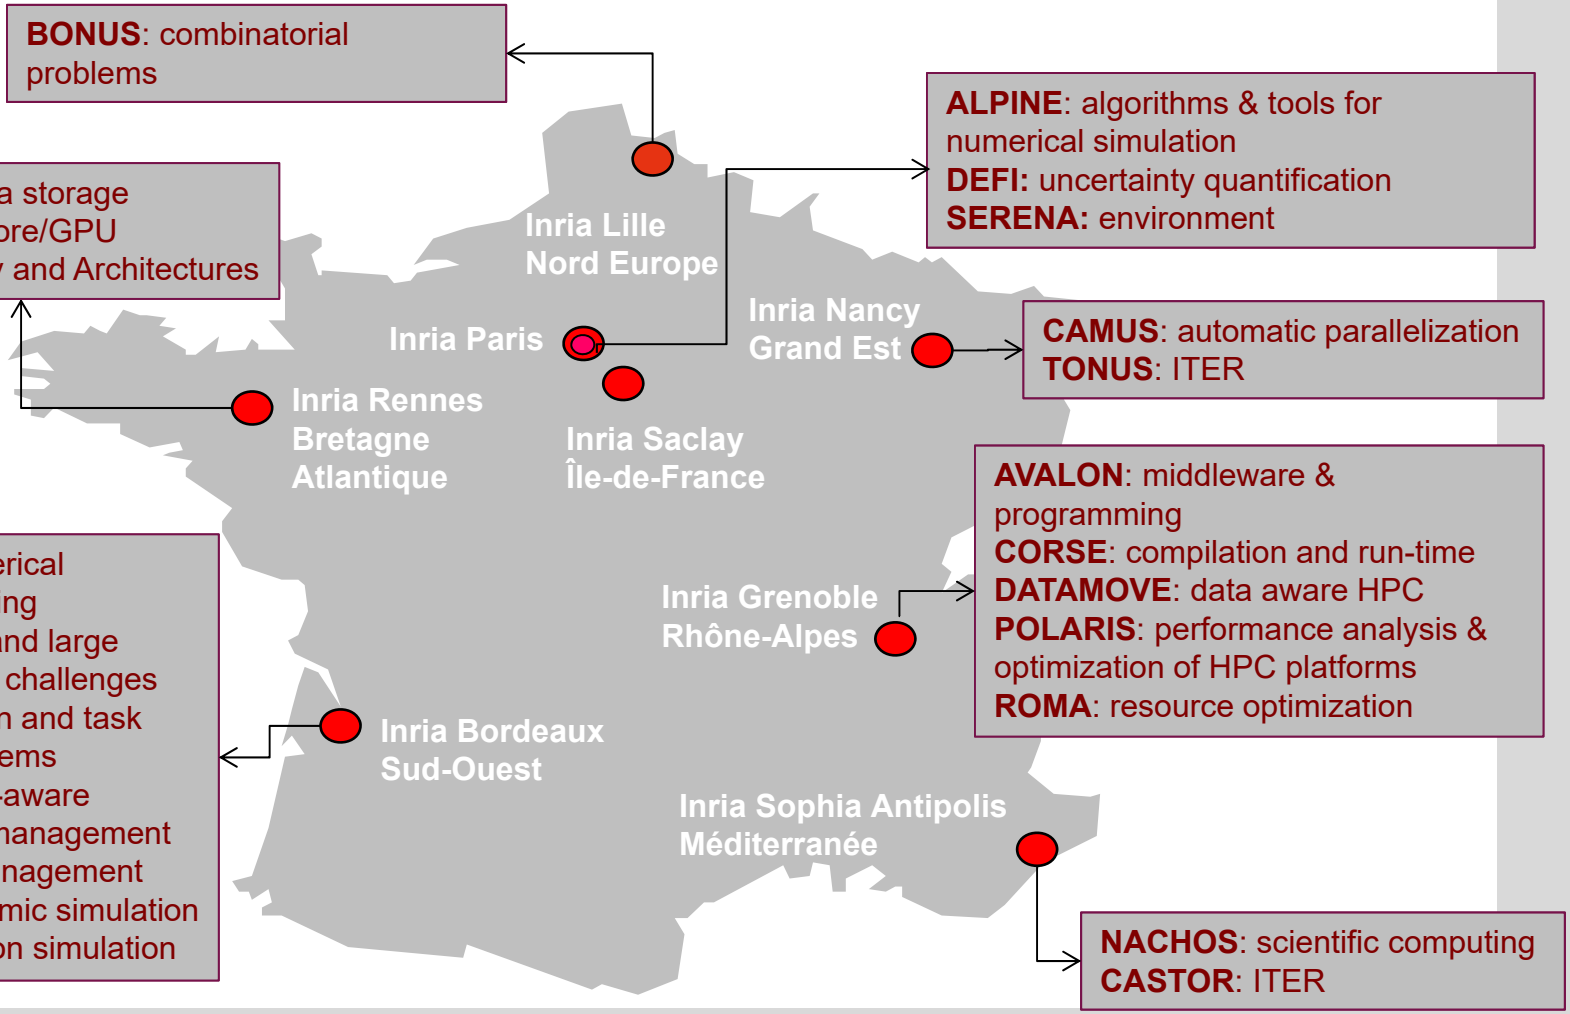

Strengths in Clouds & Data Management

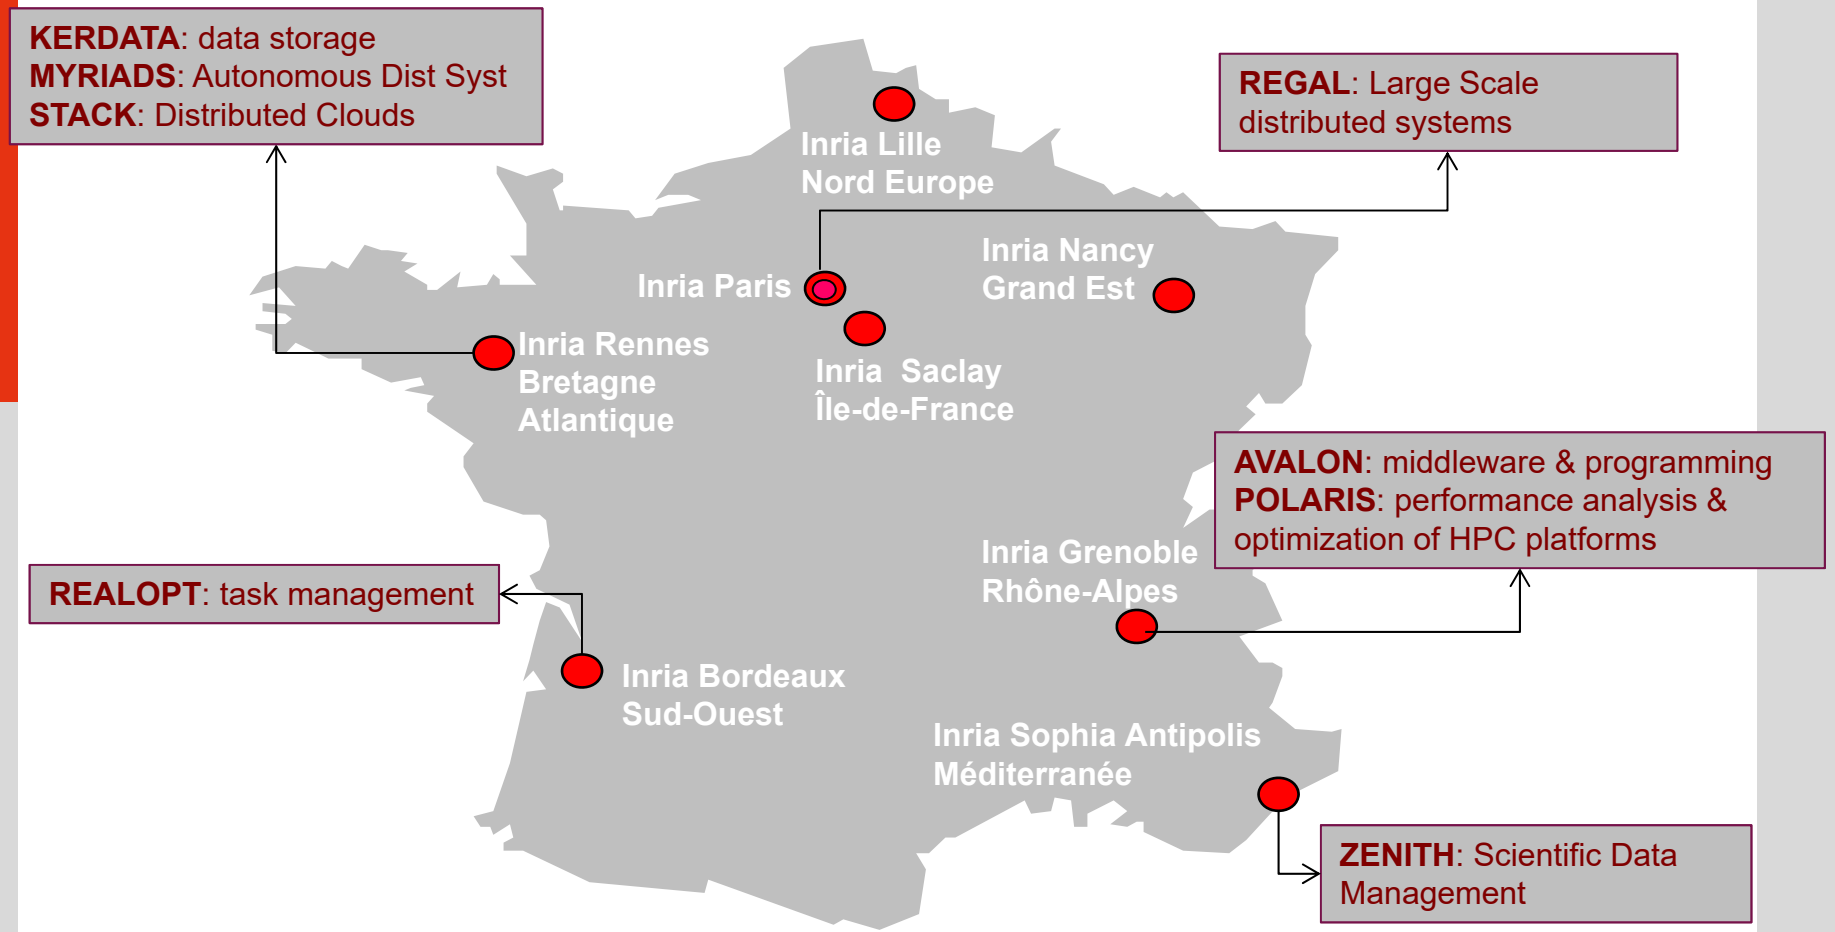

Examples of Research Projects within Inria

(IPL 4 years duration, ~1.5 m€)

CS2@Exa : *Computer and Computational Sciences at Exascale (until 2016)*

- participation to PRACE 4IP & 5IP European Projects (WP7 - Application Enabling and Support)

FRATRES : *Fusion Reactor Research and Simulation (launched in 2015)*

- fusion (ITER challenge) in relation with EoCOE (European Energy oriented Center of Excellence)

HAC-SPeCiS : *High-performance Application and Computers, Studying Performance and Correctness In Simulation (launched in 2016)*

- in relation with POP (new European Center of Excellence about Performance Optimization and Productivity), with SimGrid platform and with previous GRID5K project

DISCOVERY : *Distributed and Cooperative management of Utility Computing Infrastructures (beyond the clouds ?) (launched in 2016)*

- in relation with the I/O labs, a joint lab between Inria and Orange Labs.

HPC-BIG Data : *High Performance Computing and Big Data (launched in may 2018)*

ELCI - software environment for computation-intensive applications

Develop numerical simulation tools for HPC

- New generation of SW stack: supercomputer control, prog. & exec. Environment
- Meshing environment and tools
- Validation: better scalability, resilience, security, modularity, abstraction and interactivity of applications

Consortium

- ATOS/Bull (leader), CEA, Inria, SAFRAN, CERFACS, CORIA, CENAERO, ONERA, Univ. of Versailles and 2 SMEs (Kitware, AlgoTech)
- 36 months (terminated in September 2017)
- 150 PY

