Michael Malms

Maria Perez

Christian Garmat

What is next?

Post-H2020 Vision

Questions and Comments – 15 min at the end

ETP 4 HPC

BDV BIG DATA VALUE ASSOCIATION

ICT 2018 networking session

HPC & HPDA R&I priorities in Europe 2019 - 2022

ETP4HPC

”Building a globally competitive European world-class HPC technology value chain “

91 members
(as of July 2018)

- 56 full
- 35 associated

- 51 private
- 34 SMEs
- 17 larger companies
- 37 research organisations

Established in 2011
Officially a Dutch Association since 2012

the starting point…(1)

A new Instrument
The main Specifications

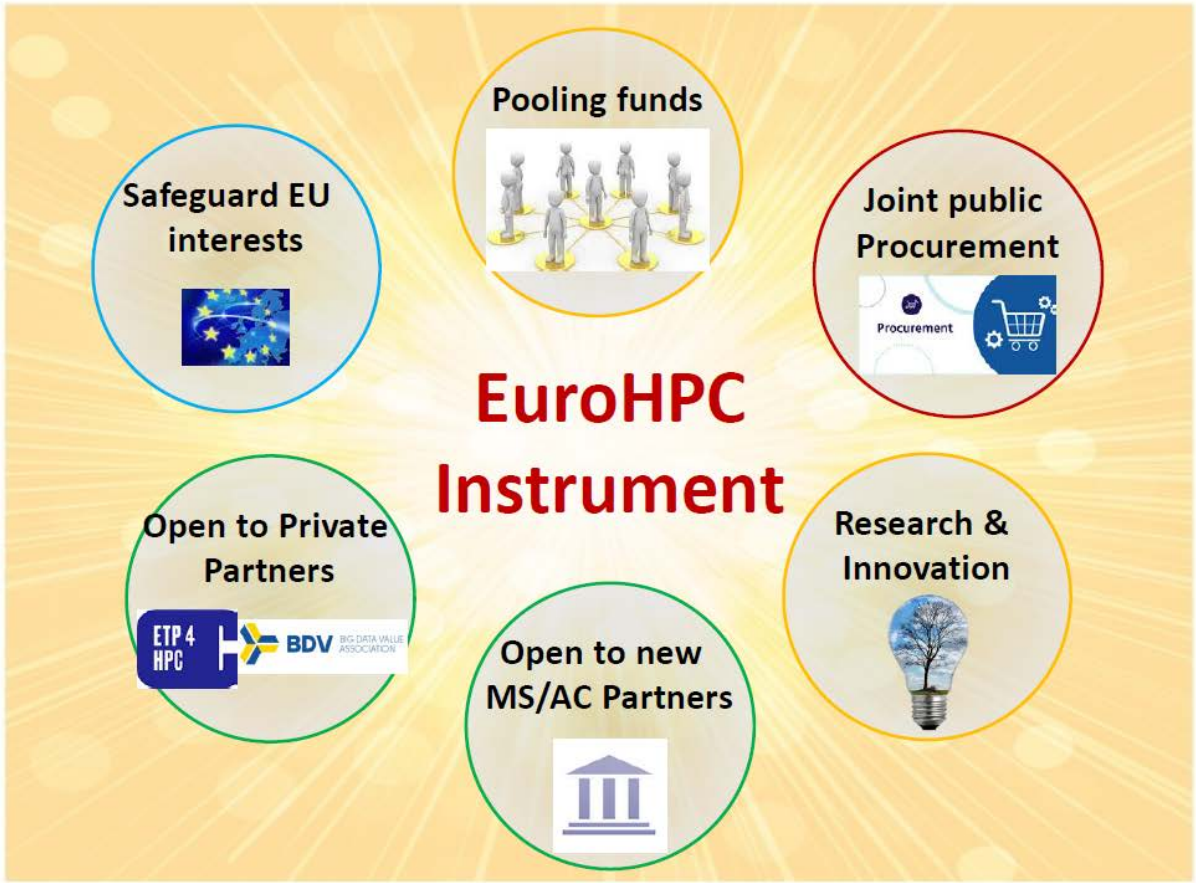

the starting point…(2)

- The central element of EC direction:

  European Sovereignty in HPC technology provisioning

  First implementation step is European Processor Initiative

- ETP4HPC and BDVA are prepared to provide input to the Research and Innovation programs supporting this direction

➢ The recommended WP19-20 sets the priorities accordingly, see next pages

the starting point…(2)

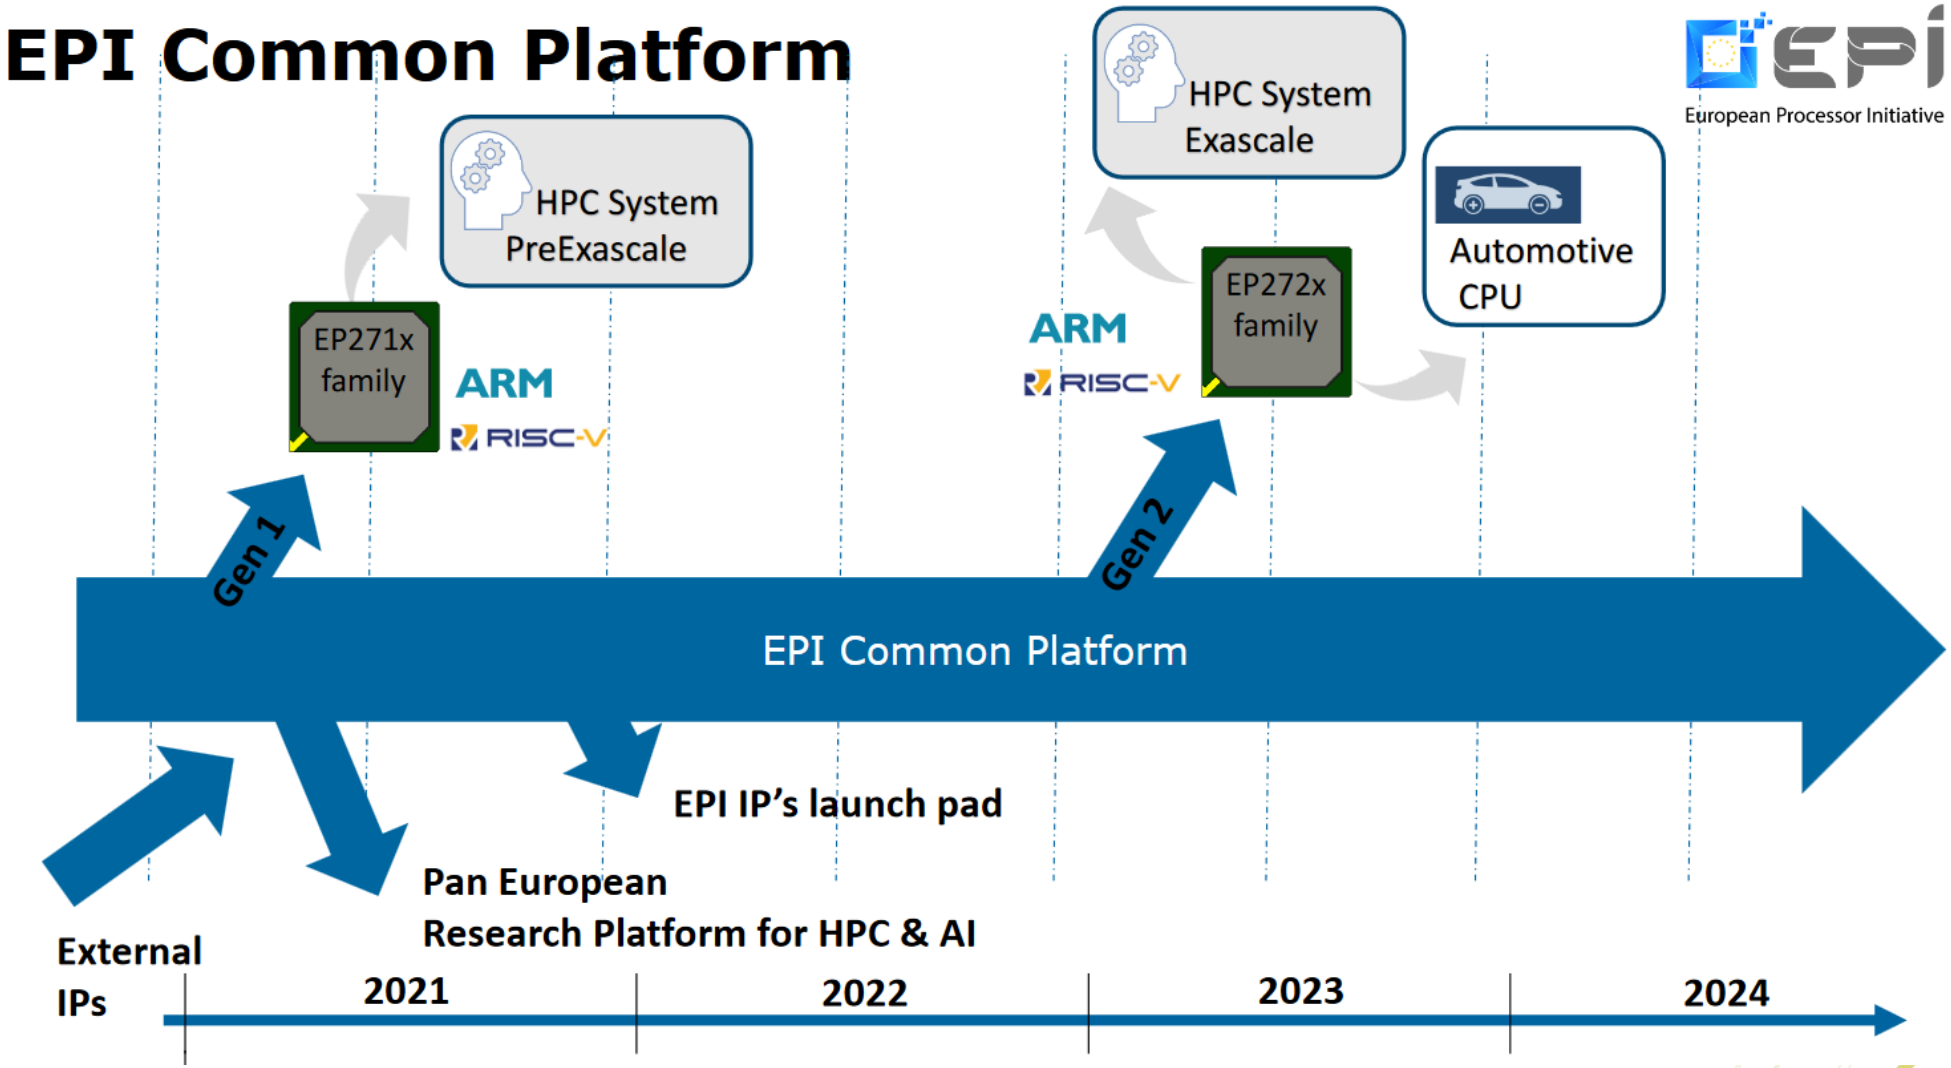

Context of SRAs and proposed WPs so far

H2020

HPC-cPPP

SRA 1 2013 | SRA 2 2015 | SRA 3 2017

Priorities based on input from stakeholders, end-users

End-to-end coverage addressing entire HPC stack

Processor platform agnostic

Open collaboration

Horizon/Digital Europe

??

EuroHPC

H2020

BDVA

Introduces a new section in its BDV PPP SRIA in 2017 for:

SRIA 4 2017

HPC and Big Data enabled Large-scale Applications

Handling massive amounts of Big Data

Large-scale HPC-enabled industrial pilot test-beds for BD applications

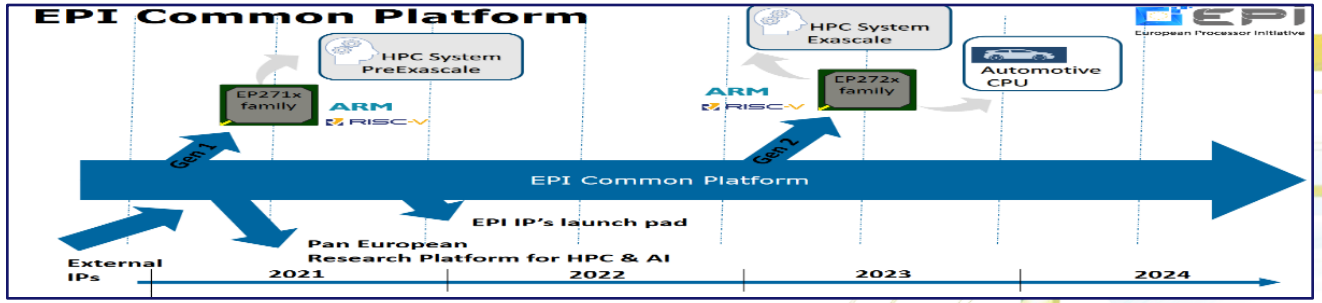

Two parallel focus areas in the next years:

EPI's focus will have to be:

chips
spec, sim, design, verification, tools, mfg

fab
7nm, yield, packing

processor integration
processor+accel+memory+network+clocks+power

eco-system
spec, tools, demo boards, appl.-doc.,SDK

Suggested focus in WP19-20

EuroHPC

?? ➡

- continued drive for system level excellence
- gain experience: HPC+BD+AI+IOT
- prepare for EPI eco-system
- WP21-22 onwards:
  - support EPI ecosystem growth
  - integrate into advanced platforms

WP19-20 – part one:

"HPC System focus"

Continued drive for system level excellence / prepare for EPI eco-system:

1A) Exascale focus:

- Projects need to demonstrate a **clear path to productization**, target a **high TRL** and focus at solving challenges in the **context of extreme scale system** designs.
- The results of the projects should both result in **supported products** at the time of exascale as well as **creation of IP** protected in Europe

  - heterogeneous and modular designs, integration of new accelerator types
  - use of advanced accelerators mixed with general-purpose CPUs , appropriate coherent HW interfaces should facilitate their interconnection
  - significant improvements of scalability of interconnects for coping with extreme scale system architectures composed of 10,000s of nodes
  - Mixed precision technologies able to execute efficiently AI-class of workloads will have to be adopted.
  - Programming models, associated run-time systems and compilers to ensure code maintainability, functionality and portability
  - ……(see proposal)

1B) Integration of upsteam technologies

- neuromorphic computing
- chip-to-chip interconnect based on photonics

Big Data: What will it take to keep going?

