



OVERVIEW OF GERMAN HPC RESOURCES & FACILITIES – HELMHOLTZ

PROF. DR. – ING. MORRIS RIEDEL, JUELICH SUPERCOMPUTING CENTRE / UNIVERSITY OF ICELAND
HEAD OF CROSS-SECTIONAL TEAM DEEP LEARNING & RDA CO-CHAIR INTEREST GROUP BIG DATA
16TH JANUARY, CIFAR-HELMHOLTZ WORKSHOP, MARS DISCOVERY DISTRICT, TORONTO, CANADA



UNIVERSITY OF ICELAND
SCHOOL OF ENGINEERING AND NATURAL SCIENCES
FACULTY OF INDUSTRIAL ENGINEERING,
MECHANICAL ENGINEERING AND COMPUTER SCIENCE



HPC & DATA SCIENCE: A FIELD OF CONSTANT EVOLUTION

Perspective: Floating Point Operations per one second (FLOPS or FLOP/s)

1.000.000 FLOP/s
~1984



- 1 GigaFlop/s = 10^9 FLOPS
- 1 TeraFlop/s = 10^{12} FLOPS
- 1 PetaFlop/s = 10^{15} FLOPS
- 1 ExaFlop/s = 10^{18} FLOPS

© Photograph by Rama, Wikimedia Commons

1.000.000.000.000.000 FLOP/s
~295.000 cores ~2009 (JUGENE)



>5.900.000.000.000.000 FLOP/s
~ 500.000 cores
~2013 → end of service in 2018

EUROPEAN UNION & COMMISSION PLANS

Supporting Artificial Intelligence & Supercomputers – Relevance of HPC & AI in Europe

“By supporting strategic projects in frontline areas such as artificial intelligence, supercomputers, cybersecurity or industrial digitisation, and investing in digital skills, the new programme will help to complete the Digital Single Market, a key priority of the Union.”

[11] COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE EUROPEAN COUNCIL, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS, EC, 2018, 2nd May 2018

You Retweeted

European Commission @EU_Commission · Apr 25

How can Europe be at the forefront of artificial intelligence, #dataeconomy, and digital healthcare?

@Ansip_EU @GabrielMariya @EBienkowskaEU @Moedas
#DigitalSingleMarket #AI



45:15 4K viewers

Digital Single Market proposals: artificial intelligence, data econ...
European Commission @EU_Commission

8 109 124

16th January 2019

Page 3

Mariya Gabriel @GabrielMariya

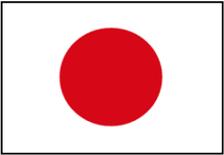
Follow

We are proud of you @fzj_jsc for the #firstclass #supercomputing facility you run. It is by efforts like yours that we reaffirm #EUaddedvalue and leadership in groundbreaking technologies. It is by cooperating that we will achieve our objectives for #EU leader in #HPC



WORLDWIDE HPC ROADMAP TO EXASCALE

Coordinated Activities



- **Flagship 2020:** Post-K

- 2020
- Fujitsu+ARM



- **TaihuLight**

- 2020
- Lenovo+ShenWei/FeiTeng CPU



- **CORAL:** 2 Exascale machines

- 2023
- Intel+Cray and IBM+NVIDIA



- **H2020 + IPCEI + EuroHPC + EU Cloud initiative**

- 2022
- **Technology and design not fixed yet**



EUROPEAN HPC STRATEGY

Coordinated Activities



- **PRACE**

- Computing infrastructures for European Users
- Operation, support and training



- **ETP4HPC**

- Industry-driven Roadmap (SRA)
- Pushing for Extreme Scale Demonstrators



- **EuroHPC**

- EU-based technology development, eg. processor
- Pushing for EU-made machine by 2022



- **H2020**

- Technology (HW+SW) development in Co-design
- FETHPC + Flagships + Quantum Computing



GERMAN GAUSS CENTRE FOR SUPERCOMPUTING

Alliance of the three national supercomputing centres HLRS (Stuttgart), JSC (Juelich) & LRZ (Munich)



16th January 2019

■ Supercomputer JUWELS @ JSC

- Juelich Wizard for European Leadership Science (JUWELS)
- Cluster architecture based on commodity multi-core CPUs
- 2,550 compute nodes: two Intel Xeon 24-core Skylake CPUs & 48 accelerated compute nodes (4 NVIDIA Volta GPUs)

■ Supercomputer SuperMUC @ LRZ

- 155,000 cores (only system in the Top20 Top500 w/o GPUs)

■ Supercomputer Hazel Hen

- 185,088 compute cores (no GPUs)

■ GCS represents Germany in PRACE



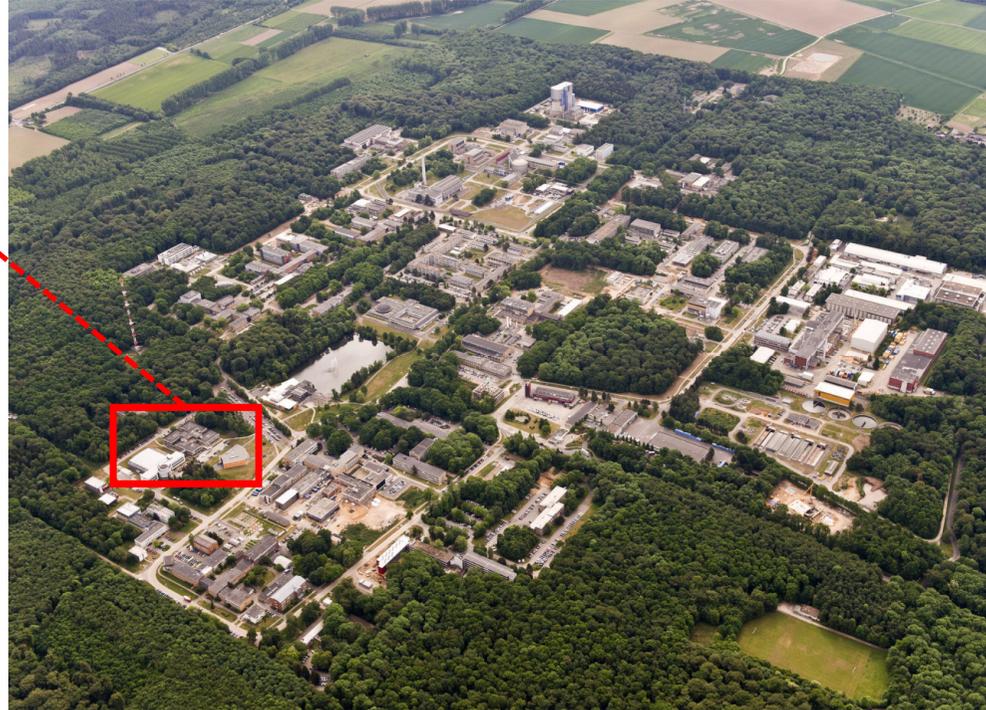
Page 6

FORSCHUNGSZENTRUM JUELICH (FZJ)

Multi-Disciplinary Research Centre of the Helmholtz Association in Germany

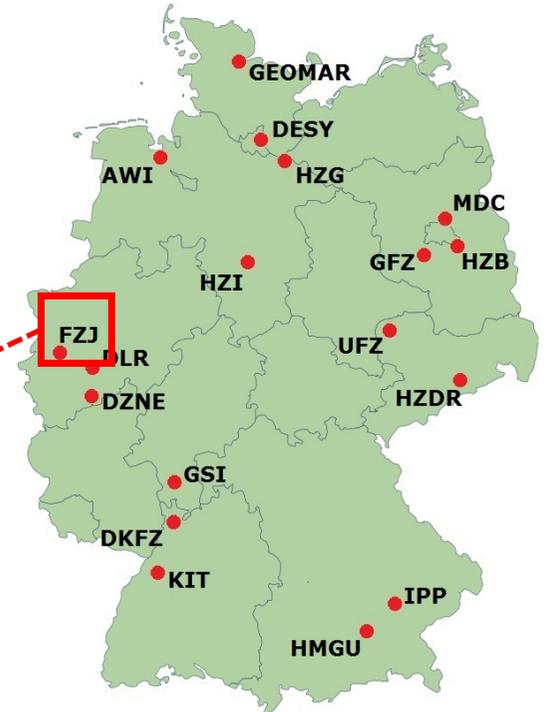


(Juelich Supercomputing Centre known as JSC)



Selected Facts

- One of EU largest inter-disciplinary research centres (~5000 employees)
- Special expertise in physics, materials science, nanotechnology, neuroscience and medicine & **information technology (HPC & Data)**



HELMHOLTZ
RESEARCH FOR GRAND CHALLENGES

[1] Helmholtz Association Web Page

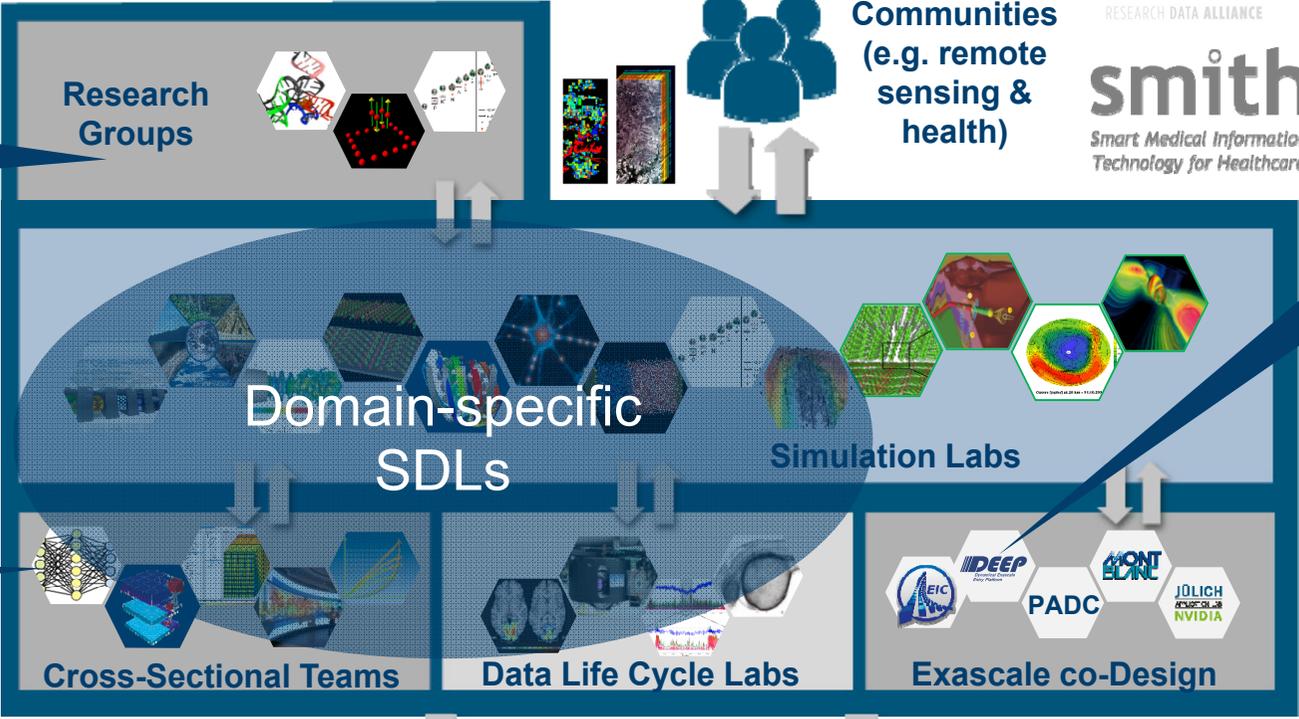
JUELICH SUPERCOMPUTING CENTRE (JSC) OF FZJ

Simulation & Data Labs (SDL) using High Performance Computing (HPC)

