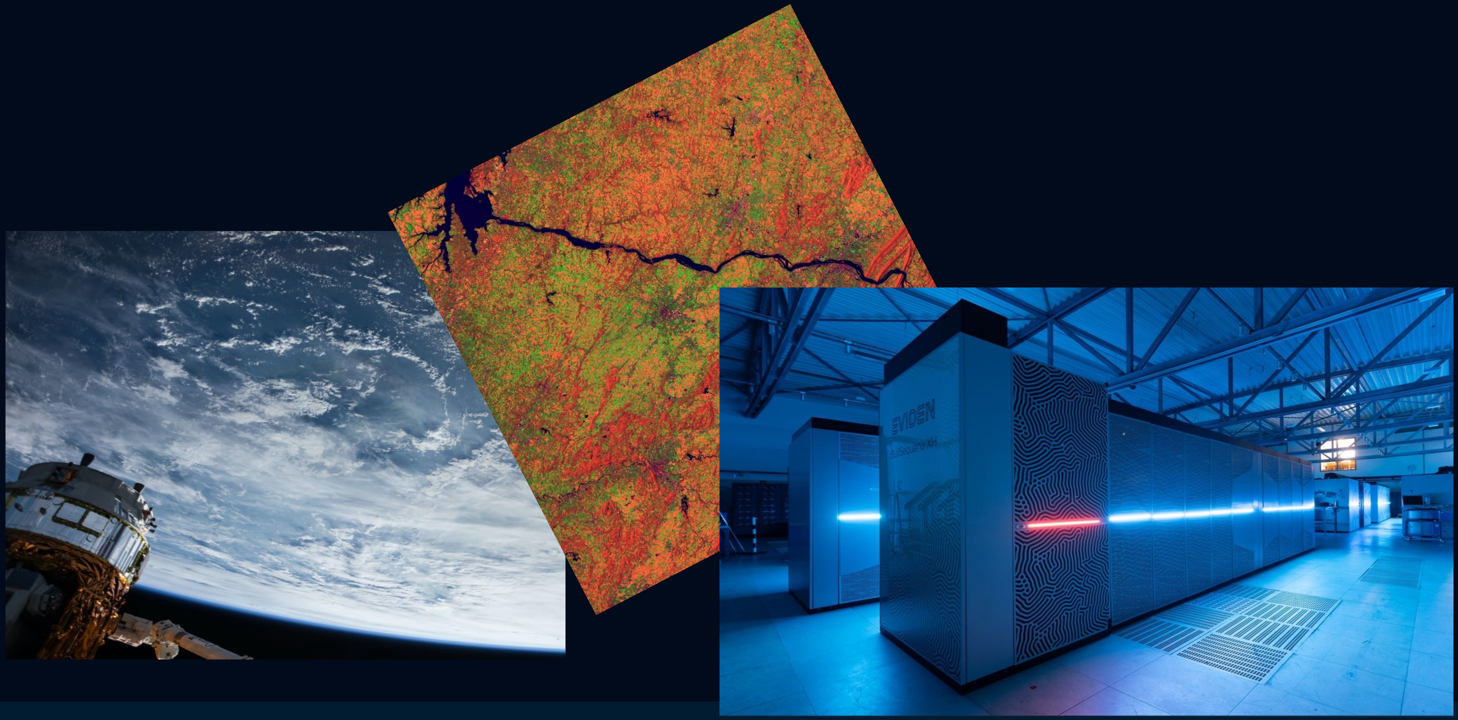

ADVANCING GEOSCIENCE THROUGH LARGE-SCALE AI WITH SUPERCOMPUTING FOR EARTH OBSERVATION AND REMOTE SENSING

4th Summer School on High-Performance and Disruptive Computing in Remote Sensing I 4 June 2024 | Santiago de Compostela (Spain)

GABRIELE CAVALLARO (WWW.GABRIELE-CAVALLARO.COM)
HEAD OF SIMULATION AND DATA LAB "AI AND ML FOR REMOTE SENSING", JÜLICH SUPERCOMPUTING CENTRE (FORSCHUNGSZENTRUM JÜLICH)
ADJUNCT ASSOCIATE PROFESSOR, SCHOOL OF ENGINEERING AND NATURAL SCIENCES, UNIVERSITY OF ICELAND
CO-CHAIR OF IEEE GRSS EARTH SCIENCE INFORMATICS TECHNICAL COMMITTEE (ESI TC)