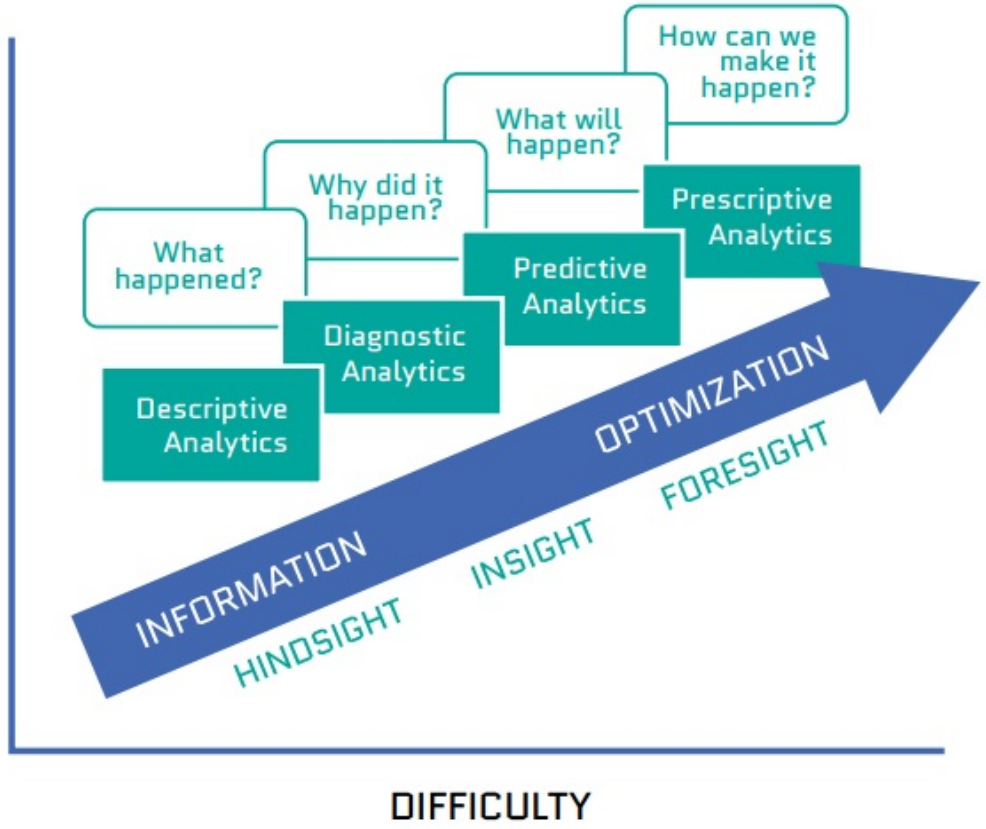

Source: National Big Data Congress 2014

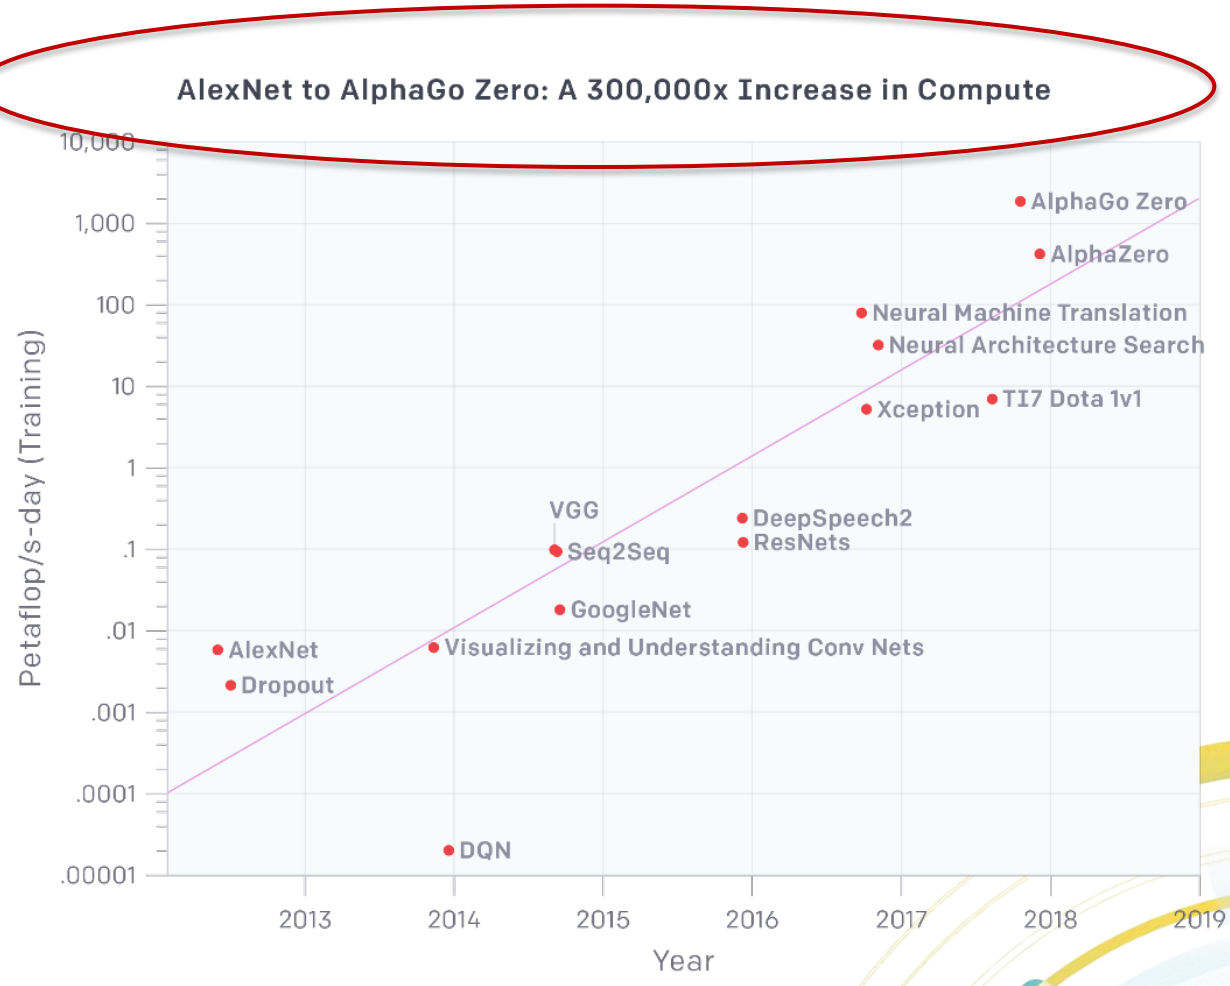

Amount of compute that is used to train a single model
Source: https://blog.openai.com/ai-and-compute/

WP19-20 – part two:

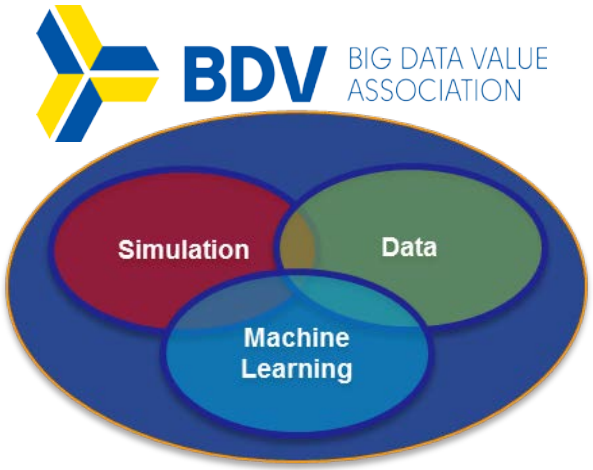

"Application focus"

- **Gain experience: HPC+BD+AI+IOT**
  - Design and integration of simulation, data analytics and other AI capabilities into HPC
  - Explore and demonstrate the technical **feasibility and the value of mixed/integrated Simulation & Data Analytics & AI IT platforms** and its application to different domains.
  - Non-technical challenges such as sharing public and private infrastructures, interoperability and standardization and efficiency in cost, energy and performance need to be tackled at the same time.
  - **CoEs or iSpaces to help with** adoption of Big Data and AI technologies that most benefit scientific and industrial modelling and simulation workloads in their workflows

HPC in Horizon/Digital Europe

Upstream technologies

- new memory/storage techn.
- nanoelectronics
- photo-electronics
- ….

Societal Challenges

“Missions”
(Horizon Europe)
e.g.: “Halving the human burden of dementia by 2030”,
e.g.: “Reach net zero greenhouse gas emissions

“Thematical Clusters”
(Horizon Europe)

- Health,
- Inclusive and Secure society,
- **Digital and Industry**,
- Cli
- Fo

Digital Europe: Supercomputing
Wide use of supercomputing both in areas of public interest and by industry (SMEs)

- …
- Artificial Intelligence & Robotics
- Advance Computing & Big Data
- NGI

Industrial users

- energy
- aviation
- automotive
- manufacturing
- pharmaceuticals
- ….

- New products & services
- Higher efficiency
- Lower cost

Application and use scenarios (Simulation, Analytics, AI, IOT)

Deployment in: HPC centre | Cloud | Fog | Edge

“non HPC”

Applications development: design, algorithms, methods, languages, tools

Technology for HPC infrastructure: architectures, hardware, software, algorithms, programming, tools

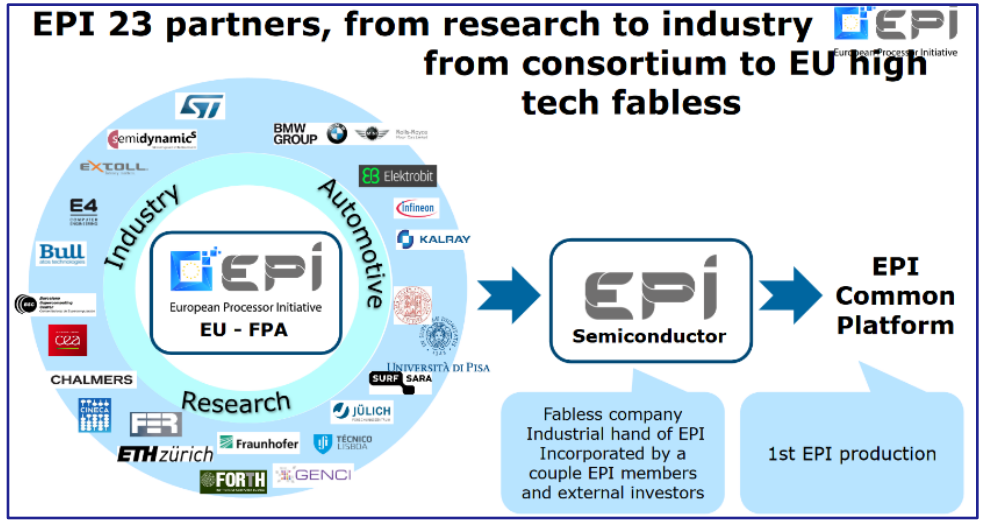

interaction for forthcoming SRAs

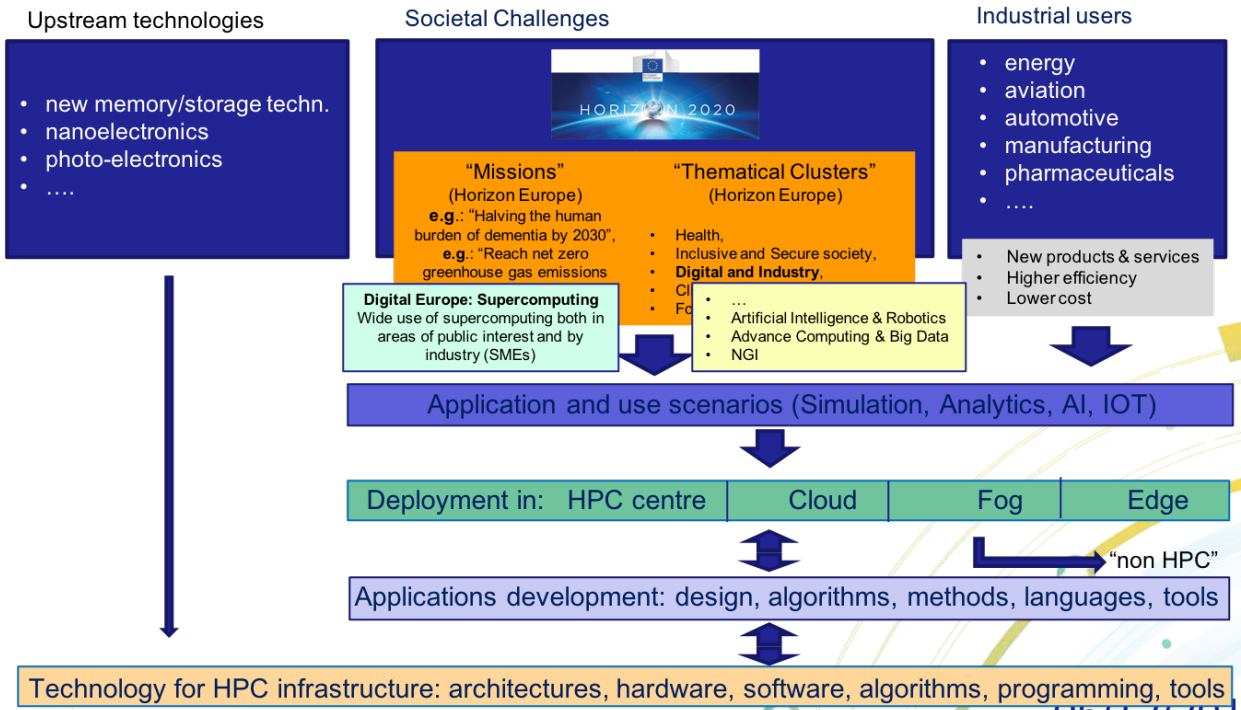

application / use case driven ecosystem growth support

ETP 4 HPC

BDV BIG DATA VALUE ASSOCIATION

THANKS!