Contacts: Emmanuel.Jeannot@inria.fr and Laurent.Lefevre@ens-lyon.fr

Experimental infrastructure

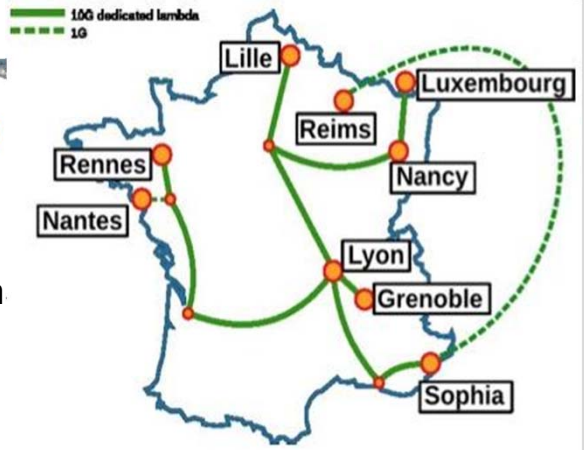

Testbed for research on distributed systems

- Born from the need of better and larger testbed
- HPC, Grids, P2P, and nowCloud computing and BigData system
- A complete access to the nodes' hardware in an exclusive mode
- Dedicated network (RENATER)
- Reconfigurable: nodes with Kadeploy and network with KaVLAN
- Current status
  - 10 sites, 29 clusters, 1060 nodes, 10474 cores
  - Technologies/resources: Intel, AMD, Infiniband, GPU clusters, energy probes
- Some Experiments examples
  - In Situ analytics, Big Data Management, HPC Programming approaches, Network modeling and simulation, Energy consumption evaluation, Batch scheduler optimization, Large virtual machines deployments

Future: SILECS

- New infrastructure based on two existing instruments (FIT and Grid'5000)
- New challenges: IoT and Clouds, New generation Cloud platforms and software stacks (Edge, FOG), Data streaming applications, Locality aware resource management, …

EoCoe: Energy oriented Center of Excellence

The Energy oriented Centre of Excellence in computing applications uses the tremendous potential offered by the ever-growing computing infrastructure to foster and accelerate the European transition to a reliable and low carbon energy supply using HPC (High Performance Computing)

- Main sites: Forschungszentrum Jülich (JUELICH, Germany) and Maison de la Simulation (CEA – CNRS – Inria, France)
- Partners : Max Planck, Fraunhofer (Germany) ; ENEA, UTN, CNR (Italy) ; CERFACS (France) ; PSNC (Poland) ; UBA (UK) ; CYI (Cyprus) ; ULB (Belgium) ; BSC (Spain)
- Strong integration in the European HPC ecosystem (PRACE, PATC network, ETP4HPC)

4 thematic pillars with

- Meteorology for Energy (solar and wind energy production)
- Materials for Energy (photovoltaic cells, batteries and super capacitors for energy storage)
- Water for Energy (geothermal and hydropower energies)
- Fusion for Energy

A transversal research to supply high-end scientific and industrial demands

- Numerical methods and applied mathematics
- Linear Algebra and scalable robust solvers
- System tools for HPC
- Advanced programming methods for Exascale, tools and services for HPC

International Joint Laboratory for Extreme Scale Computing (JLESC)
(launched in 2014)

- Director: F. Cappello (ANL), Executive Director for Inria: Y. Robert
- Partners: NCSA (US), ANL (US), Inria (FR), Jülich Supercomputing Centre (DE), BSC (SP), Riken (JP)
- Followup of the Inria/UIUC/NCSA Joint Laboratory Petascale Computing (JLPC) (2010-14)
- 9 Inria project-teams involved in the joint lab

Research around software challenges found in extreme scale High Performance Computers

- Scientific applications (big compute and big data)
- Modeling and optimizing numerical libraries
- Novel programming models and runtime systems managnement
- Resilience and fault tolerance
- I/O and visualization
- HPC & Clouds, data

Operation

- Two workshops and one summer school every year
- Short-term visits
- Long-term student exchanges

Forthcoming challenges for modeling and simulation

- Multi-scale and multi-physics modeling integrating uncertainties, model coupling
- Coupling between deterministic and probabilistic methods and models for complex multi-scale phenomena
- *Extract useful information from large data set to perform efficiently large scale simulations ; convergence between models, high performance simulation algorithms and codes and data analytics*
- Scalability of parallel numerical codes to exploit forthcoming large scale computation platforms