IEEE GRSS Geoscience and Remote Sensing Society | IEEE GRSS ESI Earth Science Informatics Technical Committee | JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE | UNIVERSITY OF ICELAND

OUTLINE

1. Evolution of Computing

2. Machine Learning drives Computational Demands

3. Supercomputing and Parallel Programming

4. Jülich Supercomputing Centre

OUTLINE

1. **Evolution of Computing**
2. Machine Learning drives Computational Demands
3. Supercomputing and Parallel Programming
4. Jülich Supercomputing Centre

COMPUTING TECHNOLOGIES HAVE EVOLVED IN RECENT DECADES

Moore's Law

Drove the semiconductor industry to cram more and more transistors and logic into the same volume

End of Dennard's Scaling

Limits in how much it is possible to shrink voltage and current without losing predictability

Multi-Core Era

Triggered by the Instruction Level Parallelism (ILP) wall

Amdahl's Law

Challenges in terms of energy efficiency, thermal management and parallelizability

New Trends

To obtain the best performance-cost-energy tradeoffs for defined tasks

MICROPROCESSOR TREND DATA

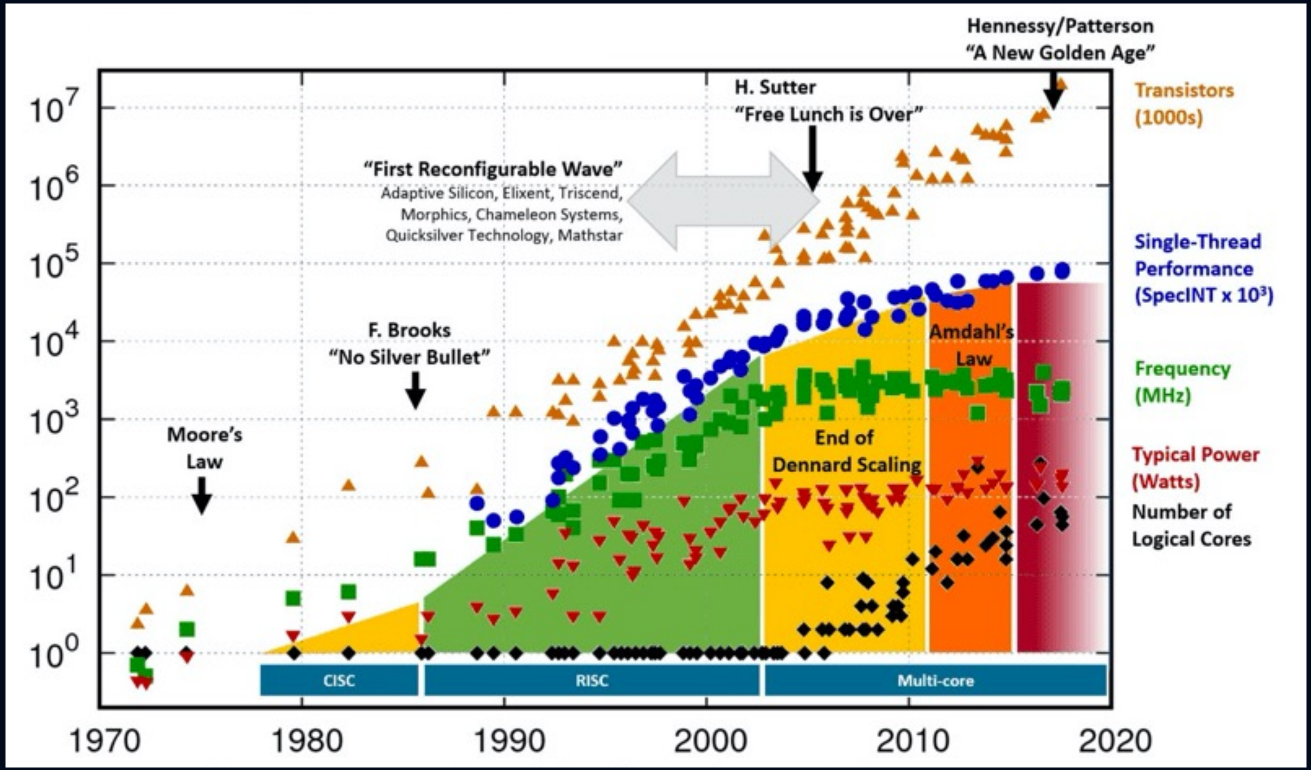

J. L. Hennessy, D. A. Patterson, "A New Golden Age for Computer Architecture", in Communications of the ACM, vol. 62 no. 2, pp. 48-60, 2019, https://doi.org/10.1145/3282307