Common priorities of Big Data and HPC beyond H2020

María S. Pérez, Universidad Politécnica de Madrid
Jim Kenneally, Intel Corporation
Chairs of BDVA's HPC-BD Group

Big Data: What will it take to keep going?

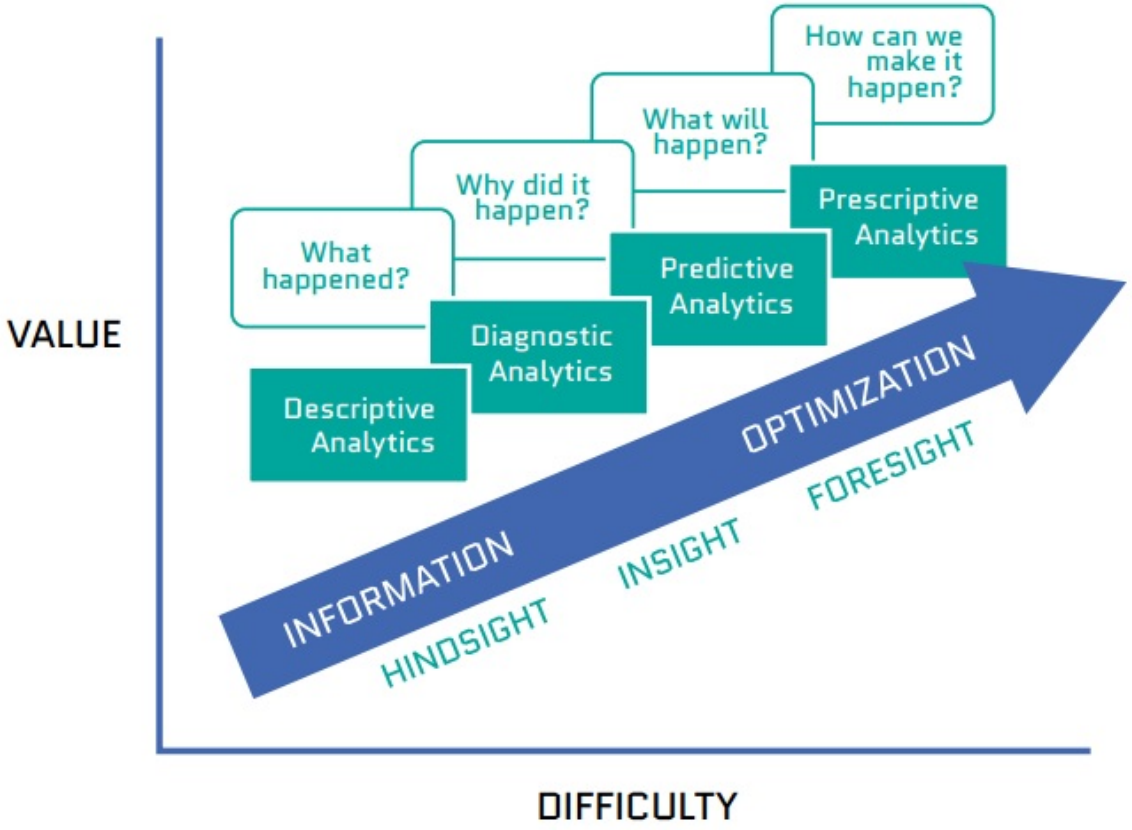

Source: National Big Data Congress 2014

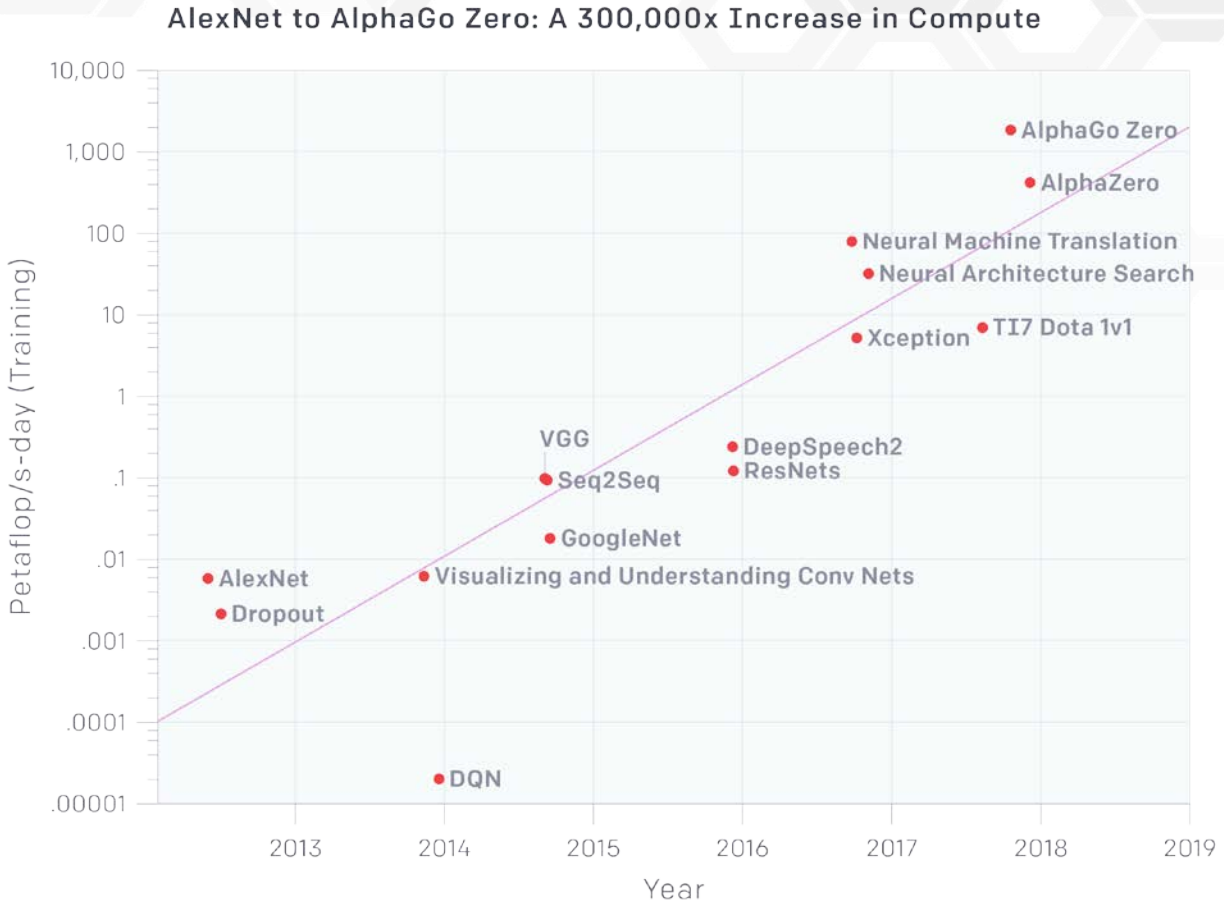

Amount of compute that is used to train a single model
Source: https://blog.openai.com/ai-and-compute/

Big Data & HPC – hybrid use cases

HPC (Numerical) Models / Simulations

+

BD Analytics (Cloud)

Navigation

IIOT

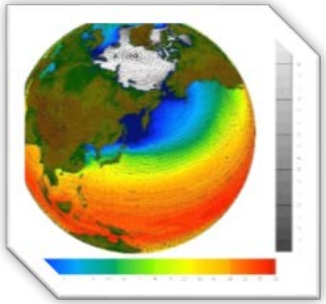

Weather

Health

ADAS

BD Analytics (Edge)

IoT / CPS / Edge /…

Digital Twin

Green Energy

Grid

www.bdva.eu

ABC (AI, Big Data & HPC intersecting)

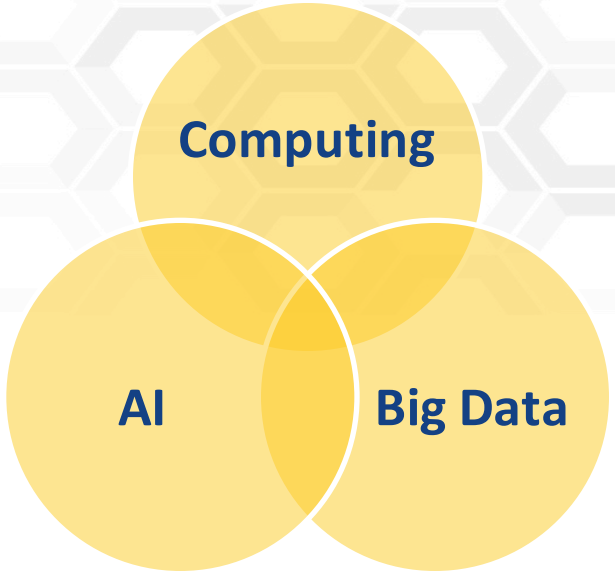

"Most of the knowledge in the world in the future is going to be extracted by machines and will reside in machines"
Yann LeCun, Director of AI Research, Facebook

ETP 4 HPC
EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

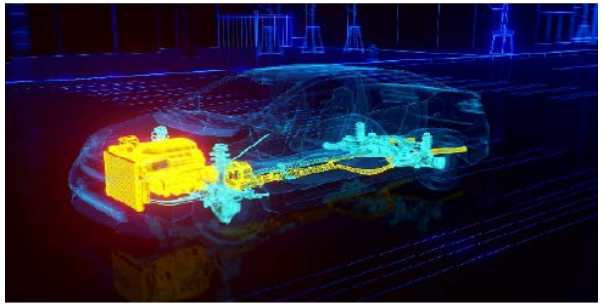

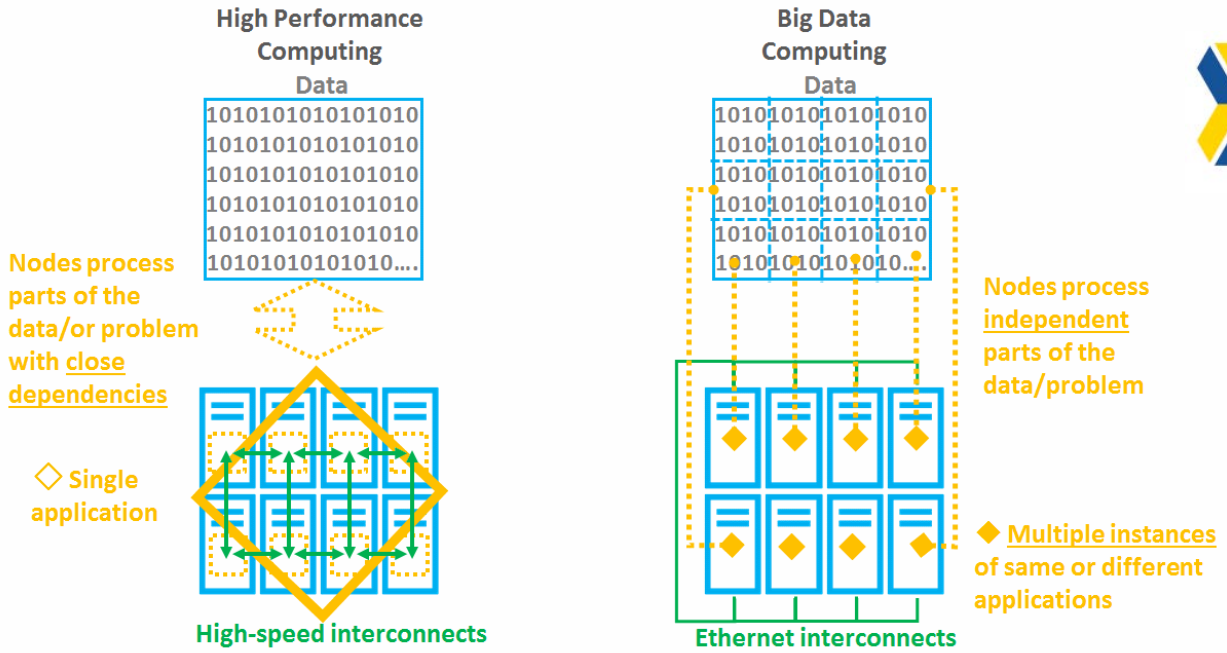

BDV BIG DATA VALUE ASSOCIATION

Source: Kenneally, Jim, and Hoppe, Hans-Christian, editors. *The technology stacks of High Performance Computing and Big Data Computing: What they can learn from each other*. 2018