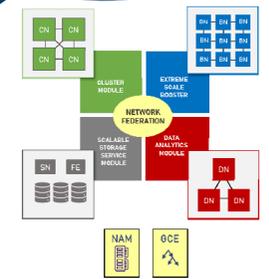


Smart Data Innovation Lab

Research Group High Productivity Data Processing



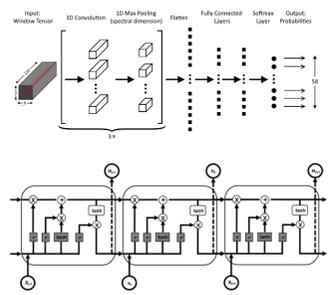
DEEP-EST EU PROJECT



UNIVERSITÄT ISLAND
UNIVERSITY OF ICELAND
SCHOOL OF ENGINEERING AND NATURAL SCIENCES
FACULTY OF INDUSTRIAL ENGINEERING,
MECHANICAL ENGINEERING AND COMPUTER SCIENCE

JÜLICH
Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

Cross-Sectional Team Deep Learning



ON4OFF
Landsvirkjun National Power Company of Iceland
SOCCERWATCH BETA

Modular Supercomputer JURECA

Facilities

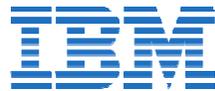
Modular Supercomputer JUWELS

Industry Relations Team

JSC & CO-DESIGN APPROACH

Drive Technology Innovation in Different Roles

■ Exascale Labs (or Competence Centres)



- Long-term collaboration with suppliers
- POWER Acceleration and Design Center
- Collaboration between Forschungszentrum Juelich, IBM and NVIDIA
- Mission statement: Provide support to scientists and engineers to target the grand challenges facing society in the fields of energy & environment, information & health care

■ Co-Design Projects

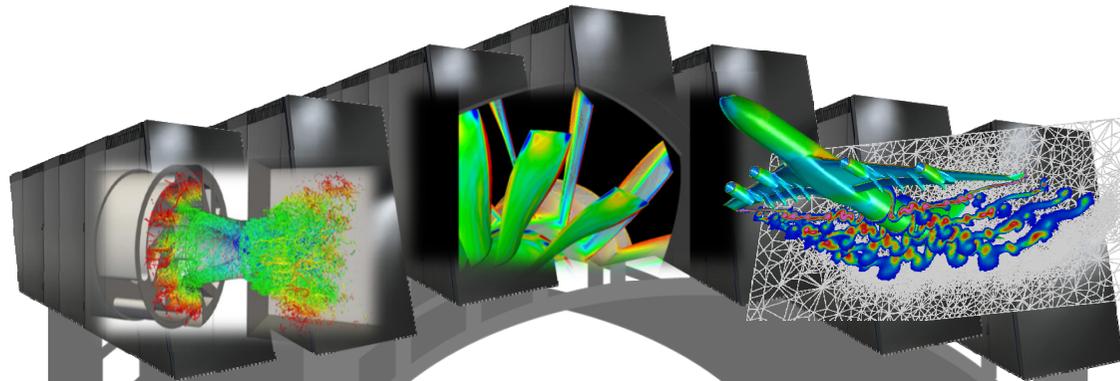
- E.g. DEEP projects & application use cases



(Selected JSC collaboration partners)

SIMULATION SCIENCES

Traditional Supercomputing Impact in Scientific Computing

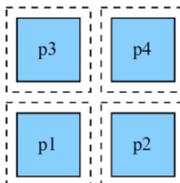
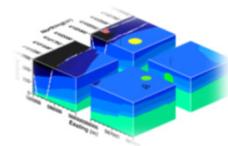
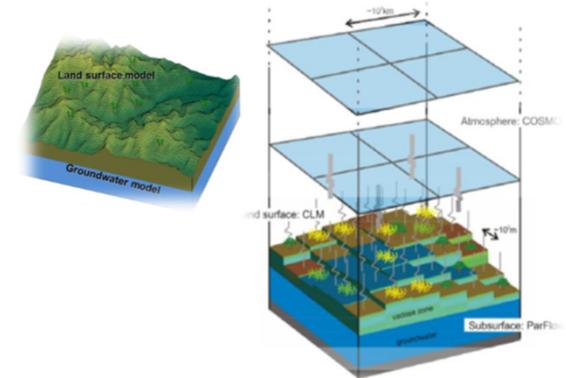


Numerical calculations... Model
...simulation over time

Experiment
'we observe
the nature'

Theory
'we create
a model
of nature'

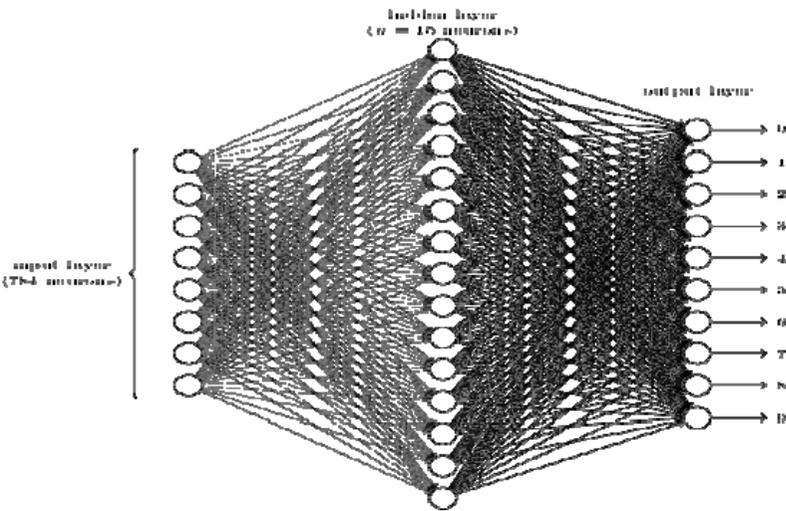
- Known physical laws
- Numerical methods
- Parallel Computing



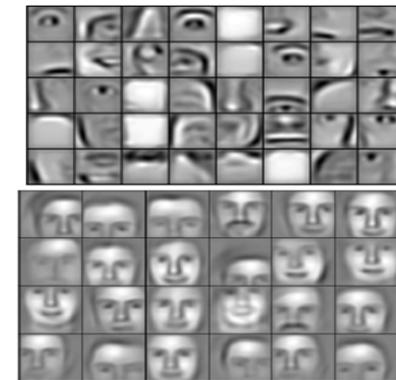
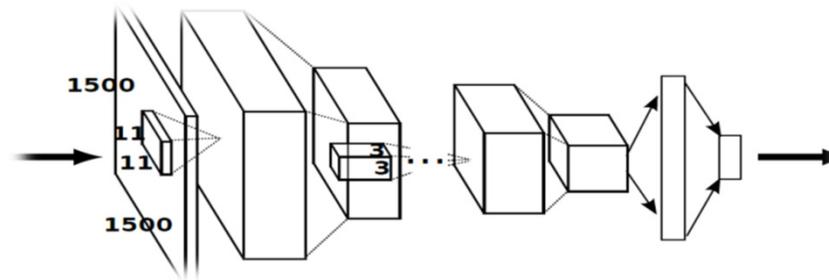
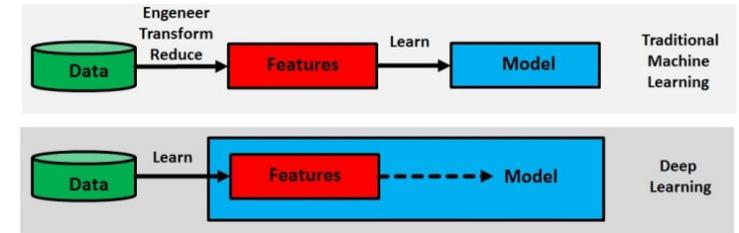
INNOVATIVE DEEP LEARNING TECHNOLOGIES

Short Overview & Role of Team Deep Learning for SIMDAS & Juelich Supercomputing Centre

