The HPC-BigData Project Lab at INRIA

Emmanuel Jeannot*
INRIA Bordeaux/LaBRI

*Many slides from Bruno Raffin (Inria Grenoble)

Bordeaux, January 2019

HPC versus Big Data

Traditional HPC:

Applications: large scale scientific computing

Machines: supercomputers

Programming: MPI+ OpenMP

Traditional Big Data:

Applications: web and business analytics - descriptive and predictive (include machine learning and deep learning)

Machines: cloud infrastructures

Programming: Map/Reduce, ML and DL dedicated libraries

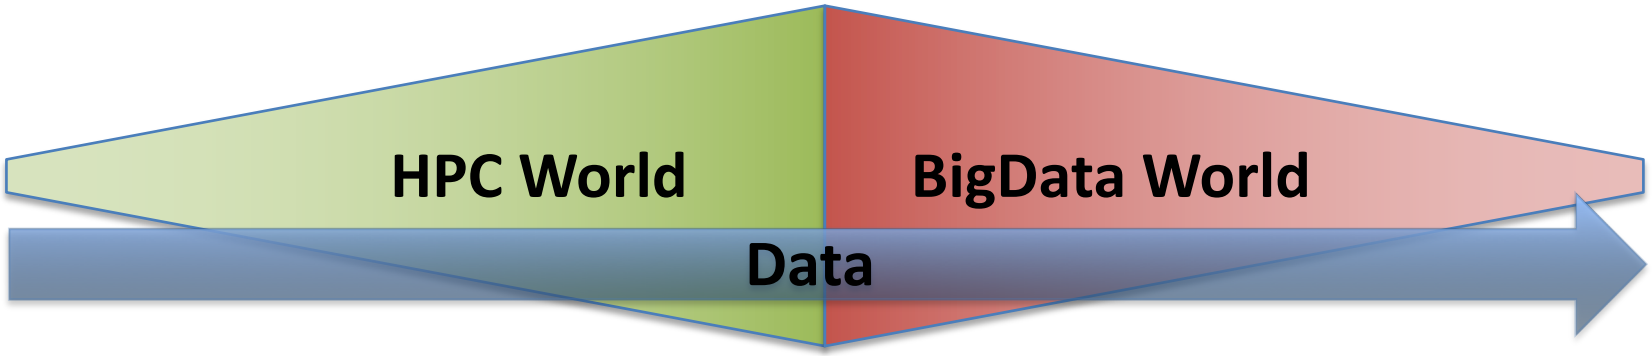

HPC and Big Data: a Closer Look at the Differences

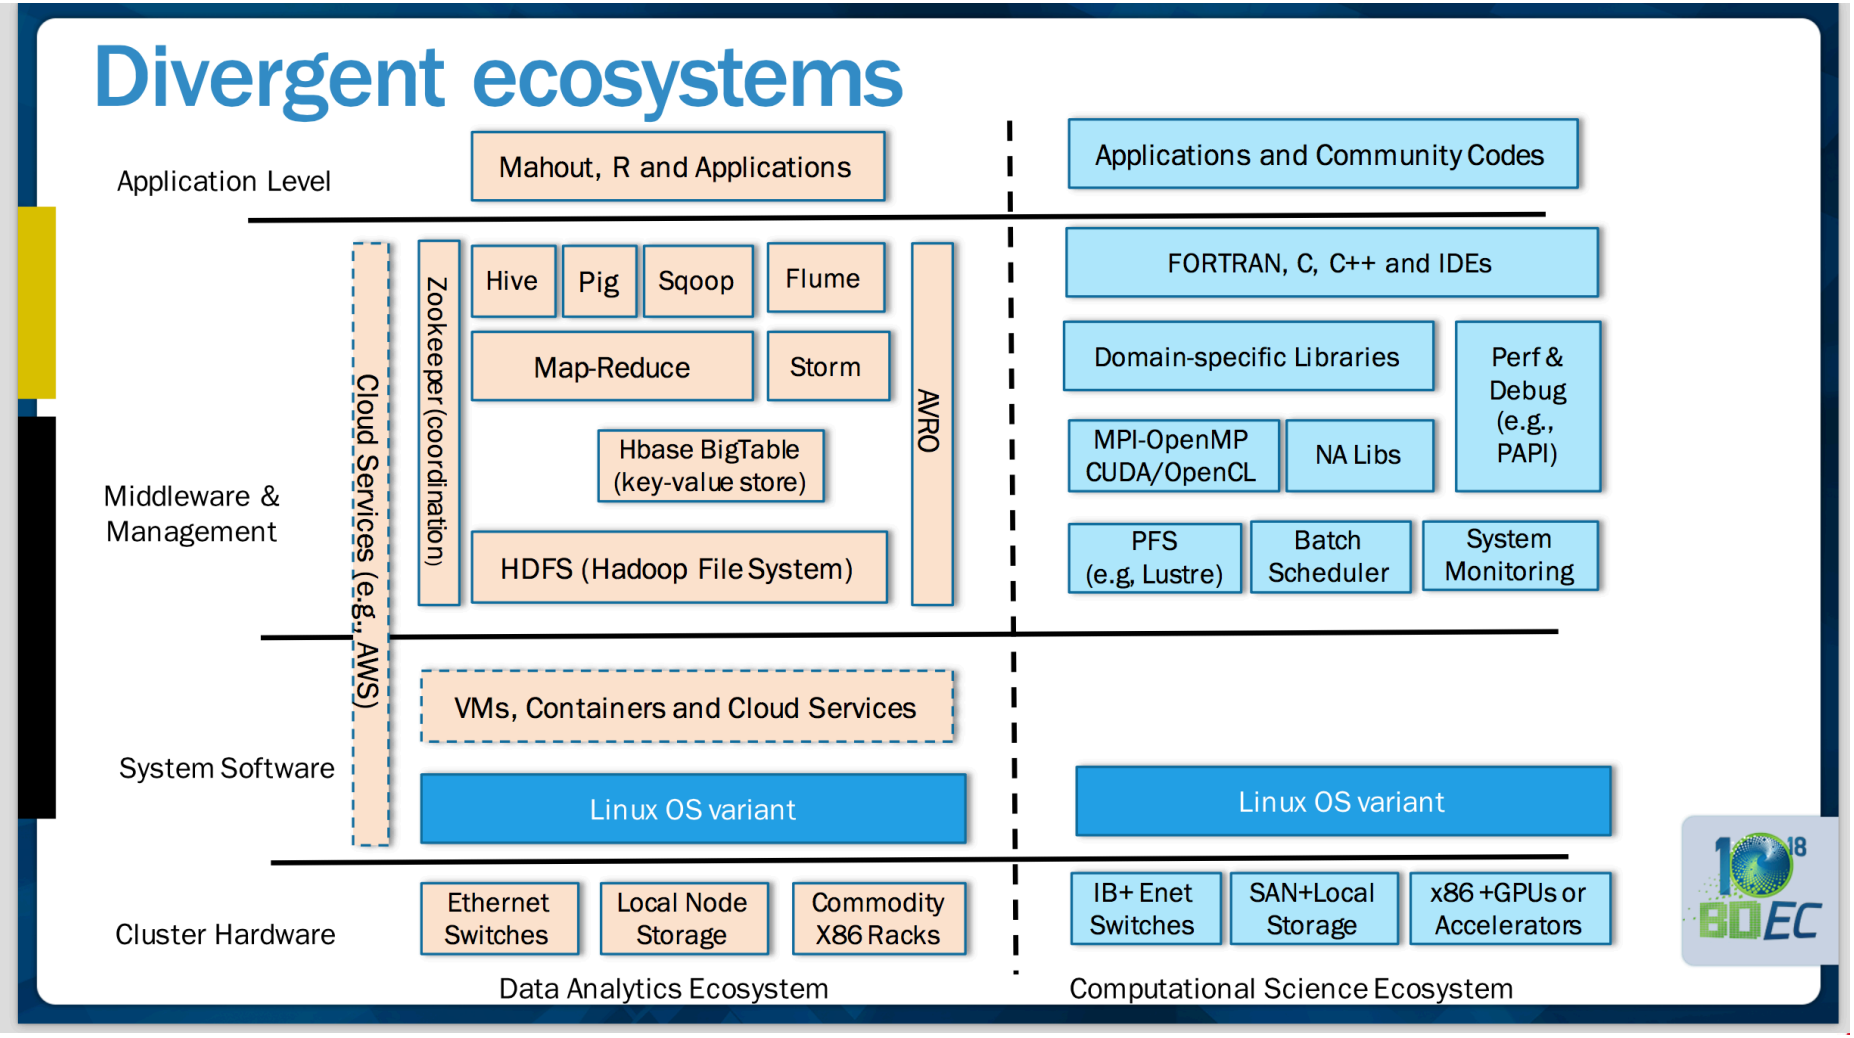

The HPC-BigData Integrated Project Lab (IPL HPC-BigData)

An INRIA funded project (2018-2022)