www.bdva.eu

Transfer capabilities & identify gaps

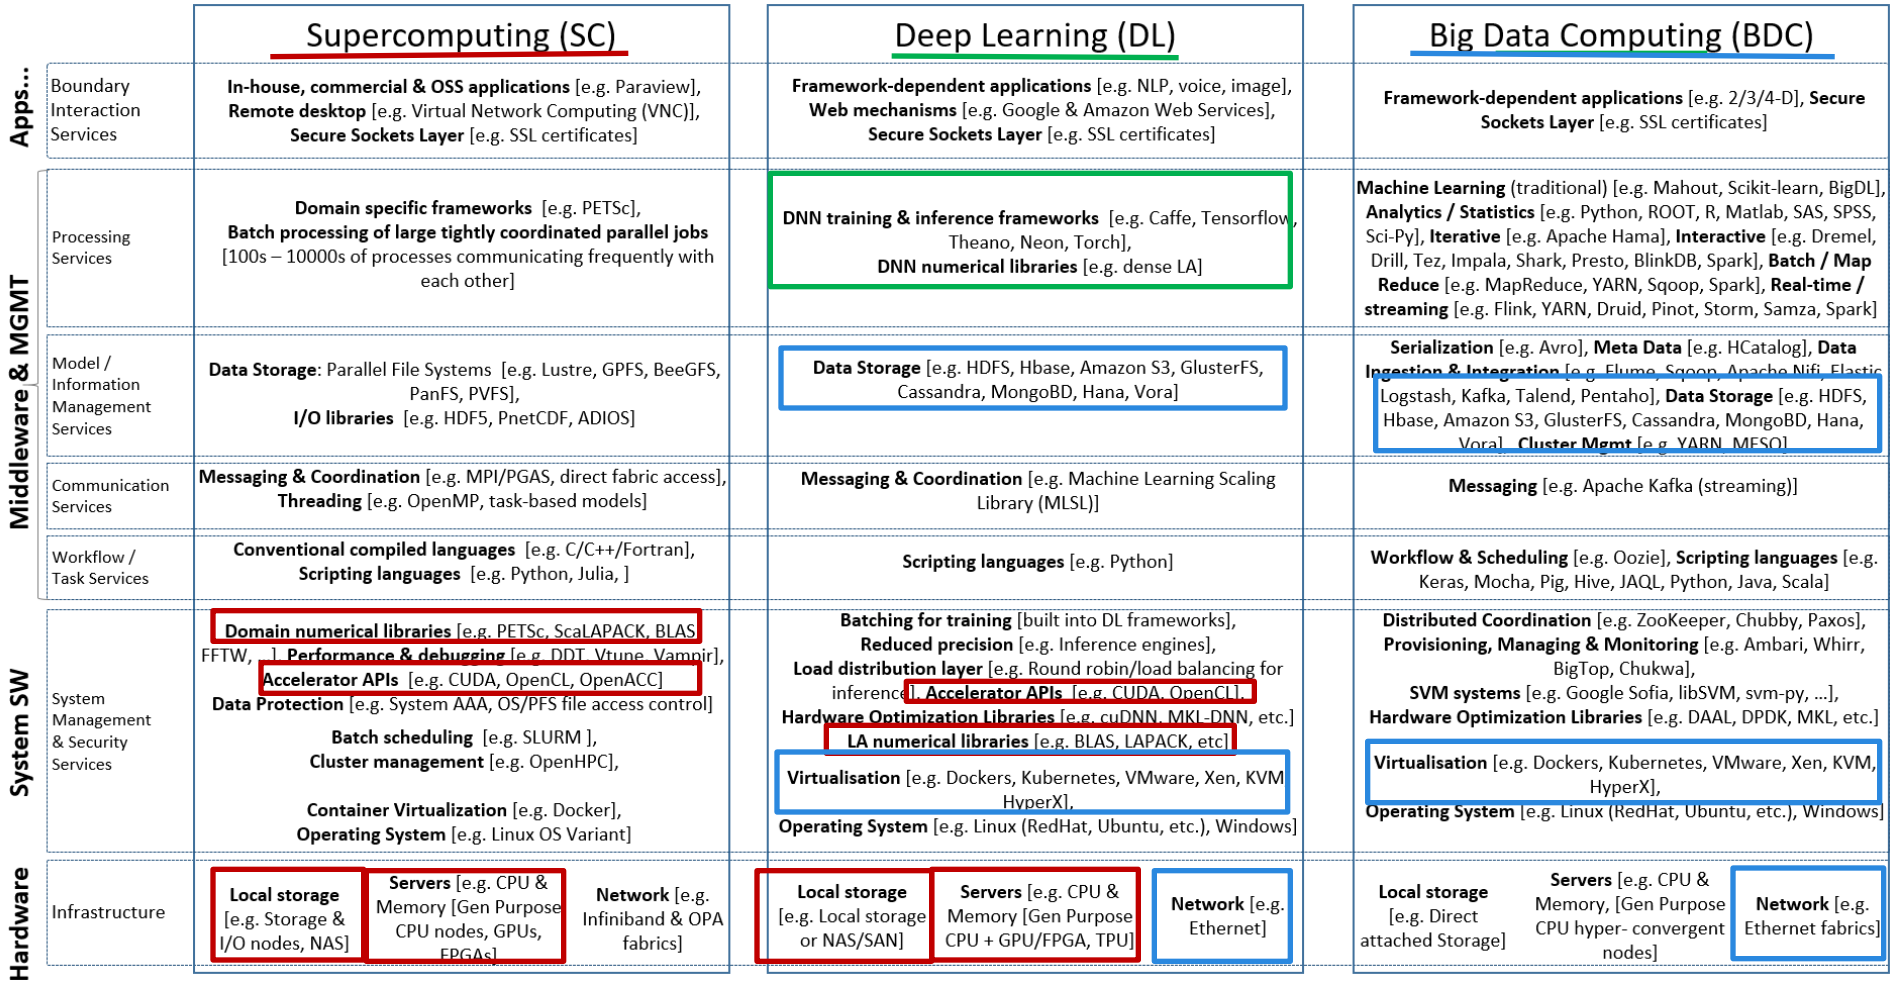

		Supercomputing (SC)	Deep Learning (DL)	Big Data Computing (BDC)
Apps…	Boundary Interaction Services	**In-house, commercial & OSS applications** [e.g. Paraview], **Remote desktop** [e.g. Virtual Network Computing (VNC)], **Secure Sockets Layer** [e.g. SSL certificates]	**Framework-dependent applications** [e.g. NLP, voice, image], **Web mechanisms** [e.g. Google & Amazon Web Services], **Secure Sockets Layer** [e.g. SSL certificates]	**Framework-dependent applications** [e.g. 2/3/4-D], **Secure Sockets Layer** [e.g. SSL certificates]
Middleware & MGMT	Processing Services	**Domain specific frameworks** [e.g. PETSc], **Batch processing of large tightly coordinated parallel jobs** [100s – 10000s of processes communicating frequently with each other]	**DNN training & inference frameworks** [e.g. Caffe, Tensorflow, Theano, Neon, Torch], **DNN numerical libraries** [e.g. dense LA]	**Machine Learning** (traditional) [e.g. Mahout, Scikit-learn, BigDL], **Analytics / Statistics** [e.g. Python, ROOT, R, Matlab, SAS, SPSS, Sci-Py], **Iterative** [e.g. Apache Hama], **Interactive** [e.g. Dremel, Drill, Tez, Impala, Shark, Presto, BlinkDB, Spark], **Batch / Map Reduce** [e.g. MapReduce, YARN, Sqoop, Spark], **Real-time / streaming** [e.g. Flink, YARN, Druid, Pinot, Storm, Samza, Spark]
	Model / Information Management Services	**Data Storage**: Parallel File Systems [e.g. Lustre, GPFS, BeeGFS, PanFS, PVFS], **I/O libraries** [e.g. HDF5, PnetCDF, ADIOS]	**Data Storage** [e.g. HDFS, Hbase, Amazon S3, GlusterFS, Cassandra, MongoBD, Hana, Vora]	**Serialization** [e.g. Avro], **Meta Data** [e.g. HCatalog], **Data Ingestion & Integration** [e.g. Flume, Sqoop, Apache Nifi, Elastic Logstash, Kafka, Talend, Pentaho], **Data Storage** [e.g. HDFS, Hbase, Amazon S3, GlusterFS, Cassandra, MongoBD, Hana, Vora], **Cluster Mgmt** [e.g. YARN, MESO]
	Communication Services	**Messaging & Coordination** [e.g. MPI/PGAS, direct fabric access], **Threading** [e.g. OpenMP, task-based models]	**Messaging & Coordination** [e.g. Machine Learning Scaling Library (MLSL)]	**Messaging** [e.g. Apache Kafka (streaming)]
	Workflow / Task Services	**Conventional compiled languages** [e.g. C/C++/Fortran], **Scripting languages** [e.g. Python, Julia,]	**Scripting languages** [e.g. Python]	**Workflow & Scheduling** [e.g. Oozie], **Scripting languages** [e.g. Keras, Mocha, Pig, Hive, JAQL, Python, Java, Scala]
System SW	System Management & Security Services	**Domain numerical libraries** [e.g. PETSc, ScaLAPACK, BLAS FFTW, …], **Performance & debugging** [e.g. DDT, Vtune, Vampir], **Accelerator APIs** [e.g. CUDA, OpenCL, OpenACC] **Data Protection** [e.g. System AAA, OS/PFS file access control] **Batch scheduling** [e.g. SLURM], **Cluster management** [e.g. OpenHPC], **Container Virtualization** [e.g. Docker], **Operating System** [e.g. Linux OS Variant]	**Batching for training** [built into DL frameworks], **Reduced precision** [e.g. Inference engines], **Load distribution layer** [e.g. Round robin/load balancing for inference], **Accelerator APIs** [e.g. CUDA, OpenCL], **Hardware Optimization Libraries** [e.g. cuDNN, MKL-DNN, etc.] **LA numerical libraries** [e.g. BLAS, LAPACK, etc] **Virtualisation** [e.g. Dockers, Kubernetes, VMware, Xen, KVM HyperX], **Operating System** [e.g. Linux (RedHat, Ubuntu, etc.), Windows]	**Distributed Coordination** [e.g. ZooKeeper, Chubby, Paxos], **Provisioning, Managing & Monitoring** [e.g. Ambari, Whirr, BigTop, Chukwa], **SVM systems** [e.g. Google Sofia, libSVM, svm-py, …], **Hardware Optimization Libraries** [e.g. DAAL, DPDK, MKL, etc.] **Virtualisation** [e.g. Dockers, Kubernetes, VMware, Xen, KVM, HyperX], **Operating System** [e.g. Linux (RedHat, Ubuntu, etc.), Windows]
Hardware	Infrastructure	**Local storage** [e.g. Storage & I/O nodes, NAS] **Servers** [e.g. CPU & Memory [Gen Purpose CPU nodes, GPUs, FPGAs] **Network** [e.g. Infiniband & OPA fabrics]	**Local storage** [e.g. Local storage or NAS/SAN] **Servers** [e.g. CPU & Memory [Gen Purpose CPU + GPU/FPGA, TPU] **Network** [e.g. Ethernet]	**Local storage** [e.g. Direct attached Storage] **Servers** [e.g. CPU & Memory, [Gen Purpose CPU hyper- convergent nodes] **Network** [e.g. Ethernet fabrics]

Disaggregated Stack Profiles (Note: some applications serve multiple functions, for illustration simplicity they were assigned to their dominant function)