■ Innovative & disruptive approach



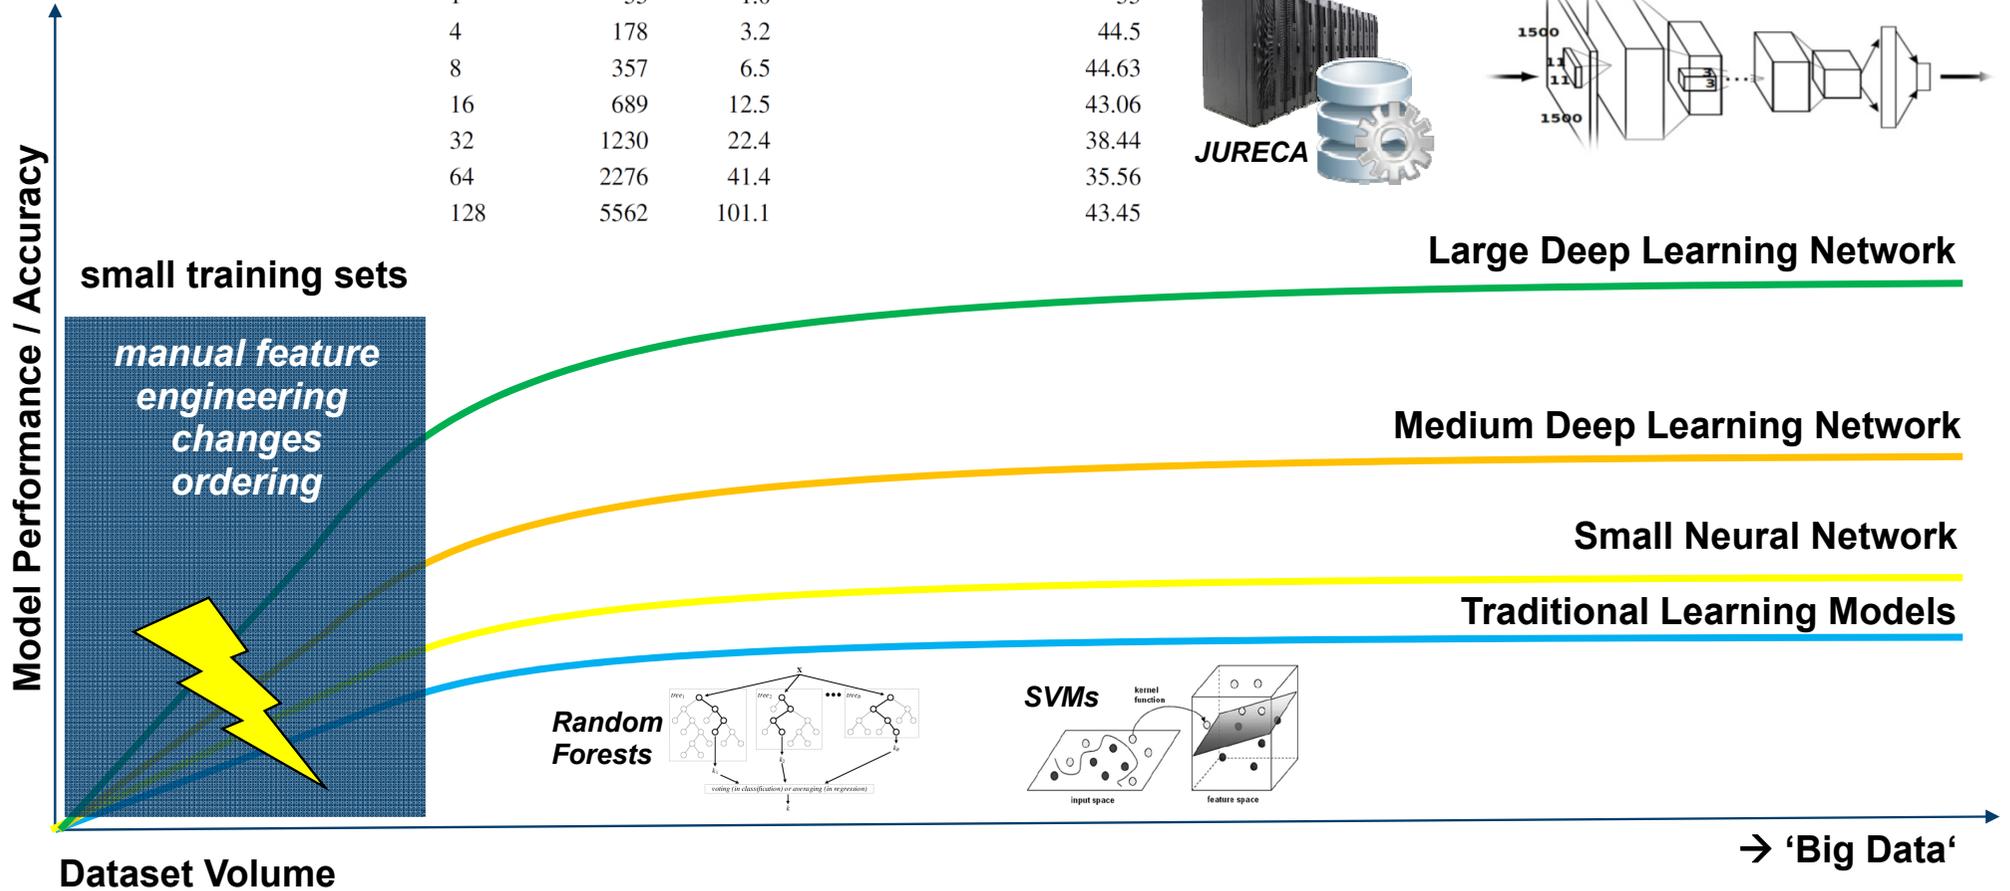
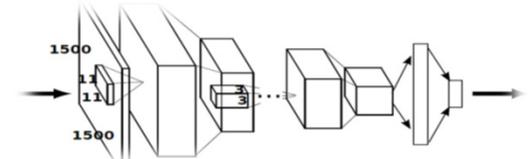
[8] M. Riedel, Invited YouTube Tutorial on Deep Learning, Ghent University



- Provide deep learning tools that work with HPC machines (e.g. Python/Keras/Tensorflow)
- Advance deep learning applications and research on HPC prototypes (e.g. DEEP-EST, etc.)
- Engage with industry (industrial relations team) & support SMEs (e.g. Soccerwatch)
- Offer tutorials & application enabling support for commercial & scientific users (e.g. YouTube)

RELATIONSHIP BIG DATA & ARTIFICIAL INTELLIGENCE (AI)

#GPUs	images/s	speedup	Performance per GPU [images/s]
1	55	1.0	55
4	178	3.2	44.5
8	357	6.5	44.63
16	689	12.5	43.06
32	1230	22.4	38.44
64	2276	41.4	35.56
128	5562	101.1	43.45