Hennessy and Patterson, Turing Lecture 2018, overlaid over "42 Years of Processors Data" https://www.karirupp.net/2018/02/42-years-of-microprocessor-trend-data/; "First Wave" added by Les Wilson, Frank Schirrmeister Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten New plot and data collected for 2010-2017 by K. Rupp

SUPERCOMPUTING

SUPERCOMPUTING EVOLUTION

Architecture paradigms

- **1940 – 1950**: first computers are Supercomputers (specialized, expensive)
- **1960 – 1980**: vector computers dominate HPC,
  while general purpose computers come to market at much lower prices
  - Focus: floating operations (linear Algebra)
  - **Special purpose technologies** (fast vector processors, parallel architectures)
  - Only few machines produced → **expensive!**

- **1990 – 2000**: cluster computers are born
  - Integrate general purpose CPUs in HPC → **more economic approach**
  - Many „computers“ connected through fast network
  - Distributed memory → MPI

- **2010 – 2020**: heterogeneous cluster systems
  - CPU + Accelerator technologies (mostly GPUs) → **more FLOPS/Watt**
  - Intel / AMD + NVIDIA / AMD / Intel
- **2020 – today**: very large GPU-based systems in HPC,
  while hyperscalers dominate AI-market, drive GPU design (and price), and build their own processors for their clouds

Credit: Bernd Mohr (Jülich Supercomputing Centre)