An example of hybrid architecture

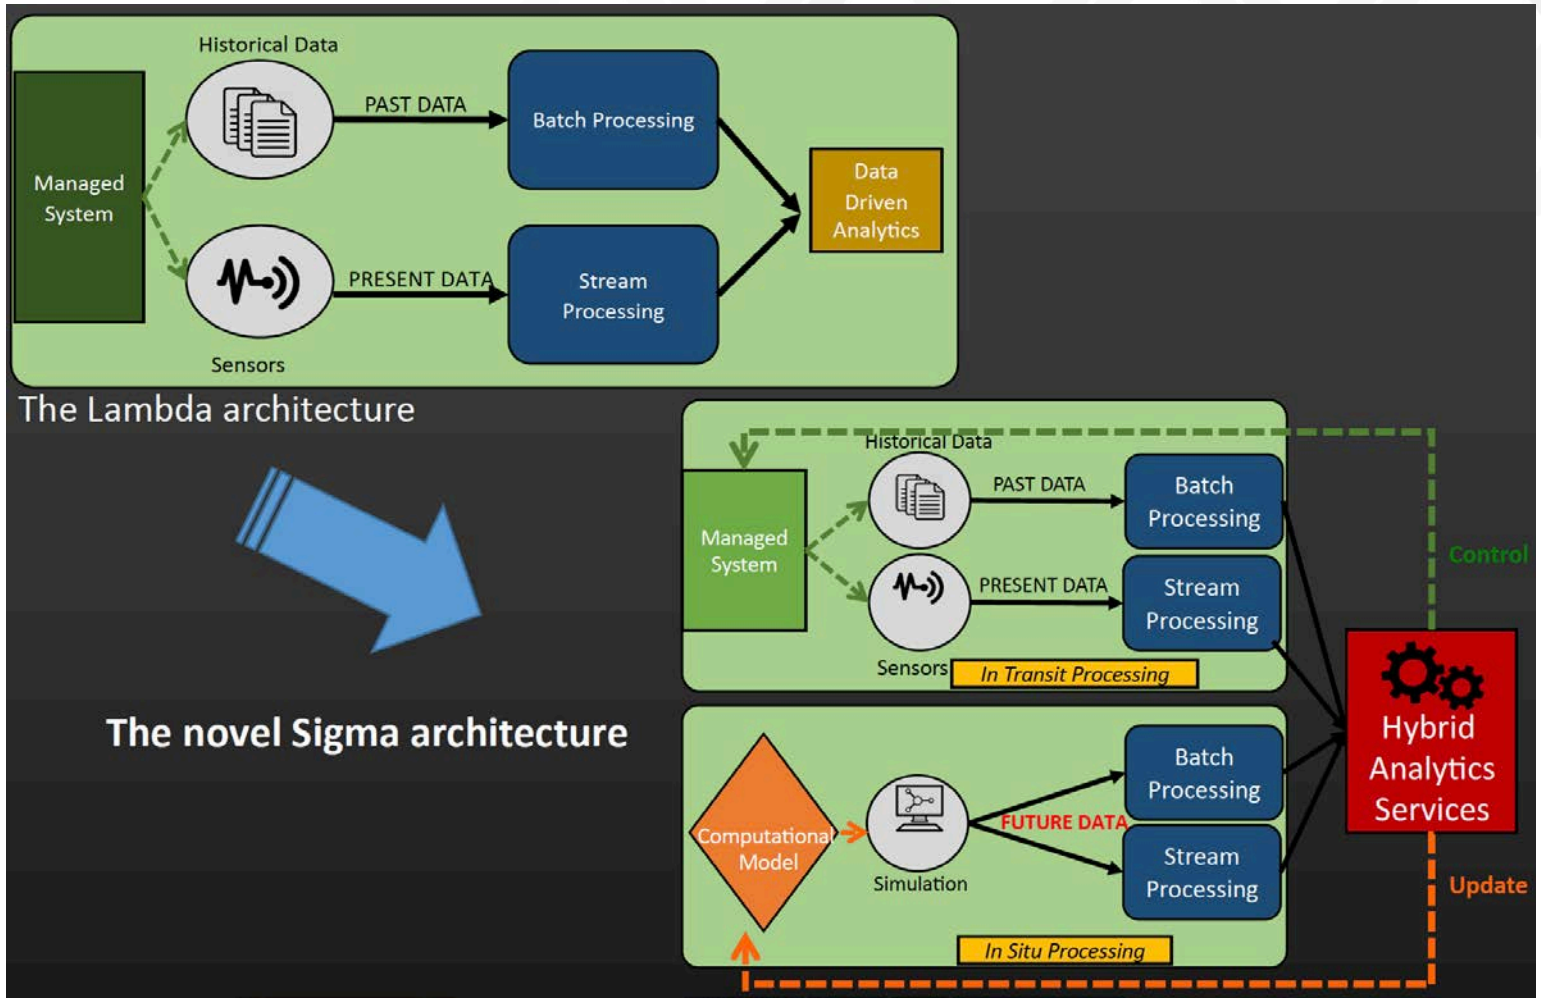

Source: Gabriel Antoniu, Alexandru Costan, Maria S. Pérez and Nenad Stojanovic, *The Sigma Data Processing Architecture: Leveraging Future Data for Extreme-Scale Data Analytics to Enable High-Precision Decisions*, BDEC2 Meeting, Nov. 2018, Indiana, USA