MODULAR SUPERCOMPUTING CO-DESIGN Towards Smart Exascale

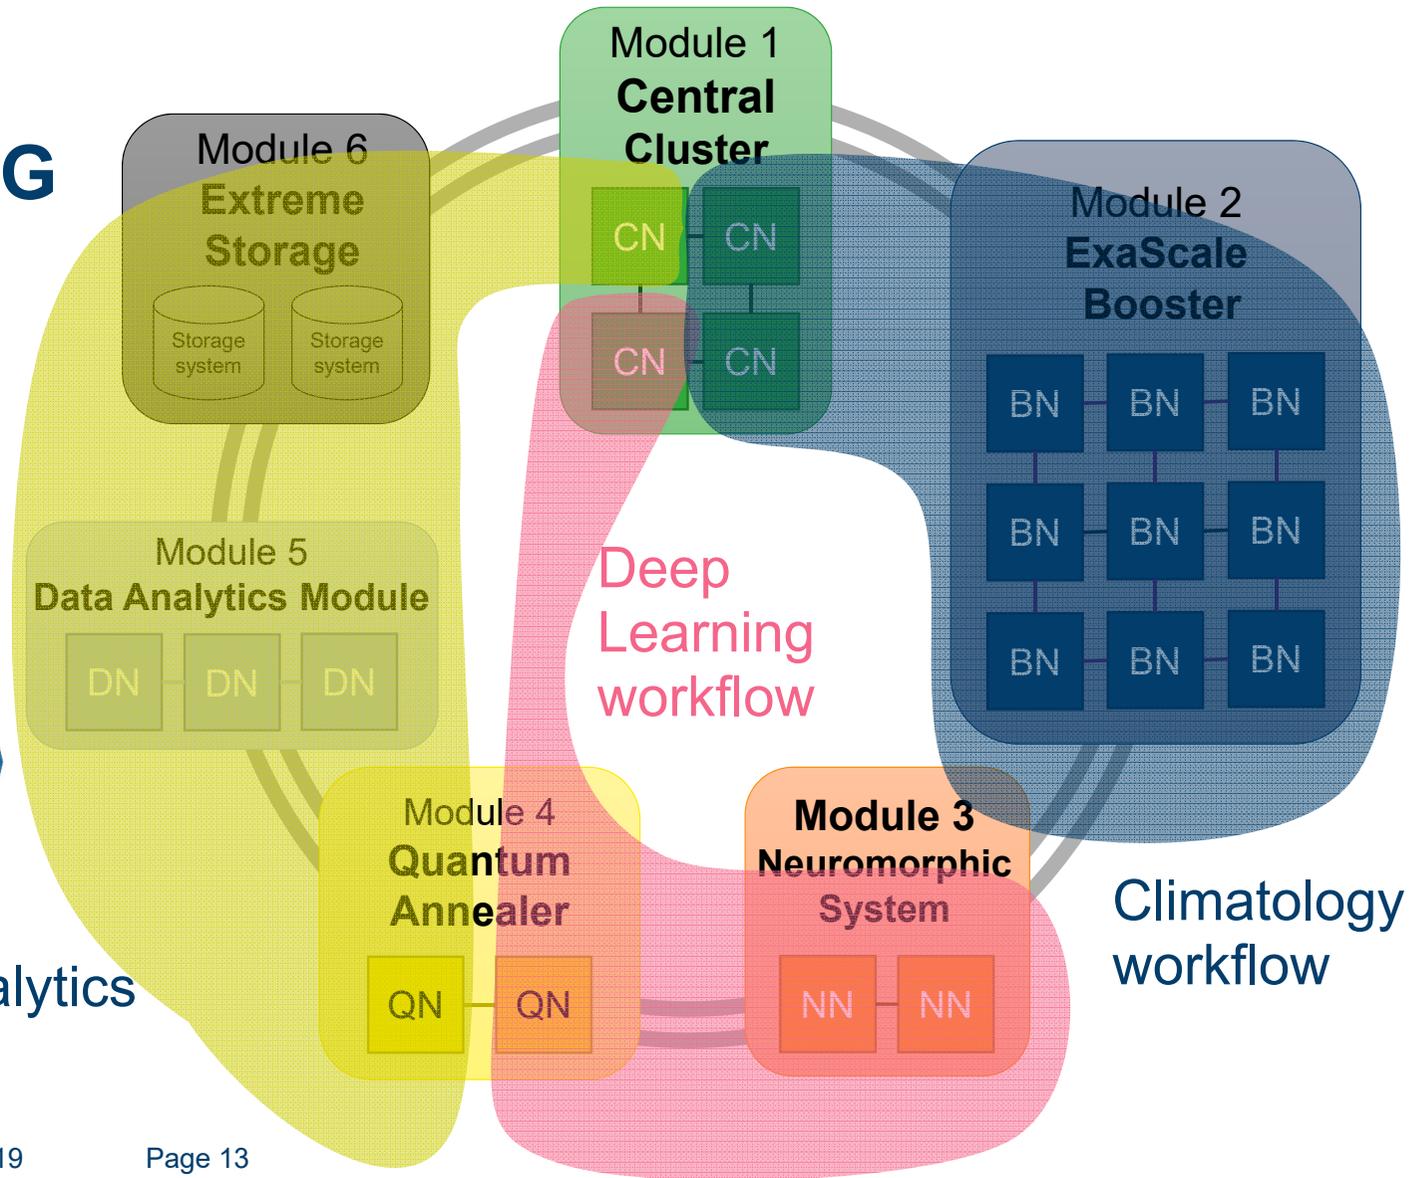


[2] DEEP Projects Web Page



HBP
Data Analytics
workflow

16th January 2019



Page 13

REQUIREMENTS OBTAINED FROM CO-DESIGN

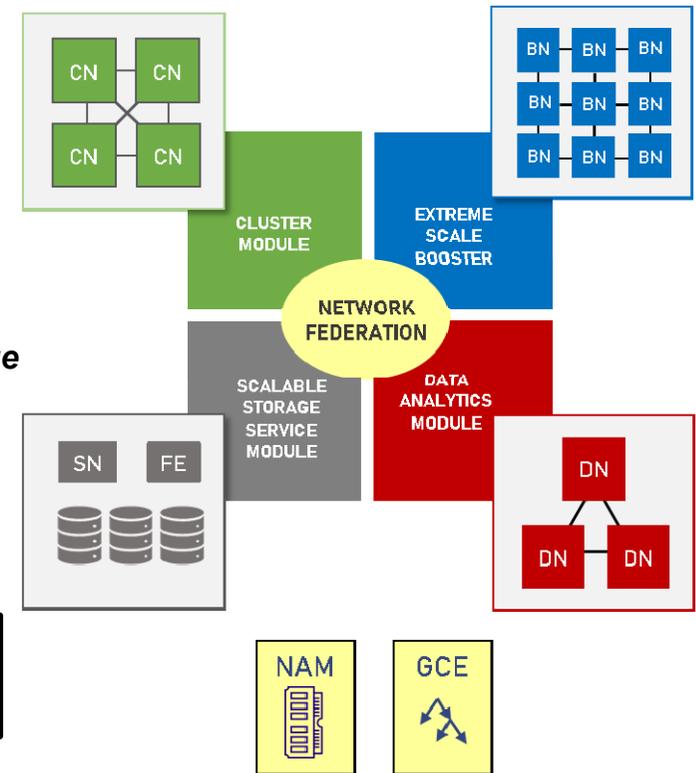


Towards a Modular Supercomputing Architecture with Network Federation

- Flop/s metric will become increasingly less(!) relevant
 - Driven by application co-design of HPC & Data Systems
 - Support for less regular computational tasks
 - Significantly larger memory footprint
 - Extreme data processing capabilities
- Improved/optimized data transport capabilities & specialized analytics
- Scalable visualisation capabilities
- Management of complex work-flows

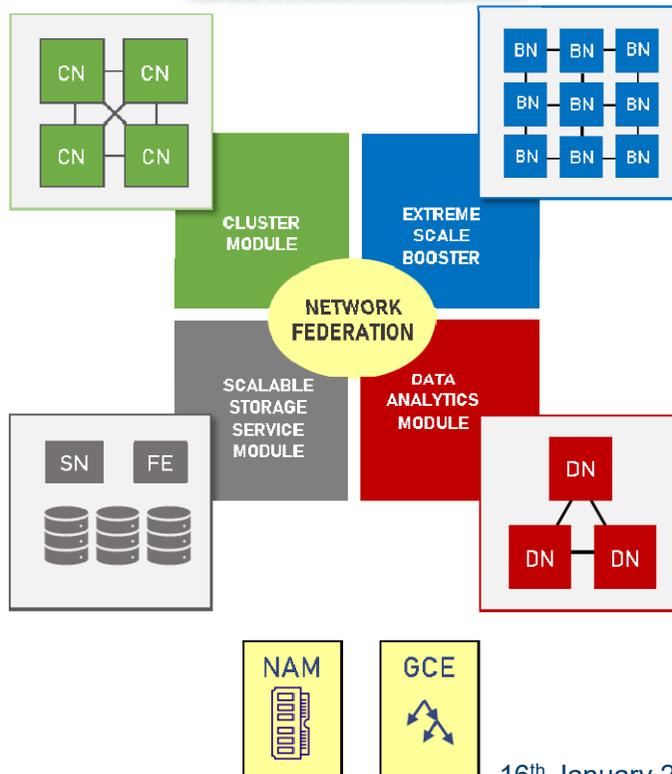
▪ One plausible answer to those facts is the modular supercomputer architecture driven by the DEEP projects – implemented in JSC

[2] DEEP Projects Web Page



EU HPC PROJECTS OVERVIEW

DEEP-EST Modular Supercomputing Architecture



[14] ETP4HPC

DEEP SERIES OF PROJECTS

EU Projects Driven by Co-Design of HPC Applications

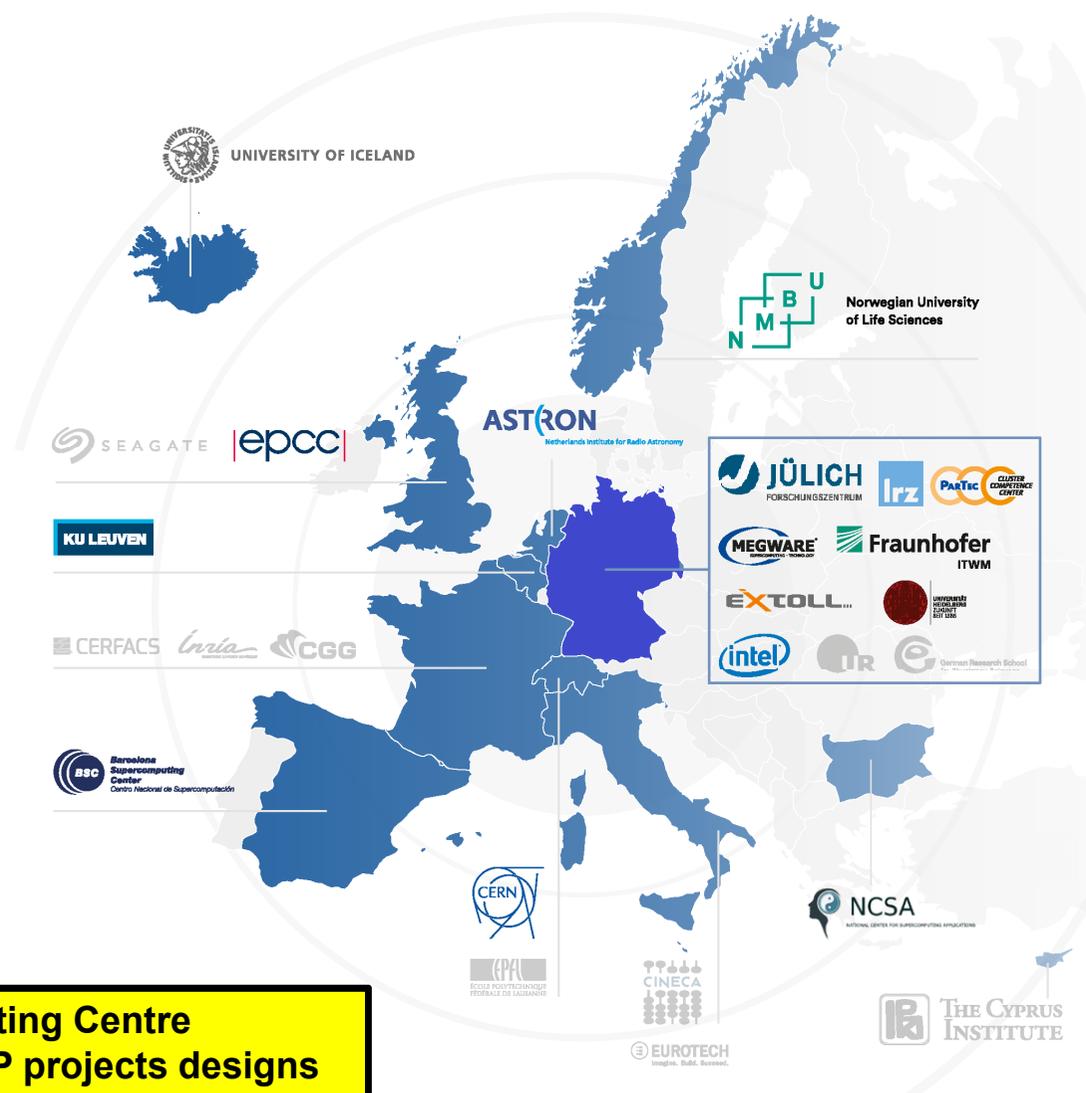


- 3 EU Exascale projects
DEEP
DEEP-ER
DEEP-EST
- 27 partners
Coordinated by JSC
- EU-funding: 30 M€
JSC-part > 5,3 M€
- Nov 2011 – Dec 2020

■ Strong collaboration with our industry partners Intel, Extoll & Megware

■ Innovative HPC hardware like Intel Nervana Neon and persistent RAMs

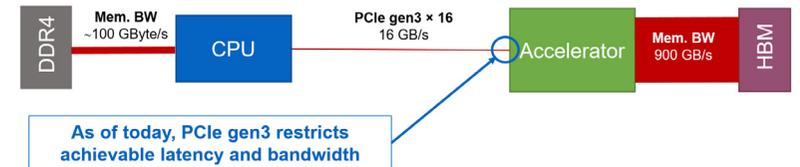
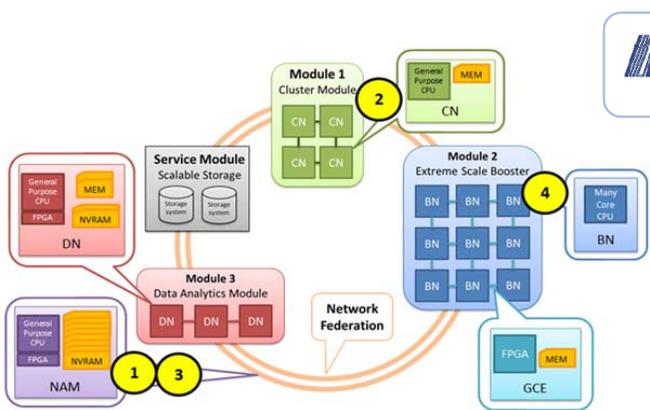
■ Juelich Supercomputing Centre implements the DEEP projects designs in its HPC production infrastructure



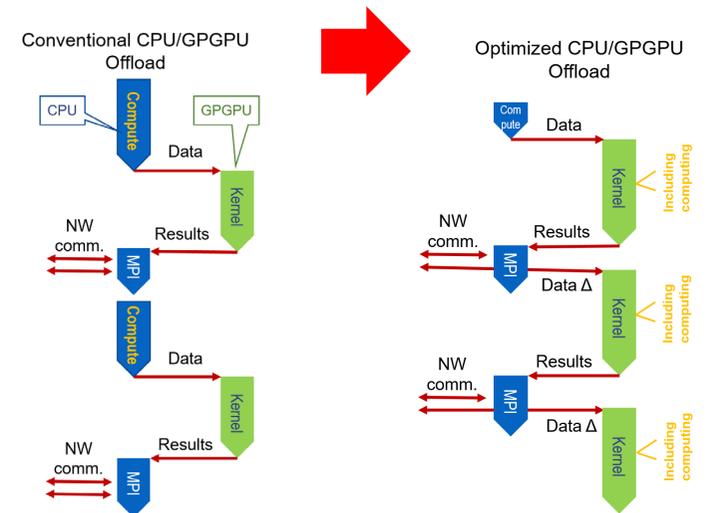
[2] DEEP Projects Web Page

DRIVING INNOVATIVE HPC HARDWARE VIA CO-DESIGN

Co-Design of Innovative HPC Memory Designs and GPU/CPU Communications in Modular Supercomputing