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

Eventually more HPC-IA than HPC-BIGDATA

Teams/People Involved

INRIA Teams

HPC:
DataMove: Bruno Raffin, Olivier Richard
KerData: Gabriel Antoniu, Alexandru Costan
Tadaam: Emmanuel Jeannot, Guillaume Aupy
RealOpt: Olivier Beaumont
Hiepacs: Olivier Coulaud
Storm: Samuel Thibault
Grid'5000: Pierre Neyron

BigData/ML:
Zenith: Patrick Valduriez et Alexis Joly
Parietal: Gael Varoquaux
Tao: Guillaume Charpiat
SequeL: Philippe Preux
Sierra: Francis Bach

Non INRIA Partners

External Research Teams:
ANL (USA): Rob Ross, Tom Peterka
LBT (CNRS Paris): Marc Baaden, Jérome Hénin -> Molecular Dynamics

Companies:
ATOS/Bull -> HPC Vendor
ESI-group -> Computational Mechanics

High Performance Analytics for Scientific Computing Apps

- In Situ/Transit analytics combining HPC (FlowVR, Damaris, Decaf) and Big Data solutions (Spark, Flink)
- On-line data streams assimilation combining ML and numerical simulation.

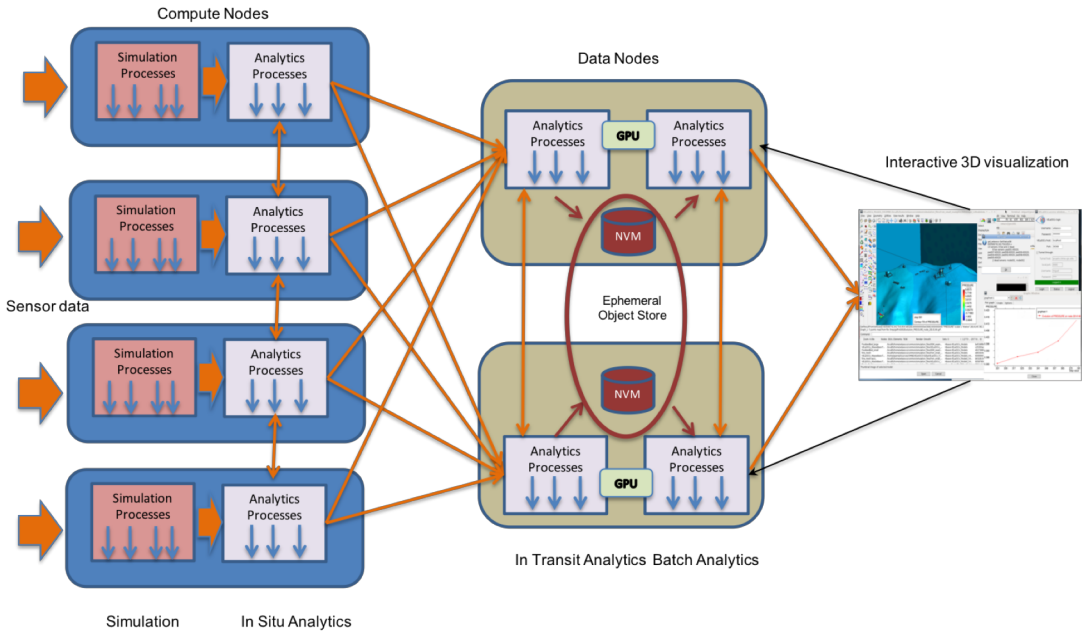

ML Framework

Toward a parallelization based on task programming with Dask (Python framework).