MARKET CAPITALIZATIONS

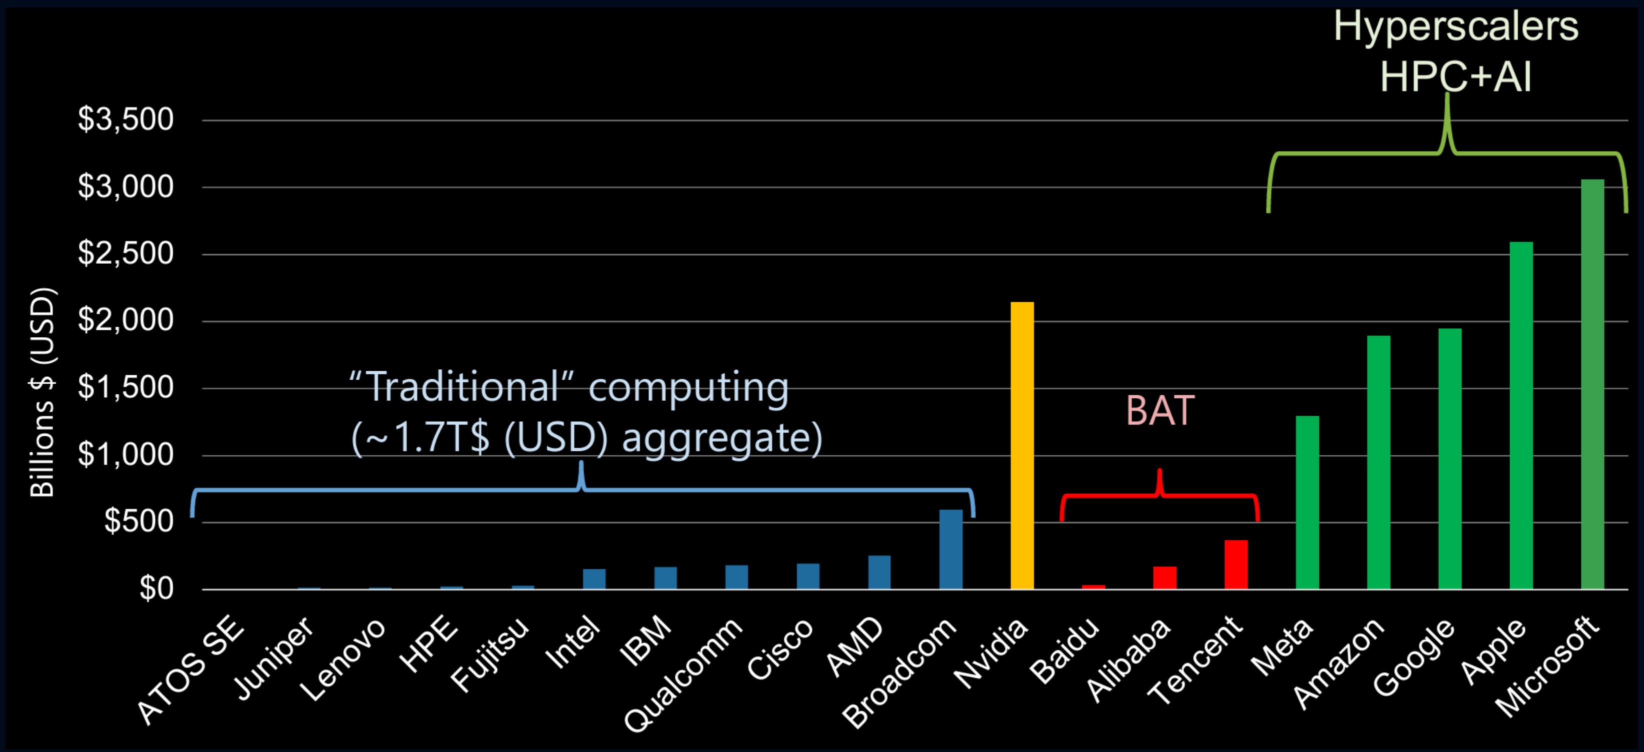

HPC looks at Cloud Computing

- Availability
- Cost vs price
- Higher level programming

Source: Reed, Gannon, Dongarra

Katherine Yelic, "Beyond Exascale Computing", ISC High Performance 2024

PERFORMANCE OF THE WORLD'S FASTEST SUPERCOMPUTERS

2008
First 1 PFLOP/s computer