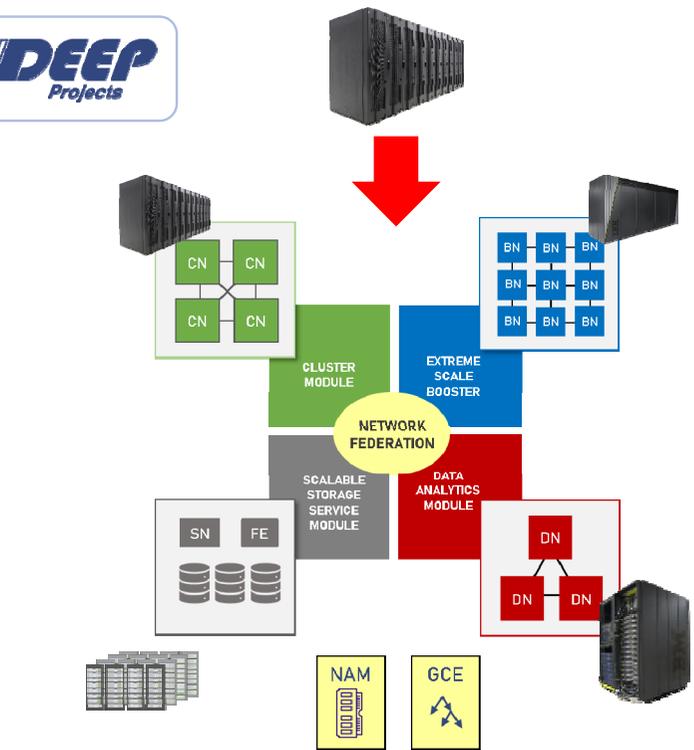
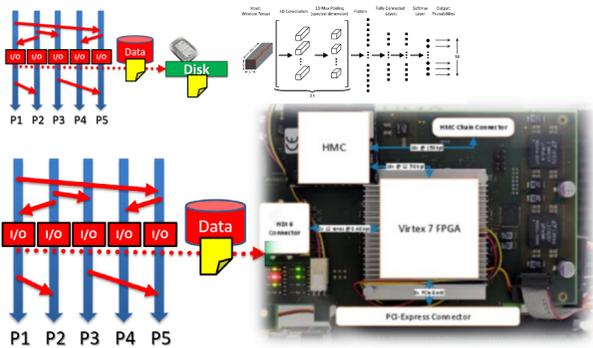


Explore scalability with NVIDIA GPUDirect beyond one node compared to NVIDIA NVLink/NVSwitch



[12] E. Erlingsson, M. Riedel et al., IEEE MIPRO Conference, 2018

Explore Network Attached Memory (NAM)

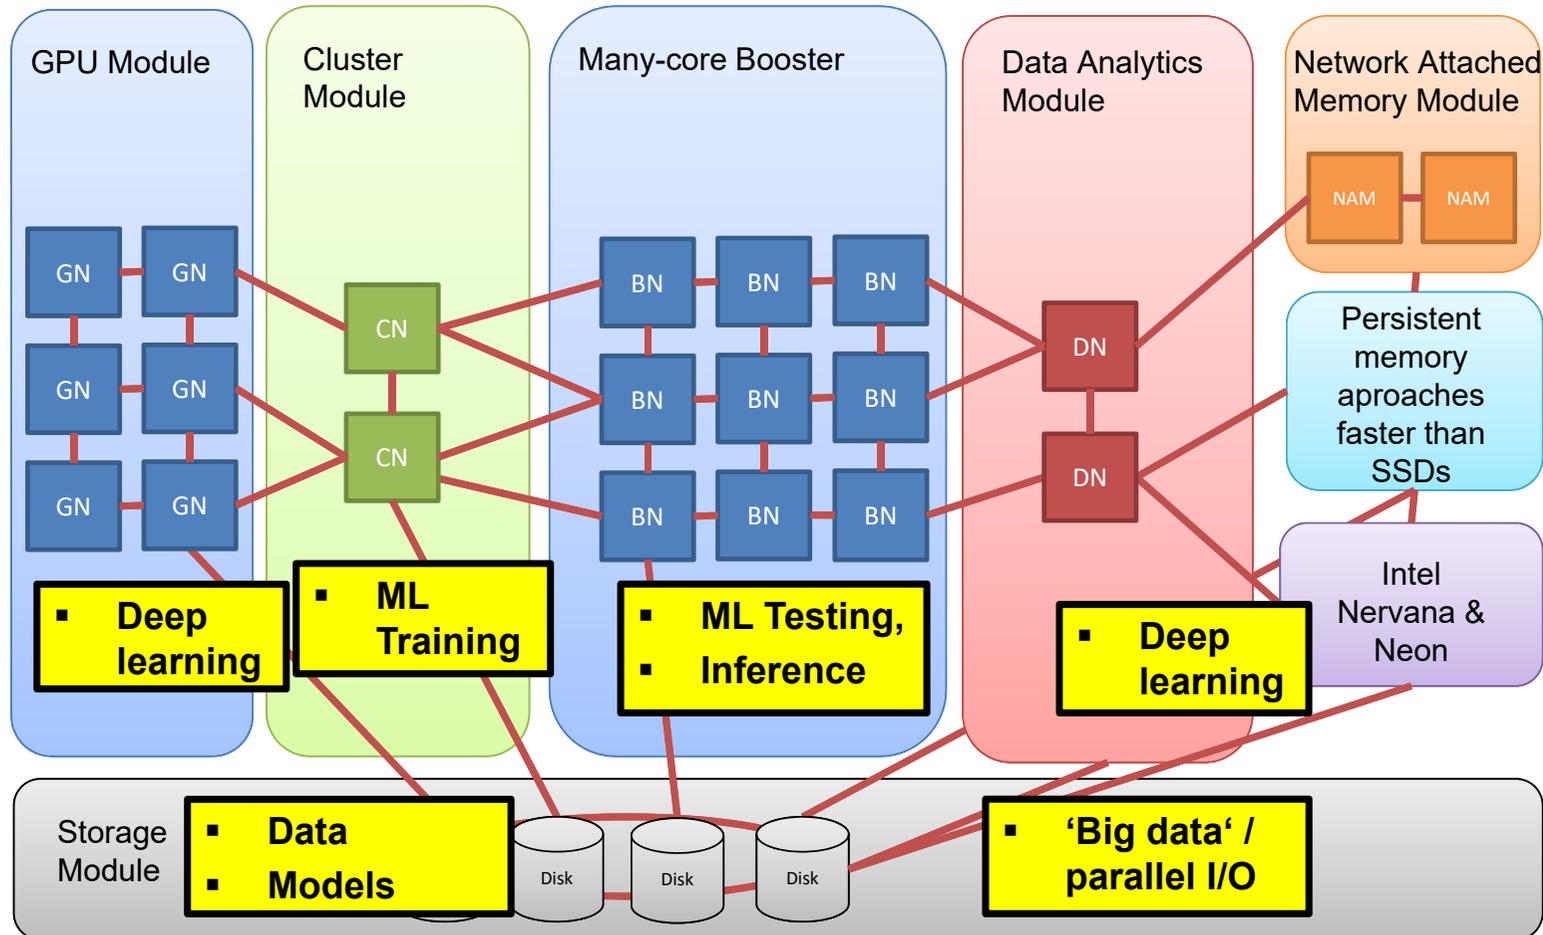


More details about the Modular Supercomputing Architecture (MSA) will be provided during my talk tomorrow: Overview of German Resources & Facilities