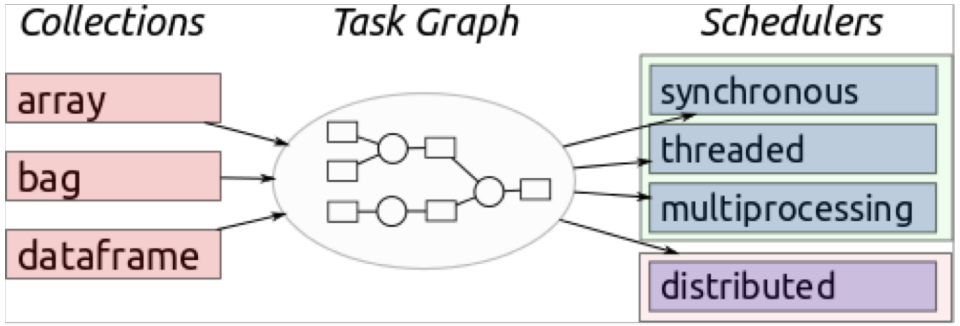

Use StarPU as a runtime

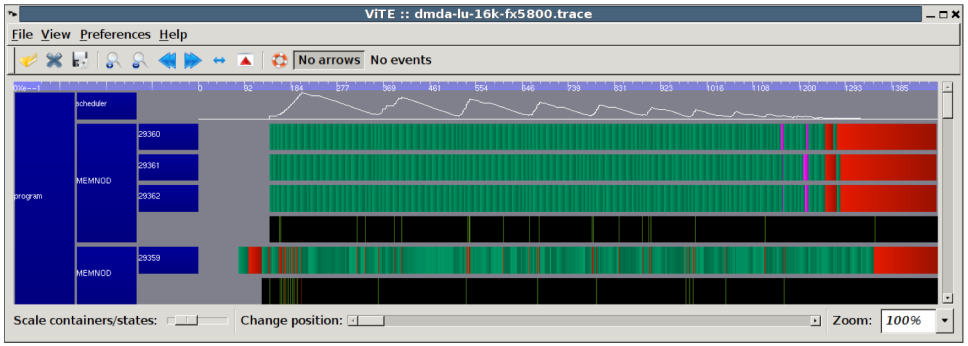

Infrastructure and Resource Management

Two different machine models: supercomputer versus cloud

- Centralized file system / distributed on-node storage
- Exclusive resources / virtual machines and multitenant apps
- High performance architectures / cost efficient architectures

Toward a converged Architectures ?

- How to handle on-node storage (elastic/ephemeral store)?
- How to perform locality-aware scheduling (dynamics and static)?
- What abstractions for convergent storage (key-value store, blob store)?
- How to leverage ML/DL approaches for scheduling?

How application can make the best possible use of the available resources

Problematic:

- Allocate data
- Partition data
- Reserve resources
- Control affinity
- Map computation
- Manage contention
- Optimize communication
- Access storage
- Perform visualization

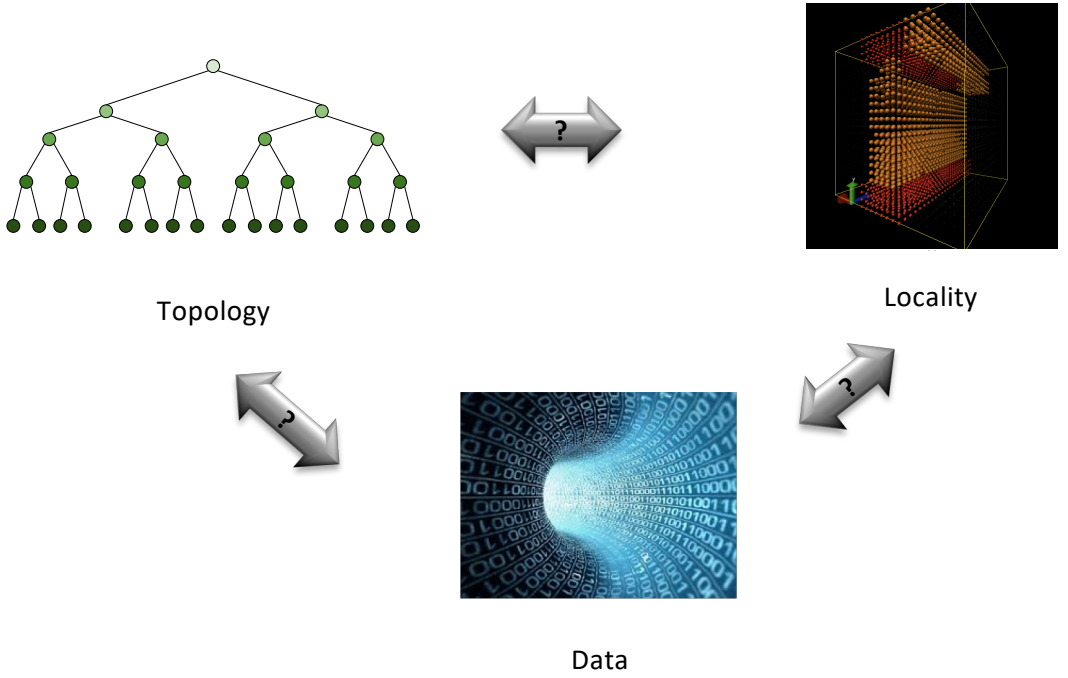

Conclusion

Follow-up of the challenge *When extreme- scale computing meets data-intensive science* for the next Inria strategic roadmap

Sparked off more interest than expected. Very motivating discussions between partners

Expect to gain momentum as we understand each other domain better:

- Start with modest direct knowledge/technology transfers
- Go up to the co-design of innovative solutions

Very important and timely topic: we will be happy to collaborate with other teams.