Use Case: Autonomous Cars

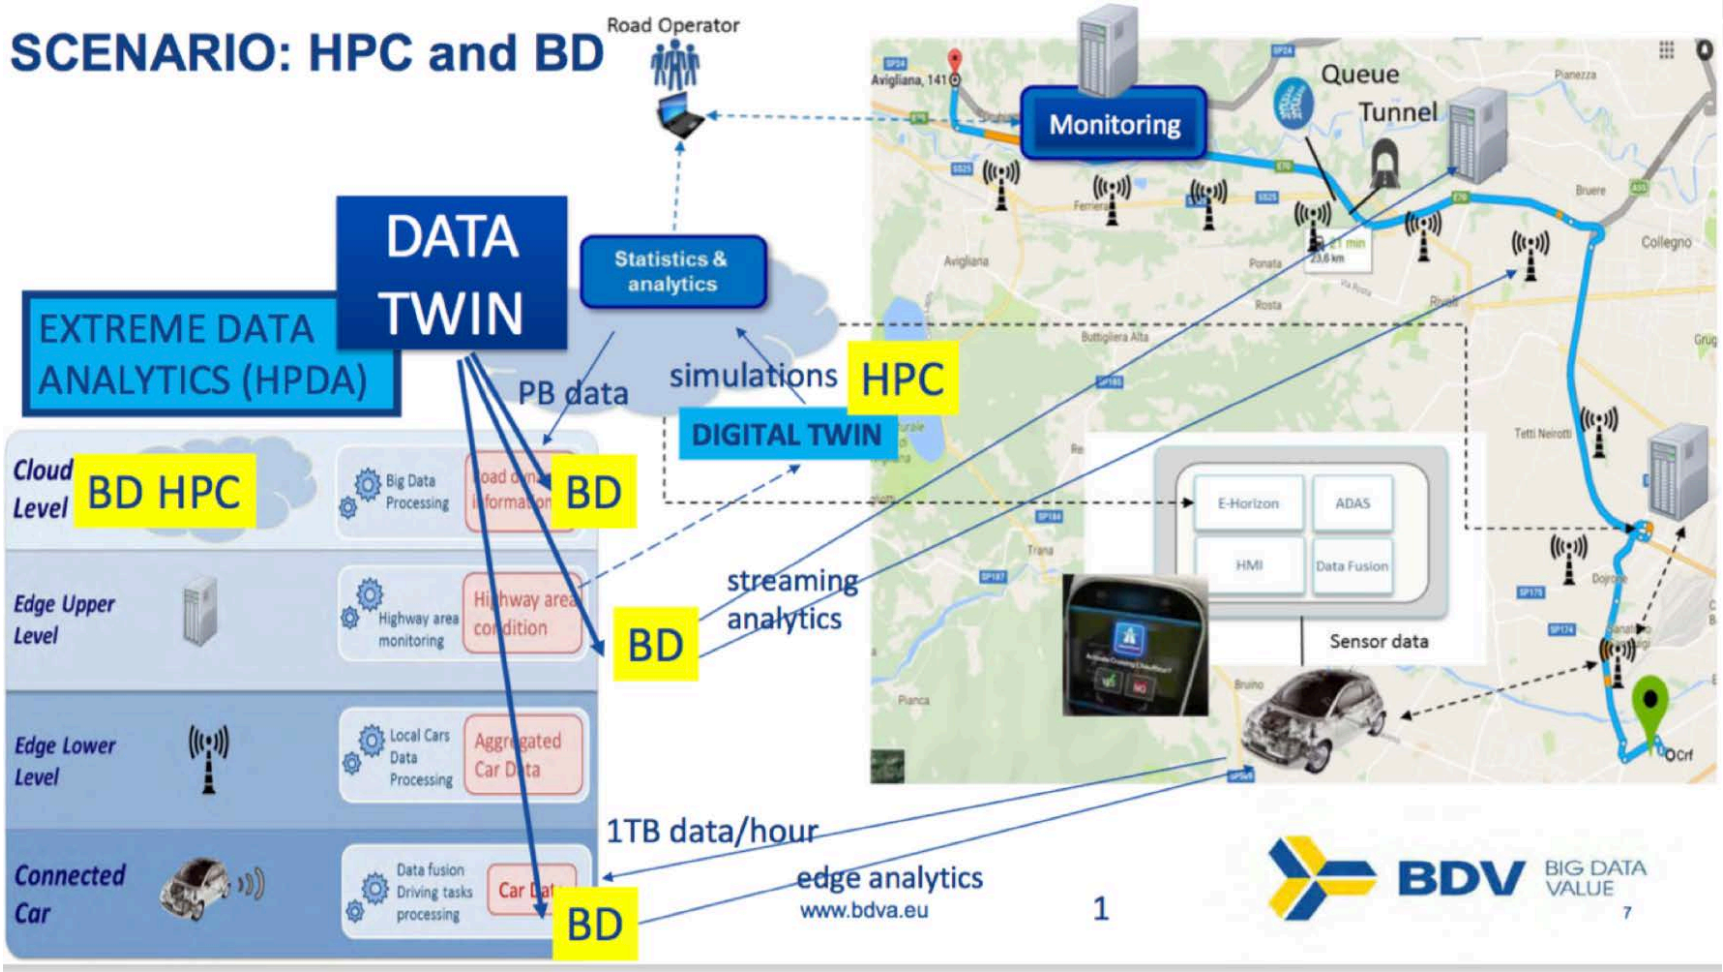

Credits: Nissatech

System of systems

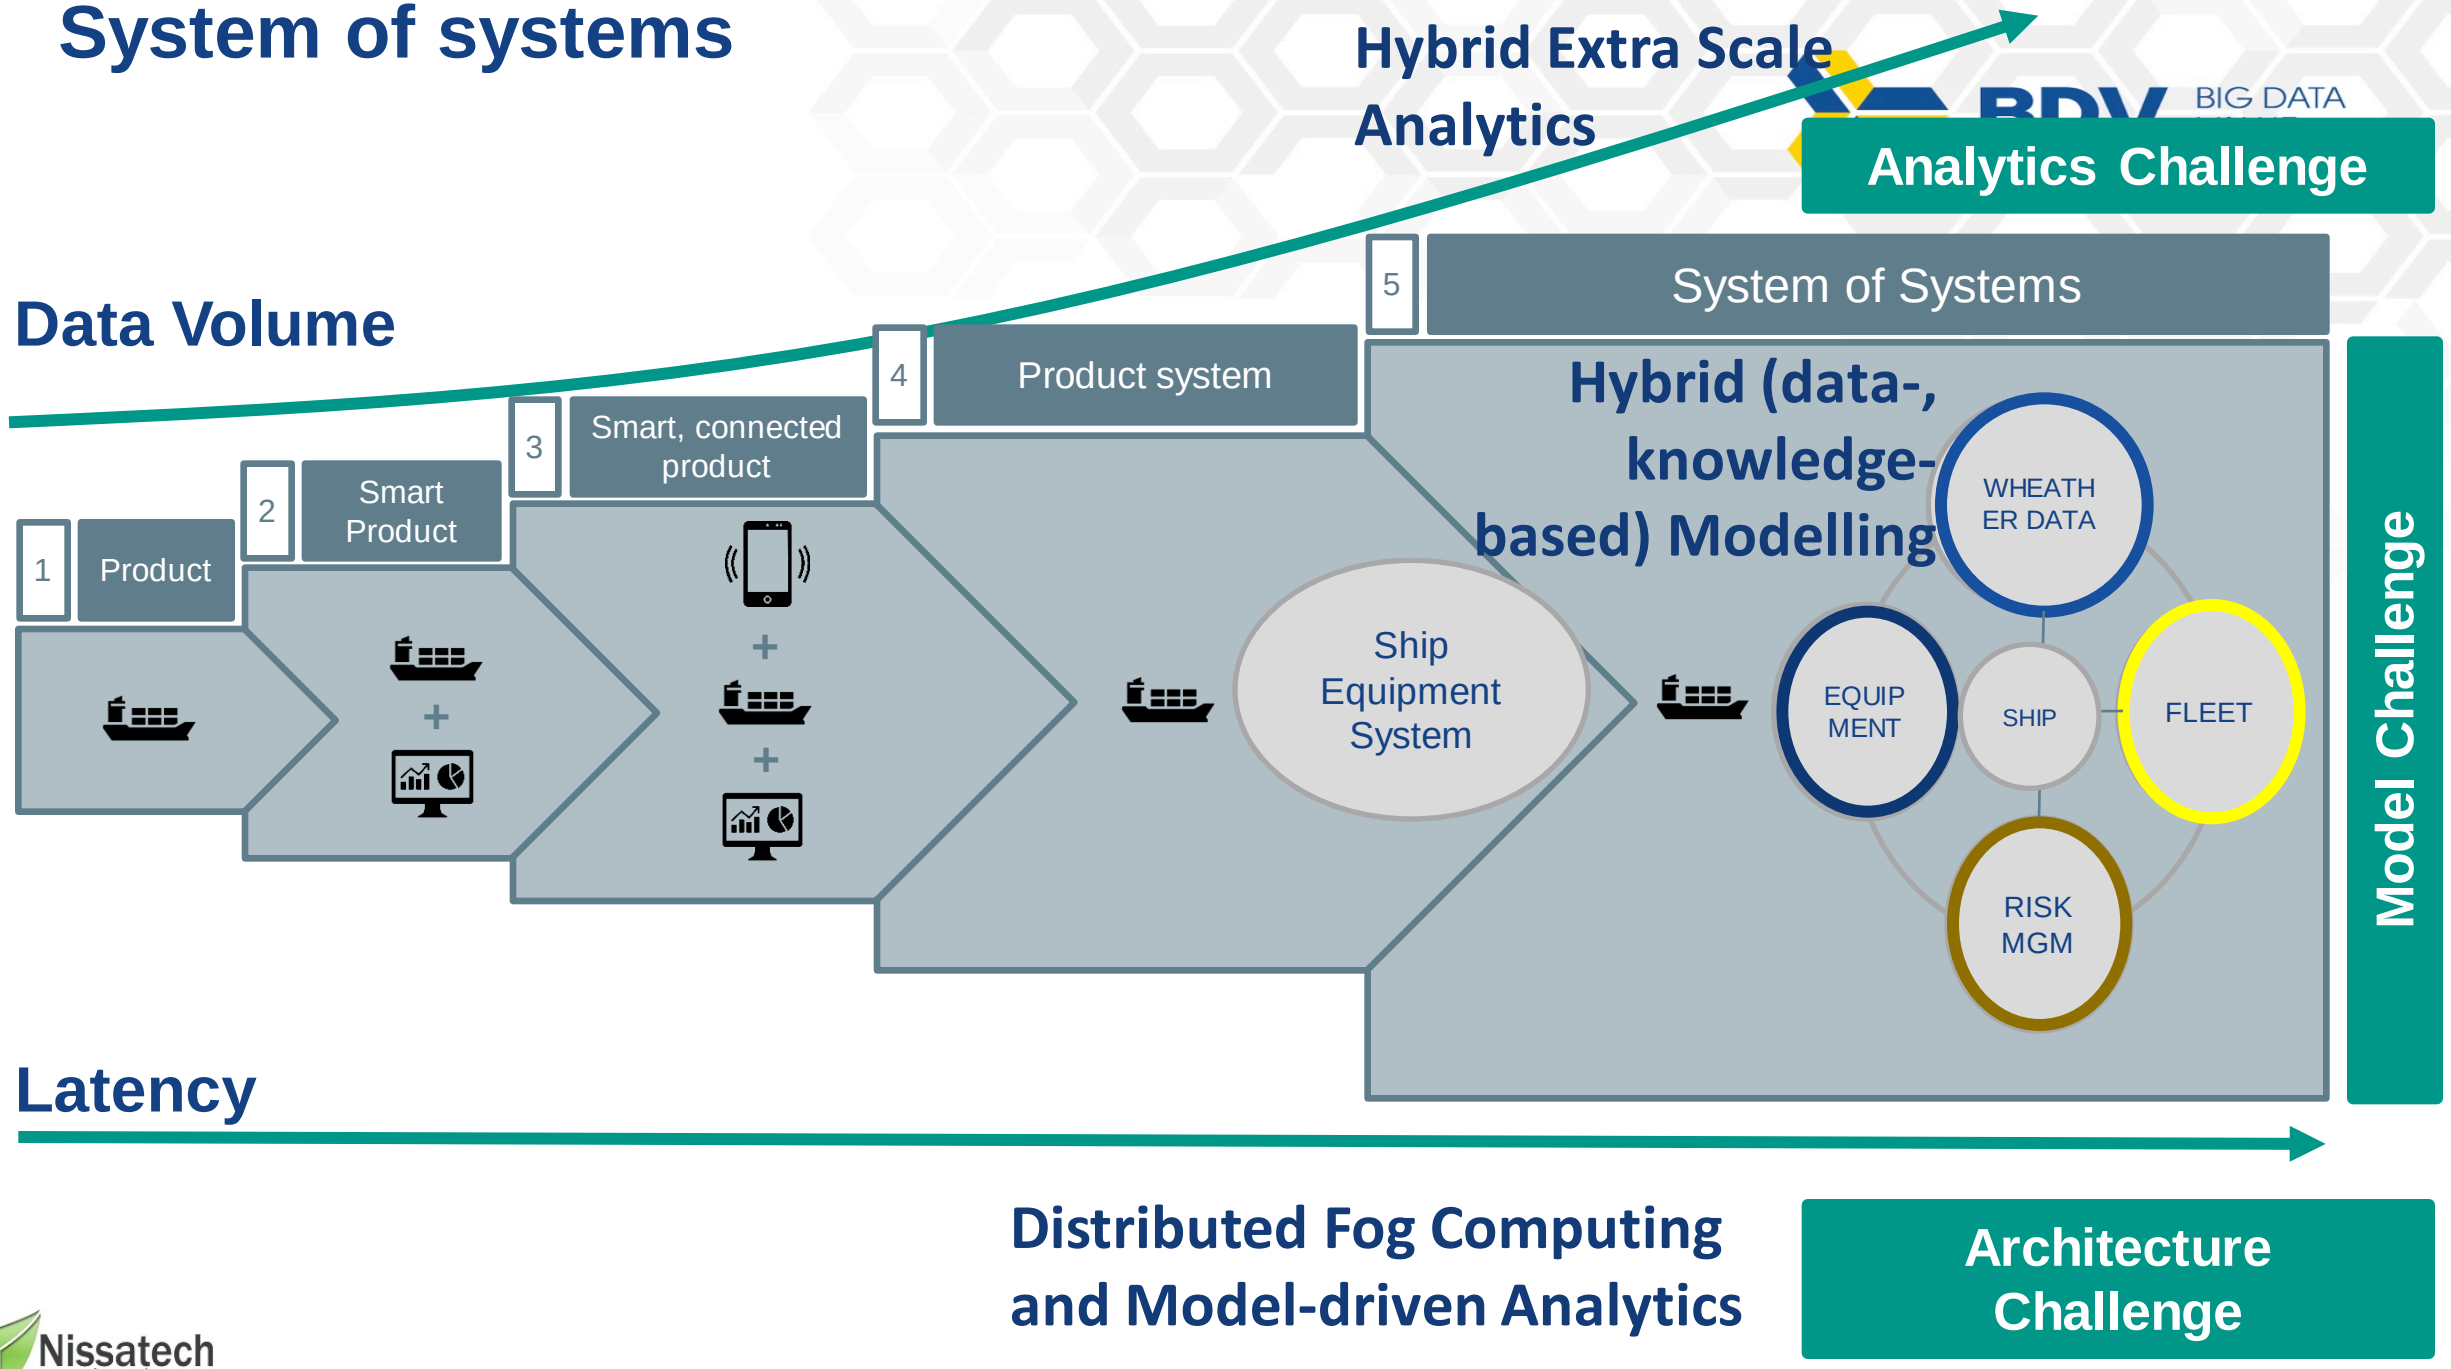

Nissatech innovation centre

‘5Ps Master Algorithm’…

…multi-objective optimization at intersection of AI, Big Data & HPC

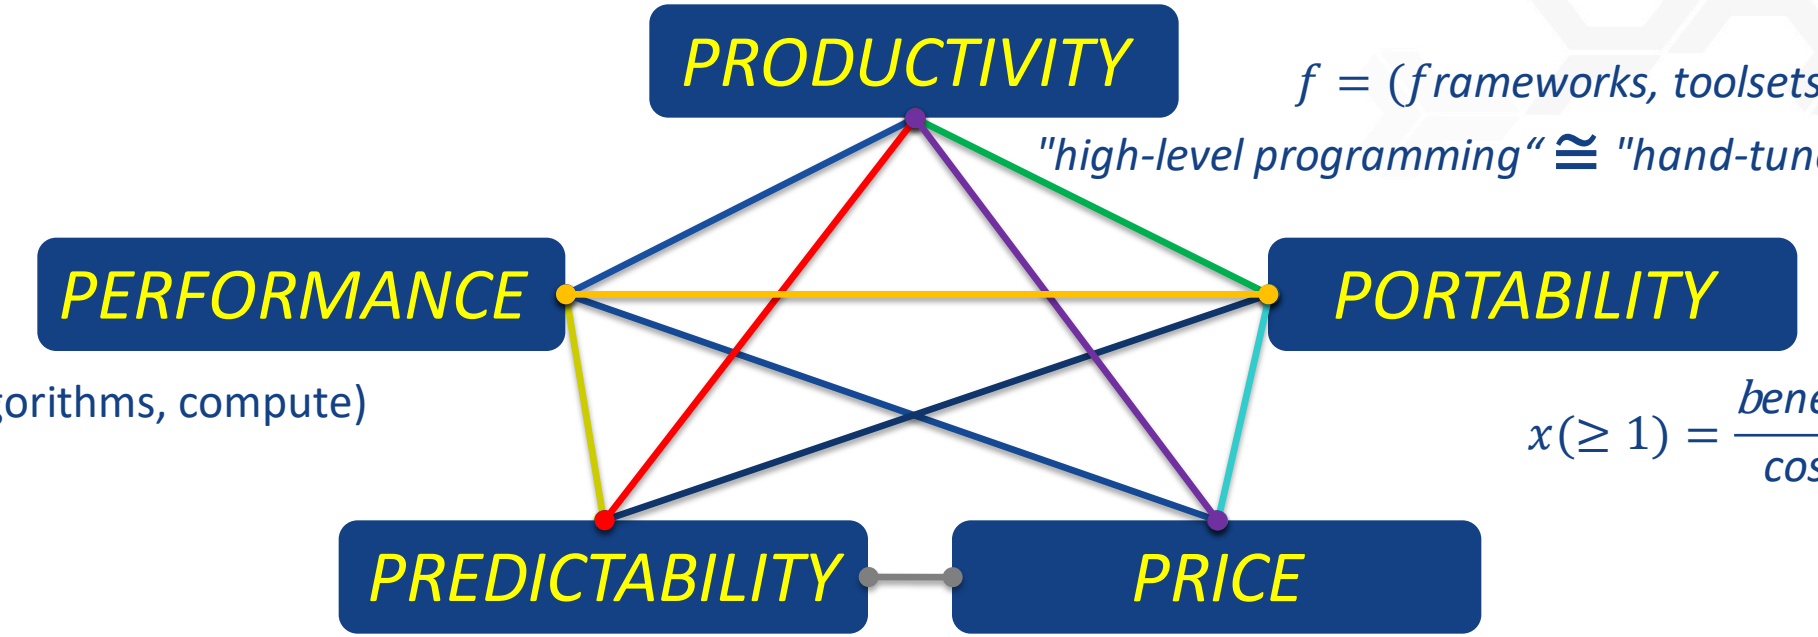

$f = (f\,\text{frameworks, toolsets, ...})$

"high-level programming“ ≅ "hand-tuned machine code"

$f = (\text{data, algorithms, compute})$

$$x(\geq 1) = \frac{\textit{benefits of transfer}}{\textit{costs of transfer}}$$

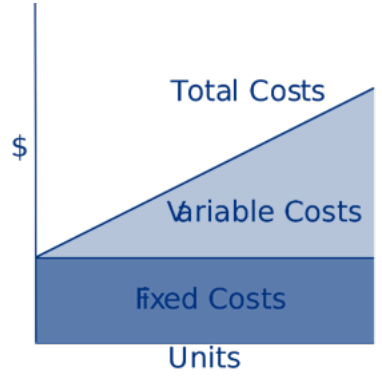

www.bdva.eu

Alignment of Research Agendas

BDVA SRIA

- Data Novel architectures for enabling new types of big data workloads (hybrid Big Data and HPC architecture)
- High Performance Data Analytics (HPDA)

ETP4HPC SRA

- Heterogeneous HPC data processing systems (HPC and Big Data hybrids), which are flexible in allowing breakthrough simulations and data analytics at the same time
- Data Centric Memory Hierarchies/Architectures
- Research in algorithms that trade computation with data accesses
- Programming models and languages for data centric computing

EuroHPC

- Design and integration of simulation, data analytics and other AI capabilities

HiPEAC Vision

- Towards flexible and efficient Exascale software couplers (direct or not, exchange of big data)
- In-situ extreme data processing and better science through I/O avoidance in HPC systems
- Declarative processing frameworks for big data analytics, extreme data fusion e.g. identification of turbulent flow features from massively parallel Exaflops and Exabytes simulations

High-Performance and Embedded Architecture and Compilation

HiPEAC Vision 2019

for Computing in 2025

Marc Duranton, Koen De Bosschere, **Christian Gamrat**, Harm Munk, Emre Ozer, Tullio Vardanega, Olivier Zendra

Common priorities of HPC, Big Data and HiPEAC for post-H2020 era, ICT 2018, Vienna

HiPEAC

=

High-**P**erformance and **E**mbedded
Architecture and **C**ompilation

HiPEAC's mission is to steer and increase the European research in the area of high-performance and embedded computing systems,

And stimulate cooperation between:

a) academia and industry

b) computer architects and tool builders.