MODULAR SUPERCOMPUTING ARCHITECTURE

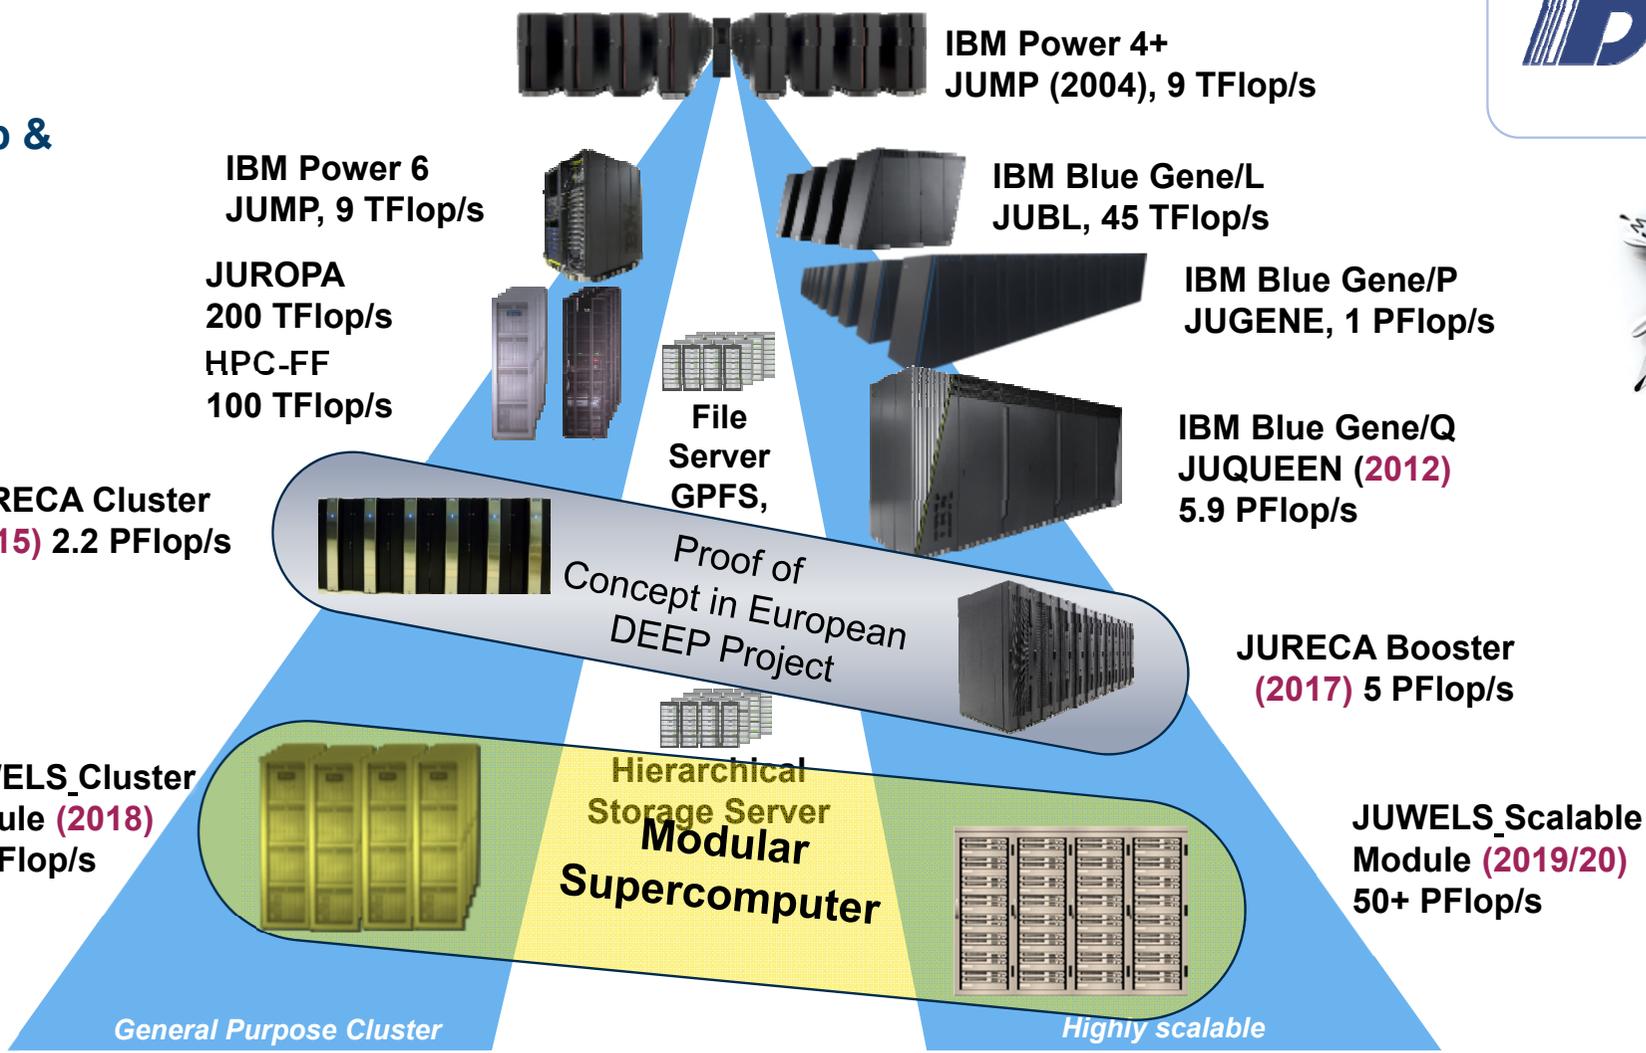


JSC Roadmap



JSC

HPC Roadmap & Key Vendors



RESEARCH FIELDS AND THEIR SHARES (2016-2018)

Supercomputing Systems Utilization → Requirement for sending in HPC project proposals to receive 'time'

Highly Scalable System



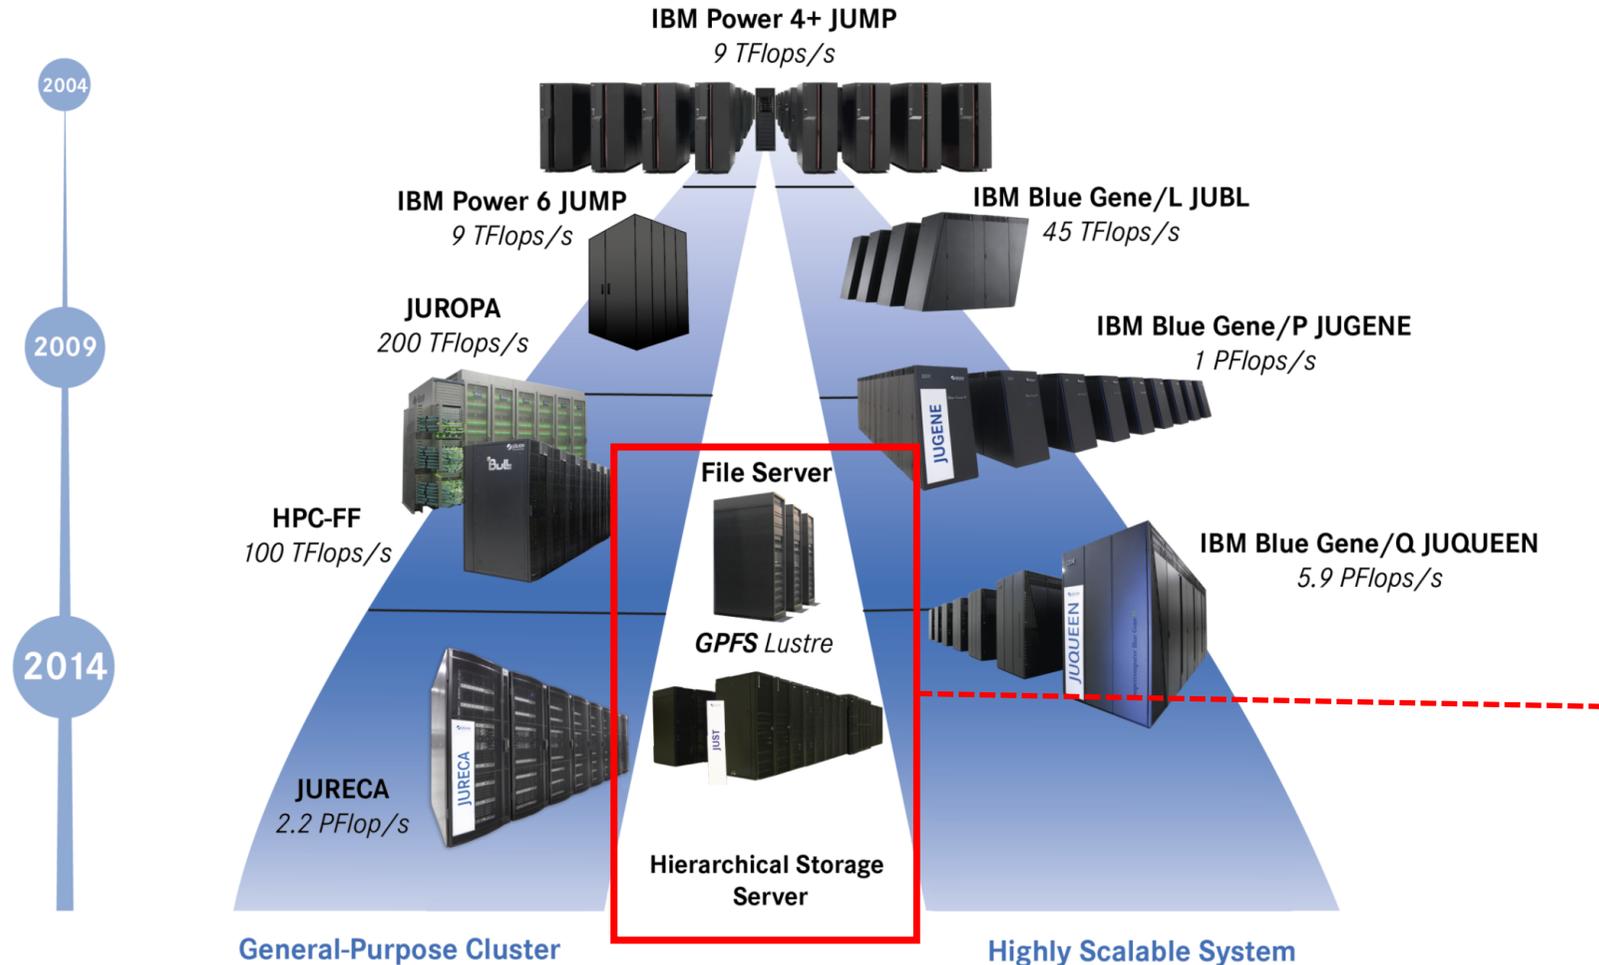
General-Purpose Supercomputer



- Astrophysics
- Plasma Physics
- Computer Science
- Biophysics
- Soft Matter
- Condensed Matter
- Chemistry
- Fluid Dynamics
- Materials Science
- Earth & Environment
- Elementary Particle Physics



HPC DRIVES 'BIG DATA' STORAGE



- Scientific computing applications can be considered as 'creators' of big data in science & engineering domains
- Traditional simulation sciences already require high capacity in data storages to store output of various application models (or checkpointing)

'BIG DATA' INFRASTRUCTURE FOR HPC & DATA SCIENCE

Multiple systems combined with whole federation of other Helmholtz centre systems

■ JUST Storage Cluster

- IBM Spectrum Scale file system (GPFS)
- 75 PB gross capacity
- 5th generation
- Parallel access

■ Tape Libraries

- Automated cartridge systems
- 300 PB
- 3 libraries (in 2 buildings)
- 60 tape drives
- 35,000 tapes



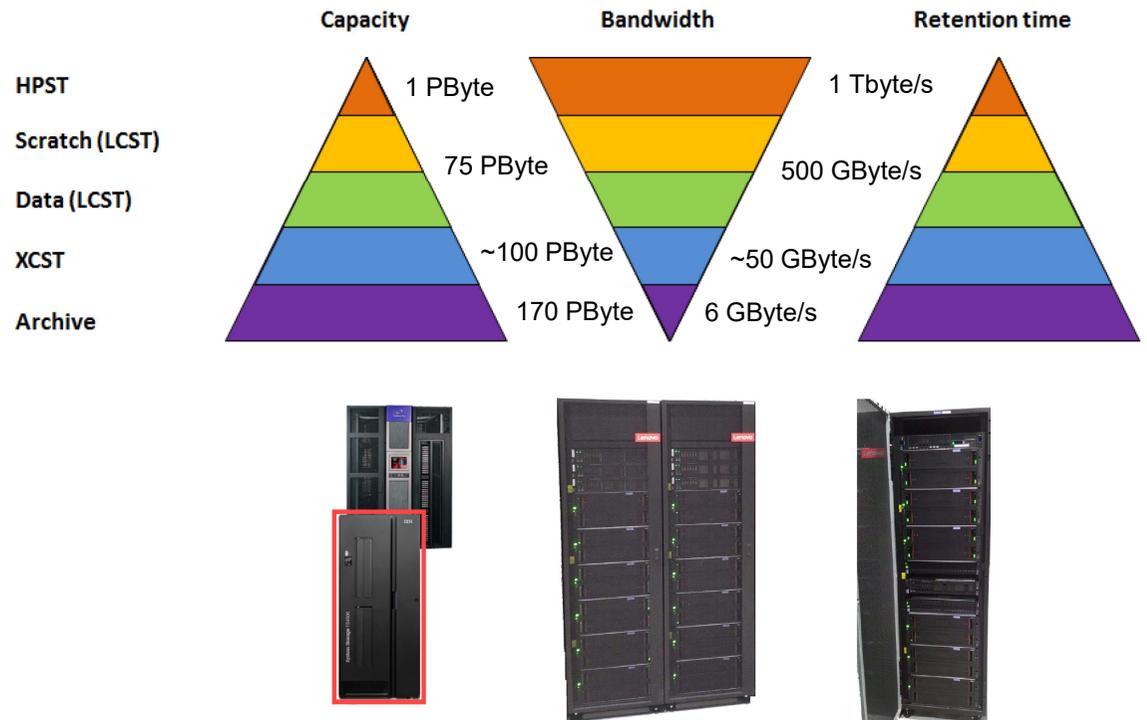
- **Journals will/do require to store data with persistent identifiers for papers**
- **Need for sharing of data in communities (e.g. turbulence flow)**