Forthcoming challenges for HPC & Clouds

Design of large scale platforms

- Fault-tolerance, More heterogeneity (Multicore, GPGPU, FPGA), Energy management, Networks

Scalability of big codes over new generation platforms

- Hierarchy of architectures, New paradigms, libraries, tools, Data movements

Convergence of HPC and Big Data paradigms and tools

- *Design of software stacks combining the strength of traditional HPC and the recent advances of data analytics*

Impact of new architecture paradigms on parallel algorithms

- 3D architectures, NVRAM, photonics

Programming large scale applications

- New programming paradigms, Links with compilers, runtime systems, middleware

Algorithm and software validation

- Methodology and Reproducibility issues

The HPC-BigData INRIA Project Lab

An INRIA funded project (2018-2022)

Gather teams from HPC, Big Data and Machine Learning to work on the convergence

HPC teams: DataMove, KerData, Tadaam, RealOpt, Hiepacs, Storm, Grid'5000

Big Data & IA teams: Zenith, Parietal, Tao, SequeL, Sierra

External partners:

- Academic: Argonne National Lab (USA), Lab Biologie Théorique (CNRS Paris)
- Industry: ATOS/Bull, ESI-group

HPC versus BigData/AI

HPC

Big Data

Parallelism for scalability

Performance comes first
Low level programming (MPI+OpenMP)
Thin software stack
Stable software libs.
HPC centers

Ease of programming comes first
High level programming (Spark, TensorFlow)
Thick software stack
Quickly changing software libs.
Cloud platforms

Jobs run a few hours on thousands of cores :

- Sensitivity Analysis : 30 000 cores for 1h30 [Terraz'17]
- Exastamp material simulation: 8000 cores for a few hours

Jobs run a few days on tens of nodes :

- Pl@ntNet learning: one week on 4 GPUs
- AlphaGo Zero ltraining: 70 hours on 64 GPU workers and 19CPU parameter [Silver'17]
- ResNet-50 on 256 GPUs in 1 hour (mini-batch training) *[Goyal 2017]*

The HPC and BigData/AI Convergence

Three Research Directions :

- Infrastructure and resource management
- HPC acceleration for AI and Big Data
- AI/Big Data analytics for large scale scientific simulations

Some of our Software Assets

scikit learn

Machine Learning in Python

Light yet Flexible
Batch Scheduler

StarPU

Task Programming for
Hybrid architectures

FlowVR, Melissa, Damaris

On-line data processing
engines for HPC

Pl@ntNet

Deep Learning based App
for plant identification

Inria inventors for the digital world

Infrastructure and Resource Management

HPC Infrastructure for AI:

New needs:

Accelerators (GPUs or other)
Large resident data sets (learning & benchmarks) (PlantNet: 10 TB of raw data)
Very long runs (days)
Fast changing software stacks (TensorFlow, PyTorch)

On-going work on identifying and experimenting AI/HPC compliant resource sharing approaches

Playground: Grid'5000, Genci experimental GPU cluster

Get data close to the compute nodes:

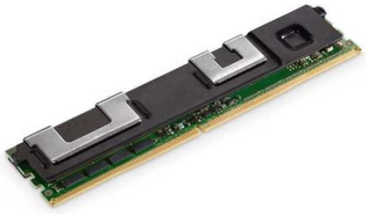

One fundamental difference on HPC versus Cloud platforms:

External file system versus on-node disks

On-node persistent storage for energy and performance (burst buffers, NVRAM):

Locality aware resource management

AI/Big Data Analytics for Large Scale Scientific Simulations

Molecular dynamics trajectory analysis with deep learning:

Dimension reduction through DL

Accelerating MD simulation coupling HPC simulation and DL

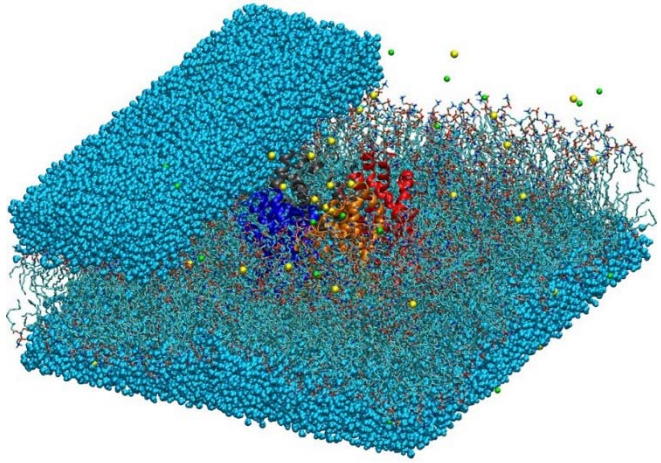

Flink/Spark stream processing for in-transit on-line analysis of parallel simulation outputs