ENABLING EDGE INTELLIGENCE

C²PS: COGNITIVE (CYBERNETIC* AND PHYSICAL) SYSTEMS

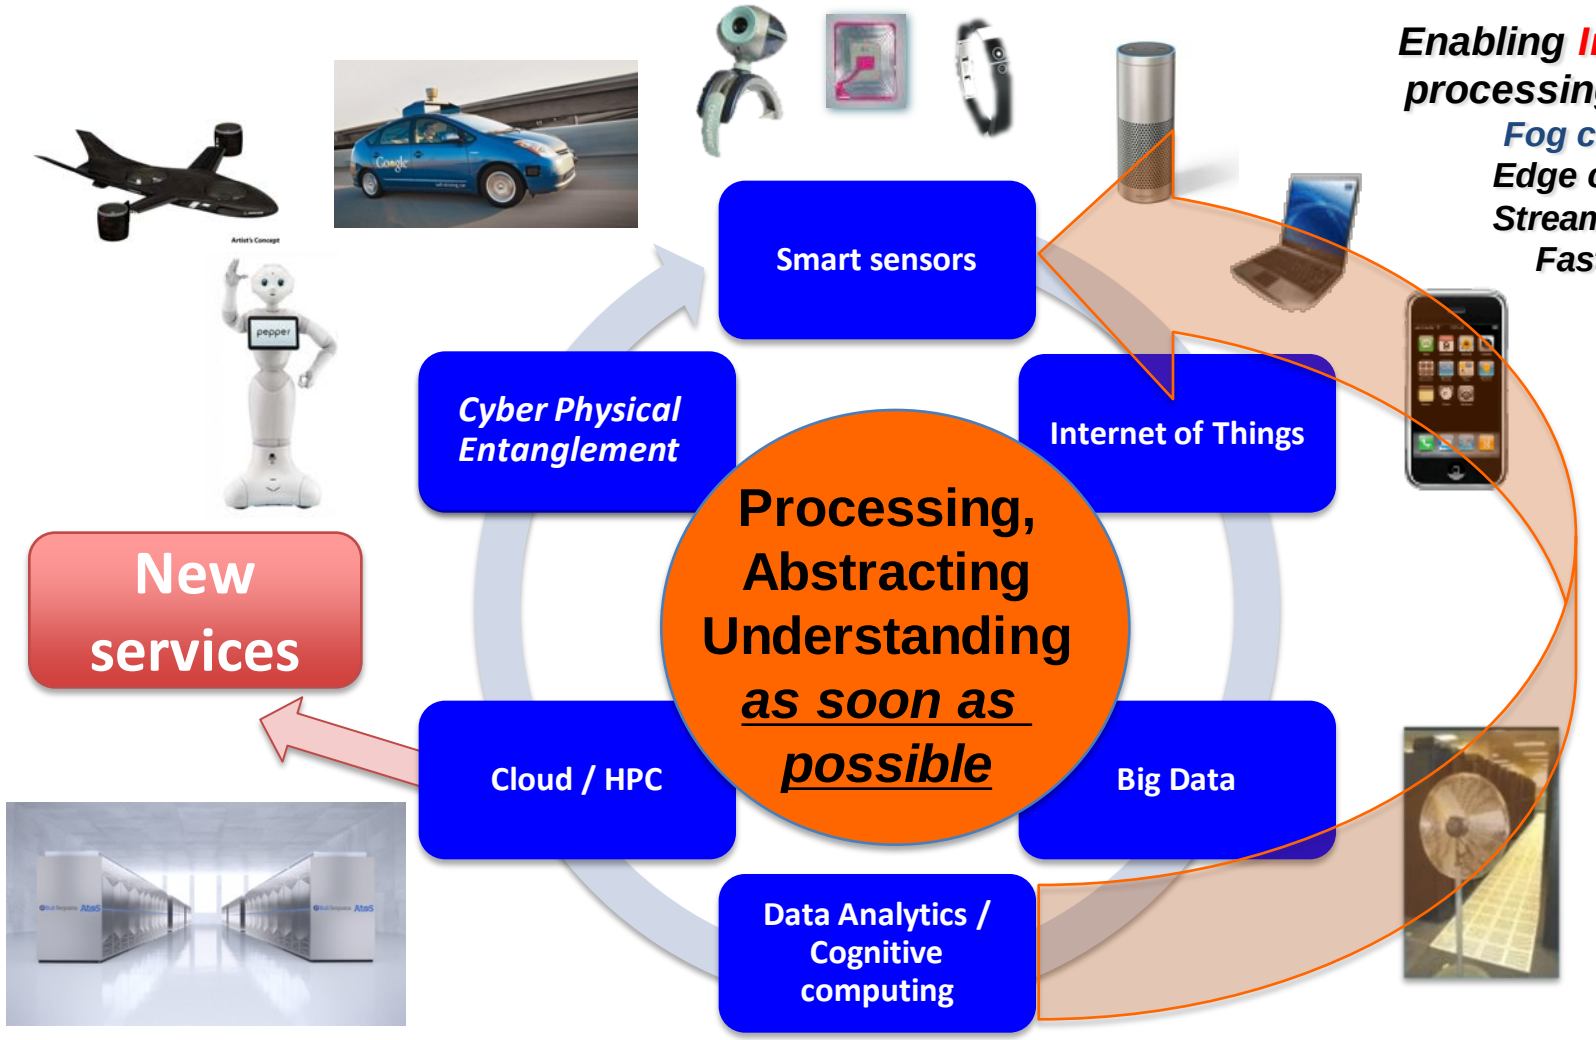

Enabling Intelligent data processing at the edge:
Fog computing
Edge computing
Stream analytics
Fast data…

True collaboration between edge devices and the HPC/cloud ensuring:

- Data security / Privacy
- Lower bandwidth
- Better use of HPC/cloud

Transforming **data** into ***information*** ***as early as possible***

** As defined by Norbert Wiener: how humans, animals and machines control and communicate with each other.*

Embedded intelligence needs local high-end computing

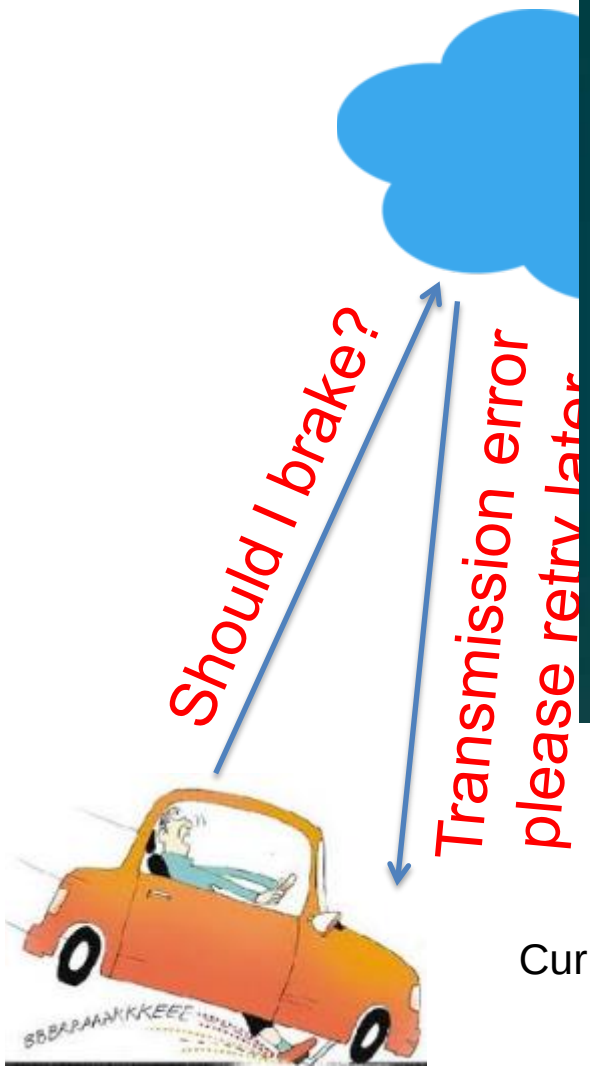

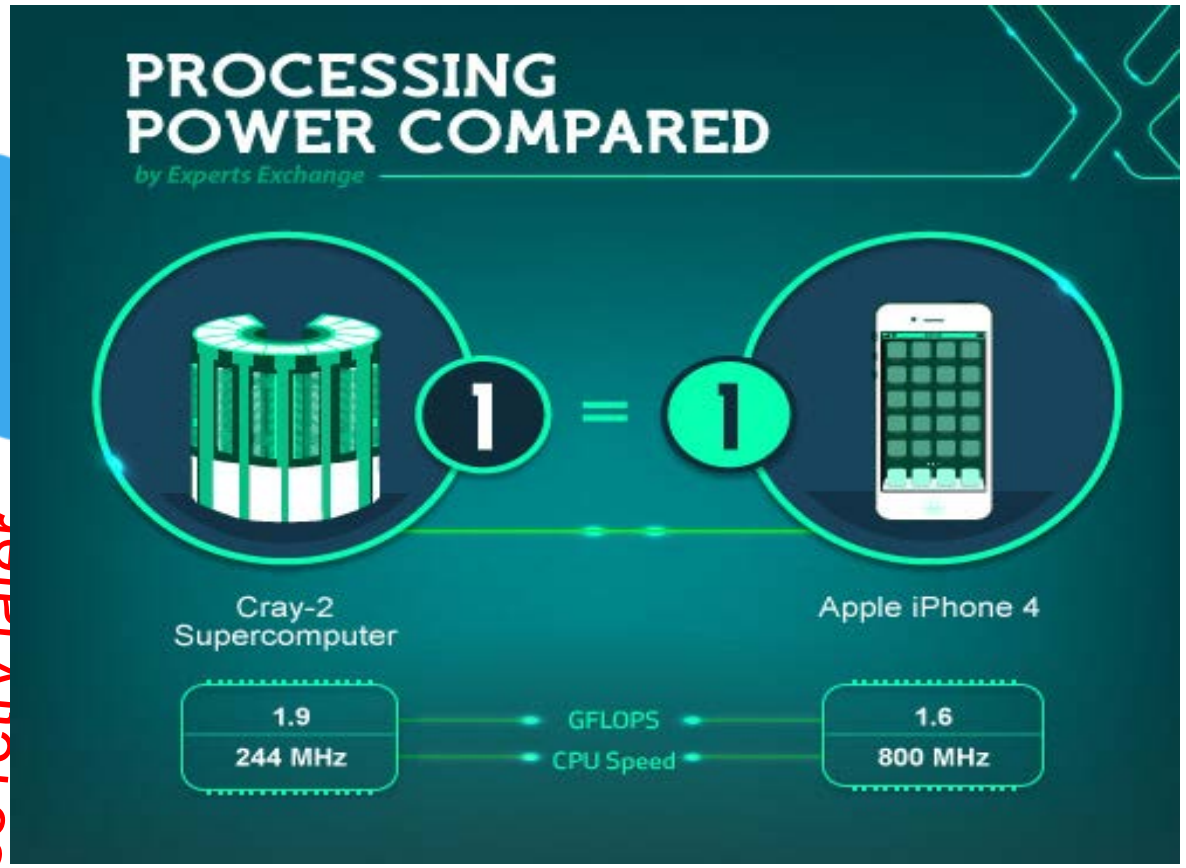

25 years

1985 → 2010

Current embedded systems are HPC systems of few decades ago

THE 3 PILLARDS OF FUTURE HPC

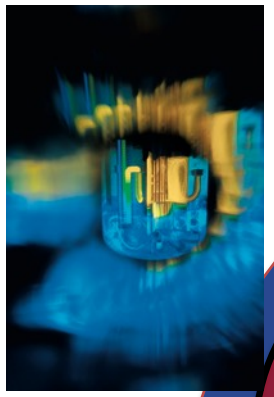

SIMULATION = MODEL FOR PREDICTION

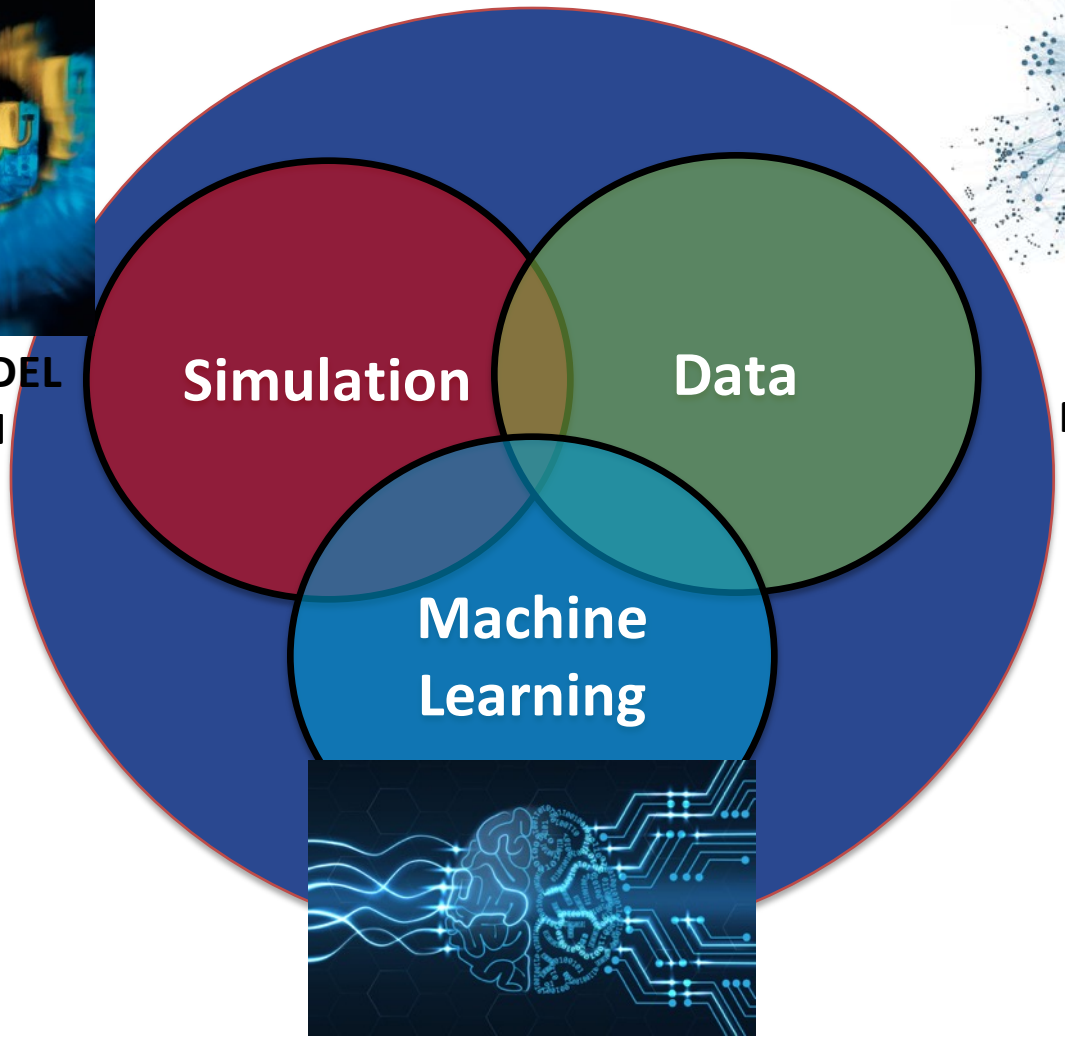

DATA IS KEY TO MACHINE LEARNING

Prediction + Machine Learning
Key to Cognitive CPS requirements

FUSING HARDWARE TECHNOLOGIES AND ARCHITECTURES

At the hardware level, the good old Von Neumann/ CMOS partnership can act as a computing substrate