ACCESS TO DATA CRUCIAL FOR APPLICATIONS

Capacity – Bandwidth – Retention Time



CPUs & GPUs

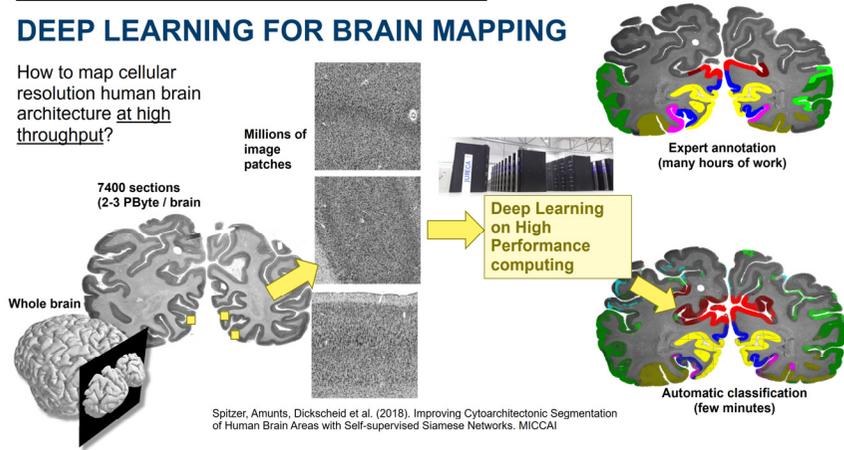


OTHER RELATED PROJECTS & APPLICATIONS

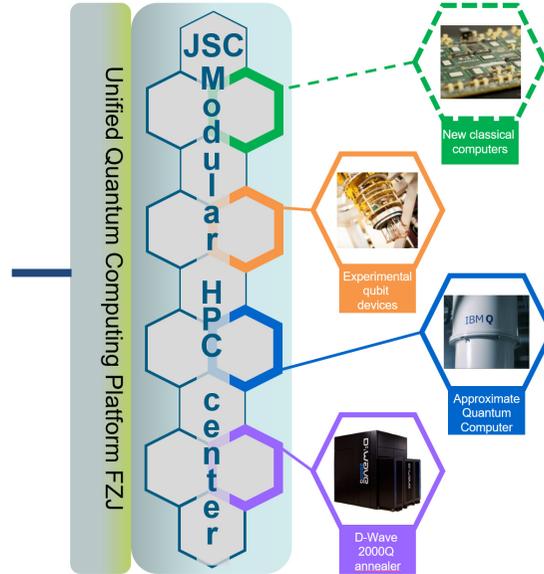
Prof. Katrin Amunts – INM-1

DEEP LEARNING FOR BRAIN MAPPING

How to map cellular resolution human brain architecture at high throughput?

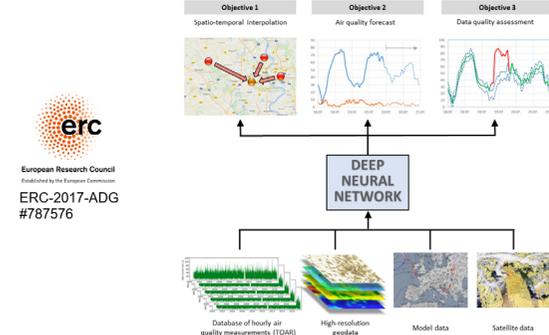


Prof. Kristel Michielsen – JSC

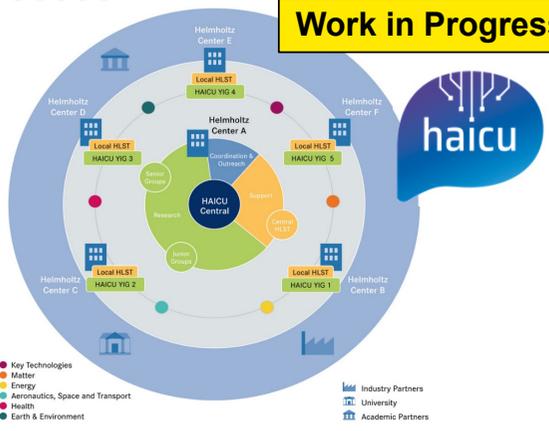


Dr. Martin Schultz – JSC

IntellIAQ: Anwendung von Deep Learning zur Interpolation, Vorhersage und Qualitätskontrolle von Daten zur Luftqualität

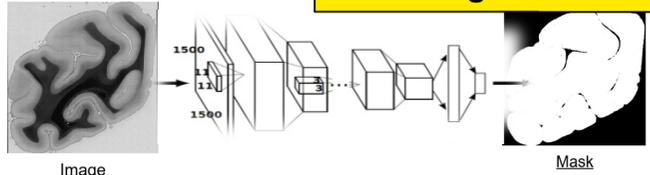


Work in Progress @ Helmholtz



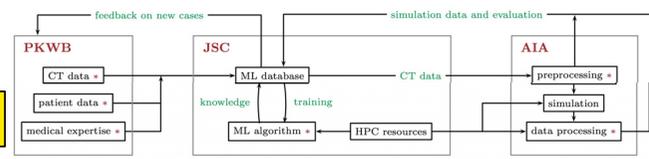
Helmholtz Artificial Intelligence Cooperation Unit

Prof. Abigail Morrison – JSC



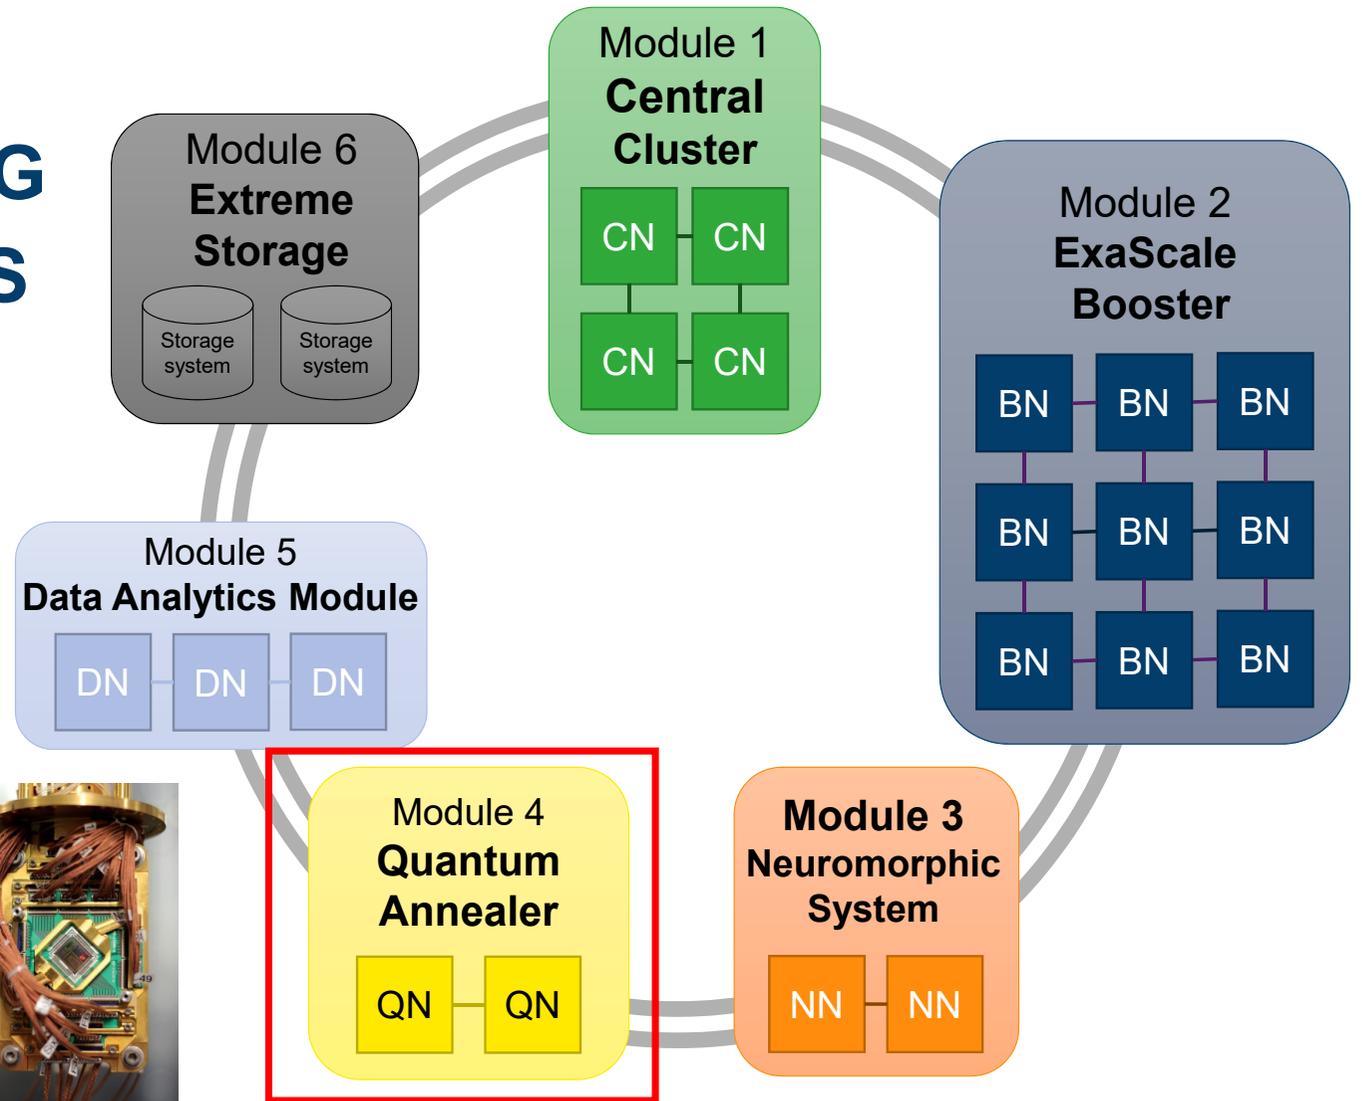
RHINO DIAGNO T JARA HPC

Dr. Andreas Lintermann – JSC/RWTH



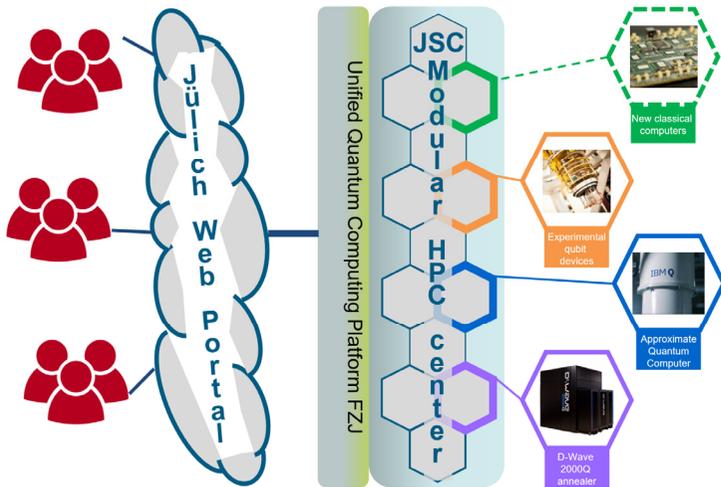
MODULAR SUPERCOMPUTING FUTURE MODULES Towards Smart Exascale