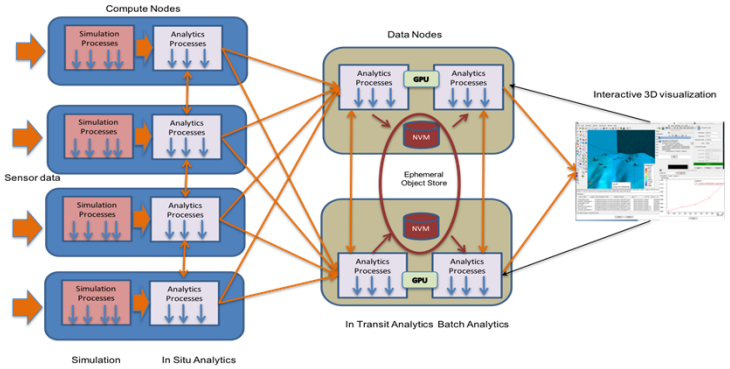

inventors for the digital world

HPC for AI and Big Data

TensorFlow graph scheduling for efficient parallel executions:

Scheduling for automatic differentiation and backpropagation

Recompute versus store frontward results

Linear algebra and tensors for large scale machine learning

Accelerating Scikit-Learn with task-based programming (Dask, StarPU)

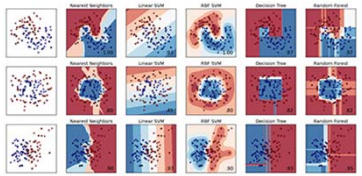

Large scale parallel deep reinforcement learning:

Self-learn to play Atari games

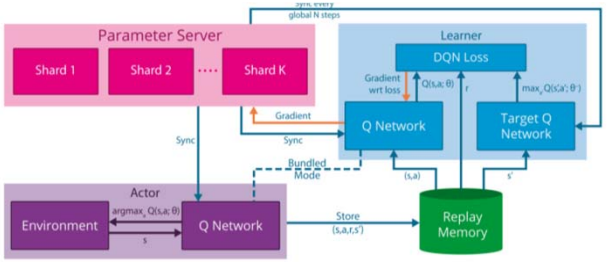

[Nair et al. 2015]

TensorFlow

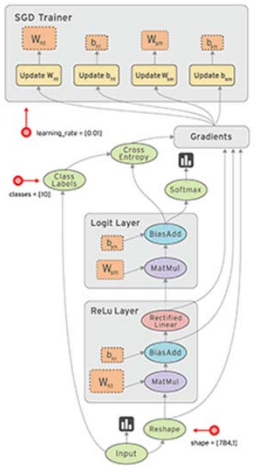

ResNet 34

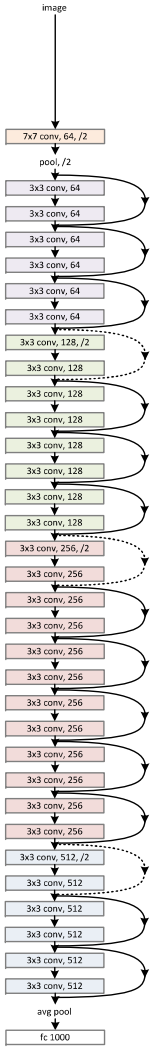

inria inventors for the digital world

More information

Isabelle RYL, Deputy CEO for Transfer and Industrial Partnership

Bruno RAFFIN, Research Director, head of DataMove project-team

Frédéric DESPREZ, Deputy Scientific Director (Frederic.Desprez@inria.fr)

- Networks, Systems and Services and Distributed Computing

Jean ROMAN, Deputy Scientific Director (Jean.Roman@inria.fr)

- Applied Mathematics, Computation and Simulation