- Acting as coordination / communication node
- Allowing Hardware / Software integration

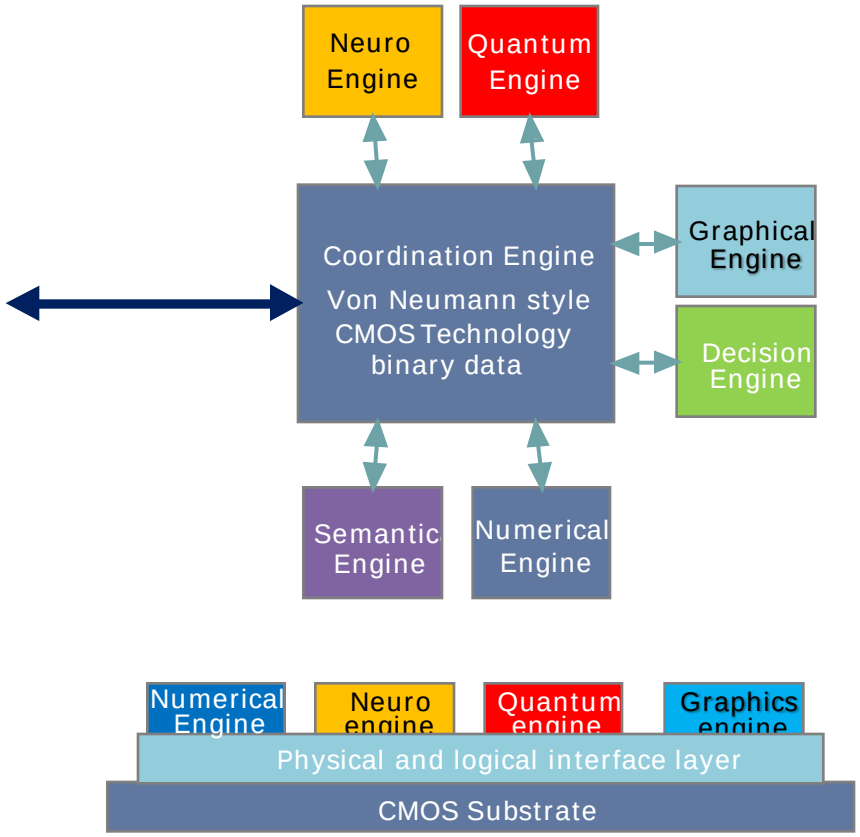

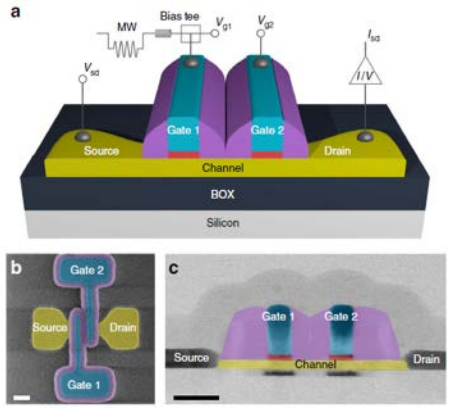

Qubits on Silicon

Maurand et al, *Nature Com.*, Jul. 2016.

NVM Synapses on Silicon

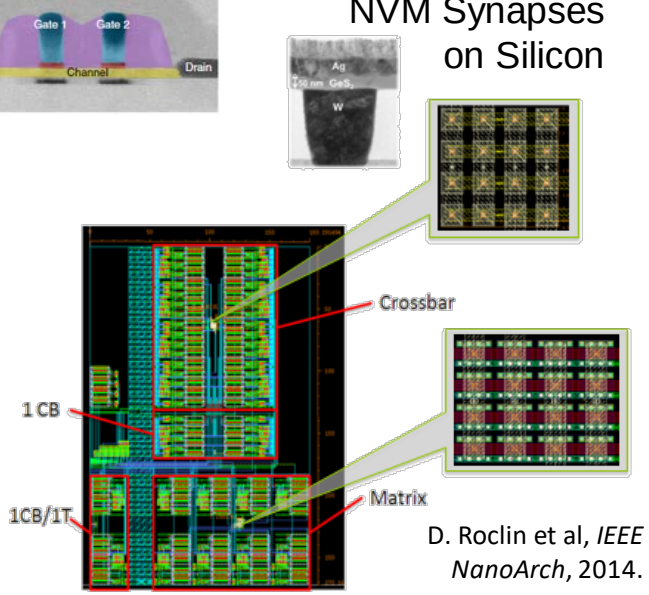

D. Roclin et al, *IEEE NanoArch*, 2014.

COORDINATE SOFTWARE PARADIGMS

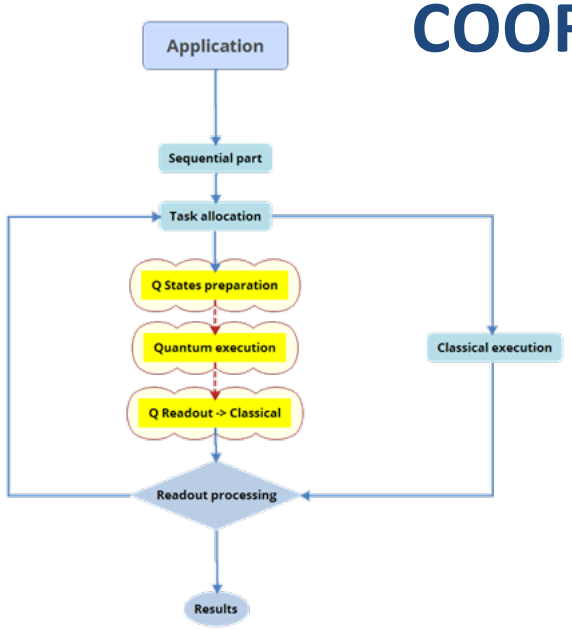

- **A sequential program (running on the coordinator) distributes tasks to engines)**

- **Distributed to computing engines**
  - Instruction streams
  - Machine Learning
  - Quantum Engine

- **Automate tasks generation and distribution**
  - To manage increased complexity
  - Provide cognitive functions

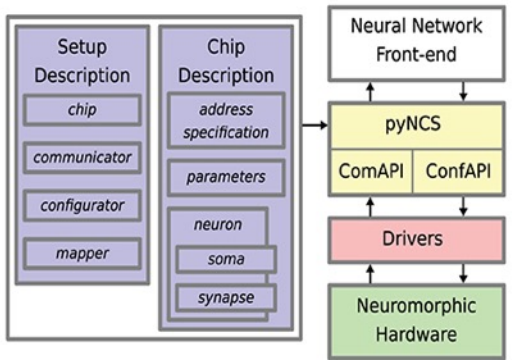

Stefanini et al. Frontiers Neuroinformatics *August 2014*

```
mycirc :: Qubit -> Qubit -> Circ (Qubit, Qubit)
mycirc a b = do
  a <- hadamard a
  b <- hadamard b
  (a,b) <- controlled_not a b
  return (a,b)
```

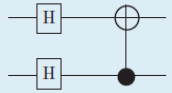

Valiron et al., "Programming the Quantum Future," *Commun. ACM*, vol. 58, no. 8, pp. 52–61, Jul. 2015.

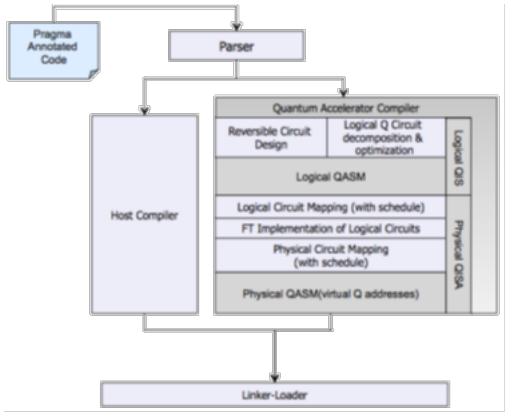

Fu et al. Proc. ACM, May 2016

HIGHLIGHTS OF THE HIPEAC VISION 2019

- **Get looking for CMOS alternatives**
  - These technologies are unlikely to supplant CMOS, but instead will complement it.
- **Develop alternative to von Neumann architectures**
  - Processing in memory, neuro-inspired computing, etc… and make them work together.
- **Accelerate, accelerate, accelerate and automate**
  - Specialize hardware for important application domains.
  - It will only be economically viable if it is automated by intelligent tools and frameworks.

HIGHLIGHTS OF THE HIPEAC VISION 2019

- **Treat ICT a continuum, from the edge to the cloud**
  - From microcontrollers, to concentrators, to micro-servers, to cloud and high-performance computing.
  - Interoperability is key; systems need to collaborate to give the best service to users.
  - We need dynamic devices which can adapt intelligently
- **Lead on the use of collective data**
  - Europe should develop the ethical use of state-owned, collective or domain data.
- **Build computers you can trust**
  - They need to be secure and safe, not harming people when they interact with them.
  - They also need to be reliable despite being increasingly complex.
  - They should be explainable enough to build trust.

HIGHLIGHTS OF THE HIPEAC VISION 2019

- **Become a leader in energy-efficient, sustainable electronics**
  - By the recycling of computing devices and modularity, prolonging the life of ICT systems.
  - Increasing the longevity of electronic systems, through certification and virtualization, modularity, specific supervision, etc.
  - Computing should also be used to find solutions to the sustainability crisis we are facing.

- **Develop a robust digital ethics framework**
  - We should start thinking about whether everything that can be done should, in fact, be done.
  - Ethics will become an important part of ICF

SUMMARY: THE HIPEAC VISION 2019

- Better and Smarter Computers: HPC is on board!
  - CMOS alternatives
  - Accelerate and Automate
- Big data from the edge to the cloud
  - Treat ICT a continuum
  - Lead on the use of collective data
- Trustable and Sustainable ICT
  - Build computers you can trust
  - Develop a robust digital ethics framework

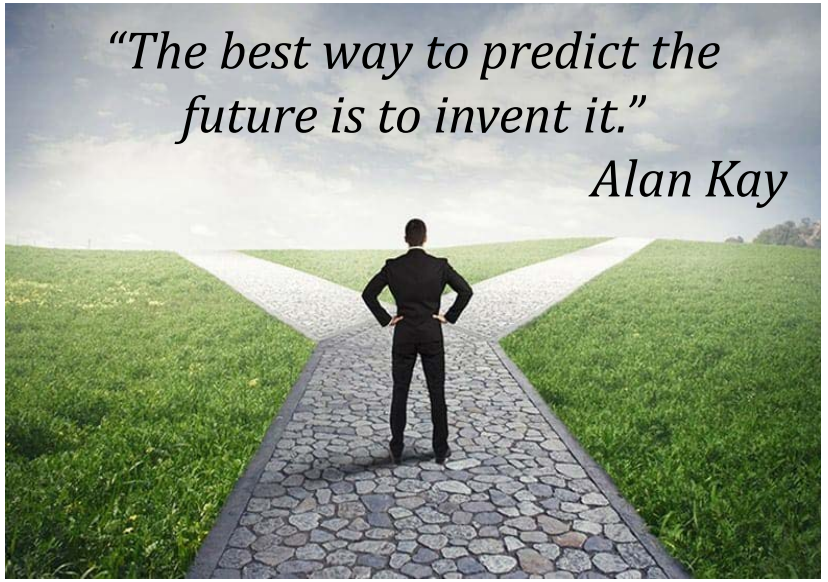

Together:

Thank you

2008 2009 2011 2013 2015 2017

https://www.hipeac.net/vision/

The HiPEAC project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement number 779656.

What is next?

Questions?
Comments?