D:wave
The Quantum Computing Company™



EXPLORING INNOVATIVE COMPUTING PARADIGM

New Computing Paradigms like Quantum Computing



'quantum annealer and universal gate quantum computing based on same concepts, but are useful for different tasks and different problems'



'certain problems like optimization problems can be solved quadratically or polynomially faster with quantum algorithms, specifically quantum annealer (i.e. D-Wave)'

- Quantum Computing Short Summary [13] www.big-data.tips, 'Qubit'
 - System are based on circuits are called **qubits / quantum bits**
 - Qubits can not represent just 0 or 1 (transistors used today)
 - **Superposition**: a qubit can represent 0 and 1 simultaneously
→ N qubits = 2^N bits information
 - E.g. 2 qubits = 4 states, 3 qubits = 8 states
 - Calculations of qubits exploit superposition and jointly work with other qubits via **entanglement**
 - Connection retains a **very specific correlation** in their energy states (combinations of 0 and 1)
 - **Correlation represents multiple combinations of values simultaneously** instead of just one combination at a time

OTHER INTERESTING POINTERS

Helmholtz Association Activities & Research Data Alliance

■ Helmholtz Data Federation (HDF)

- Federation and extension of multi-topical data centers with new storage- and analysis hardware
- Usage of innovative data management solutions & excellent user support

■ Helmholtz Analytics Framework (HAF)

- Common components for data analytics
- Applied parallel machine learning methods

■ Research Data Alliance (RDA)

- Research data sharing without boundaries
- Interest groups and working groups

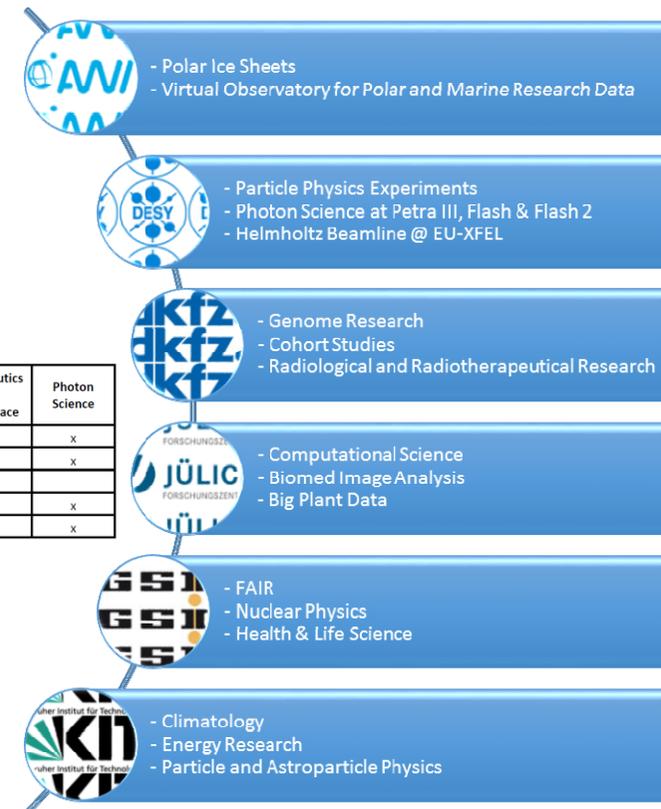
	Earth System Modeling	Structural Biology	Neuro-science	Aeronautics and Aerospace	Photon Science
Stochastik & Statistics	x		x	x	x
Optimisation	x			x	x
Supervised learning	x		x	x	
Unsupervised learning		x	x		x
Image Analysis			x		x



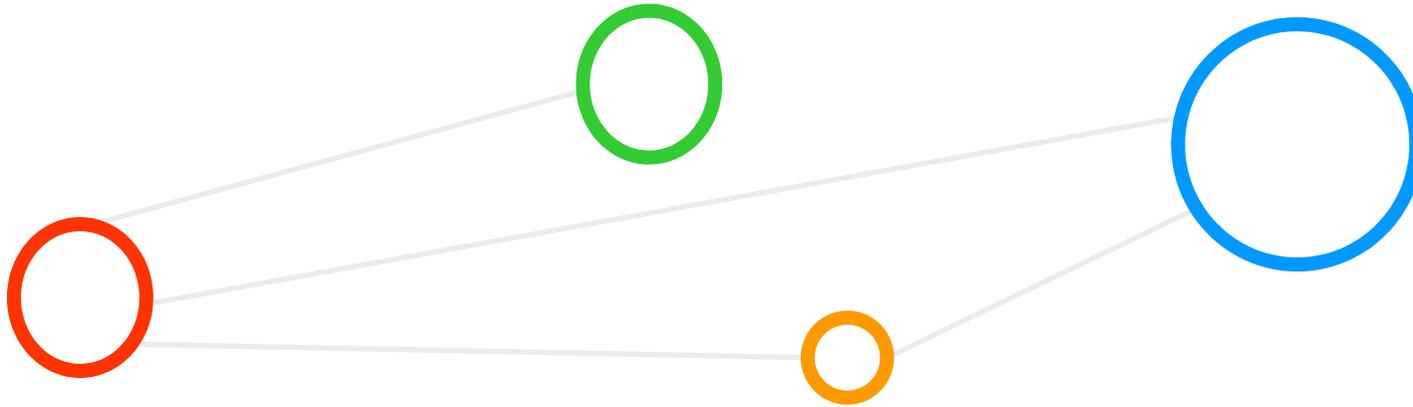
HELMHOLTZ
RESEARCH FOR GRAND CHALLENGES



NextGen
@HELMHOLTZ



REFERENCES



REFERENCES (1)

- [1] Helmholtz Association Web Page,
Online: <https://www.helmholtz.de/en/>
- [2] DEEP Projects Web Page,
Online: <http://www.deep-projects.eu/>
- [3] SMITH Projects Web Page,
Online: <http://www.smith.care>
- [4] Alfred Winter et al., 'Smart Medical Information Technology for Healthcare (SMITH) – Data Integration based on Interoperability Standards', submitted to Journal of Methods, 2018, to appear
- [5] AIXCAPE Web Page,
Online: <http://www.aixcape.org/association>
- [6] JSC Industry Relations Team (IRT) @ Juelich Supercomputing Centre,
Online: http://www.fz-juelich.de/ias/jsc/EN/Expertise/IndustryRelations/_node.html
- [7] OpenFOAM Web Page,
Online: <https://www.openfoam.com/>
- [8] M. Riedel, 'Deep Learning using a Convolutional Neural Network', Ghent University, Invited YouTube Tutorial,
Online: https://www.youtube.com/watch?v=gOL1_YlosYk&list=PLrmNhuZo9sgZUdaZ-f6OHK2yFW1kTS2qF

REFERENCES (2)

- [9] SoccerWatch.TV,
Online: <https://soccerwatch.tv/>
- [10] Smart Data Innovation Lab (SDIL),
Online: <https://www.sdil.de/en/>
- [11] European Commission, COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE EUROPEAN COUNCIL, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS, EC, 2018, 2nd May
Online: https://ec.europa.eu/commission/sites/beta-political/files/communication-modern-budget-may2018_en.pdf?utm_source=POLITICO.EU&utm_campaign=e3a8a86cc6-EMAIL_CAMPAIGN_2018_05_02&utm_medium=email&utm_term=0_10959edeb5-e3a8a86cc6-189710085
- [12] E. Erlingsson, G. Cavallaro, A. Galonska, M. Riedel, H. Neukirchen, 'Modular Supercomputing Design Supporting Machine Learning Applications', in conference proceedings of the 41st IEEE MIPRO 2018, May 21-25, 2018, Opatija, Croatia
- [13] Big Data Tips, 'Qubit',
Online: <http://www.big-data.tips/qubit>
- [14] ETP4HPC,
Online: <https://www.etp4hpc.eu/>

ACKNOWLEDGEMENTS

Previous & current members of the High Productivity Data Processing Research Group



PD Dr.
G. Cavallaro



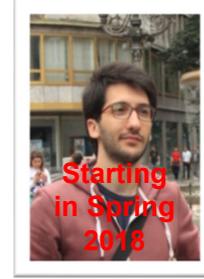
Senior PhD
Student A.S. Memon



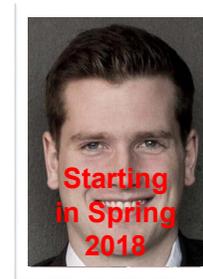
Senior PhD
Student M.S. Memon



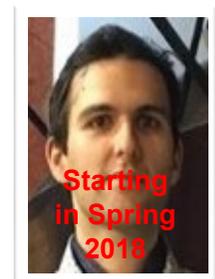
PhD Student
E. Erlingsson



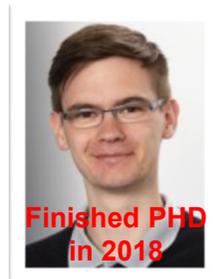
PhD Student
S. Bakarar



PhD Student
A. Pálsson



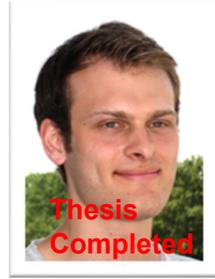
PhD Student
R. Sedona



Dr. M. Goetz
(now KIT)



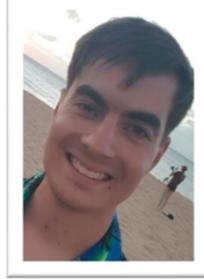
MSc M.
Richerzhagen
(now other division)



MSc
P. Glock
(now INM-1)



MSc
C. Bodenstein
(now Soccerwatch.tv)



MSc Student
G.S. Guðmundsson
(Landsverkjun)



This research group has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 763558 (DEEP-EST EU Project)