(Roadrunner)

1997
First 1 TFLOP/s computer:
(Intel ASCI Read/9152)

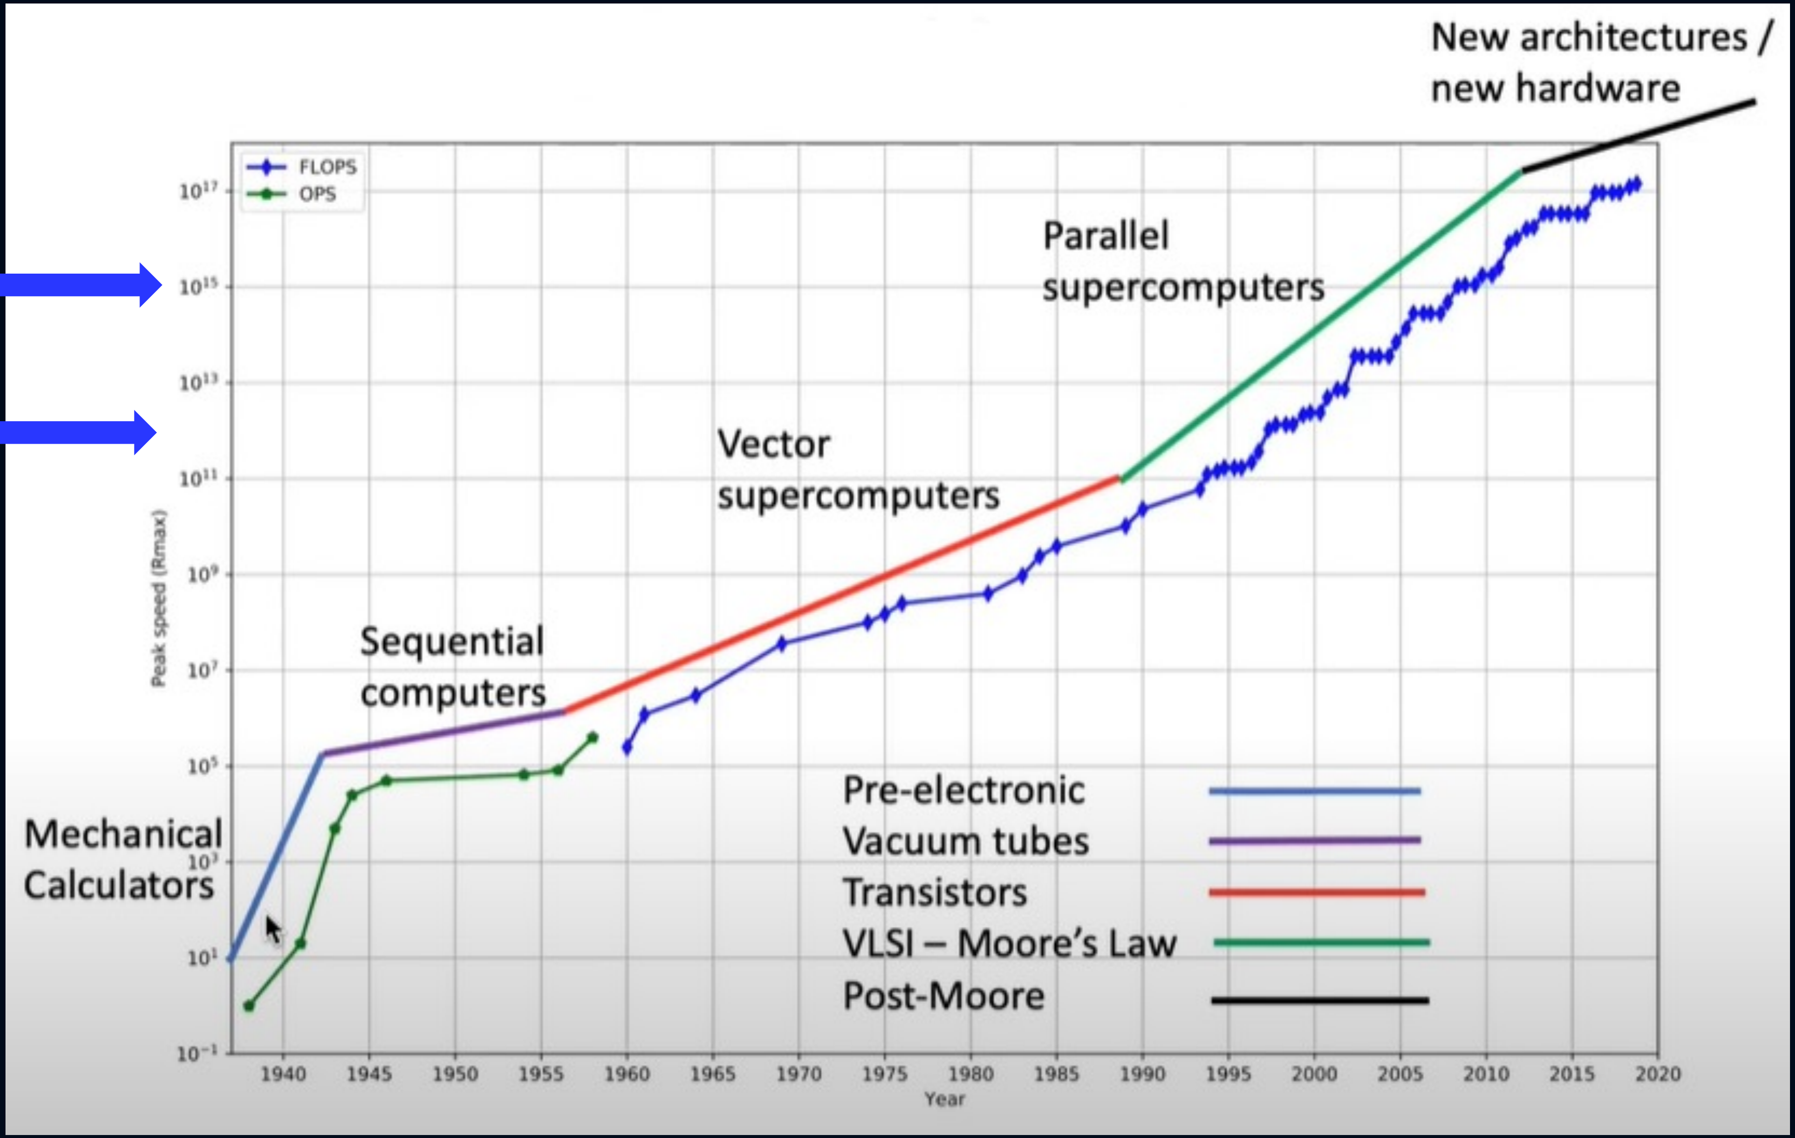

*floating point operations per second (FLOPS, flops or flop/s)

Rob Schreiber, High Performance Computing: Beyond Moore's Law https://www.youtube.com/watch?v=LOf57fdxln4&t=258s

FROM PETASCALE TO EXASCALE COMPUTING

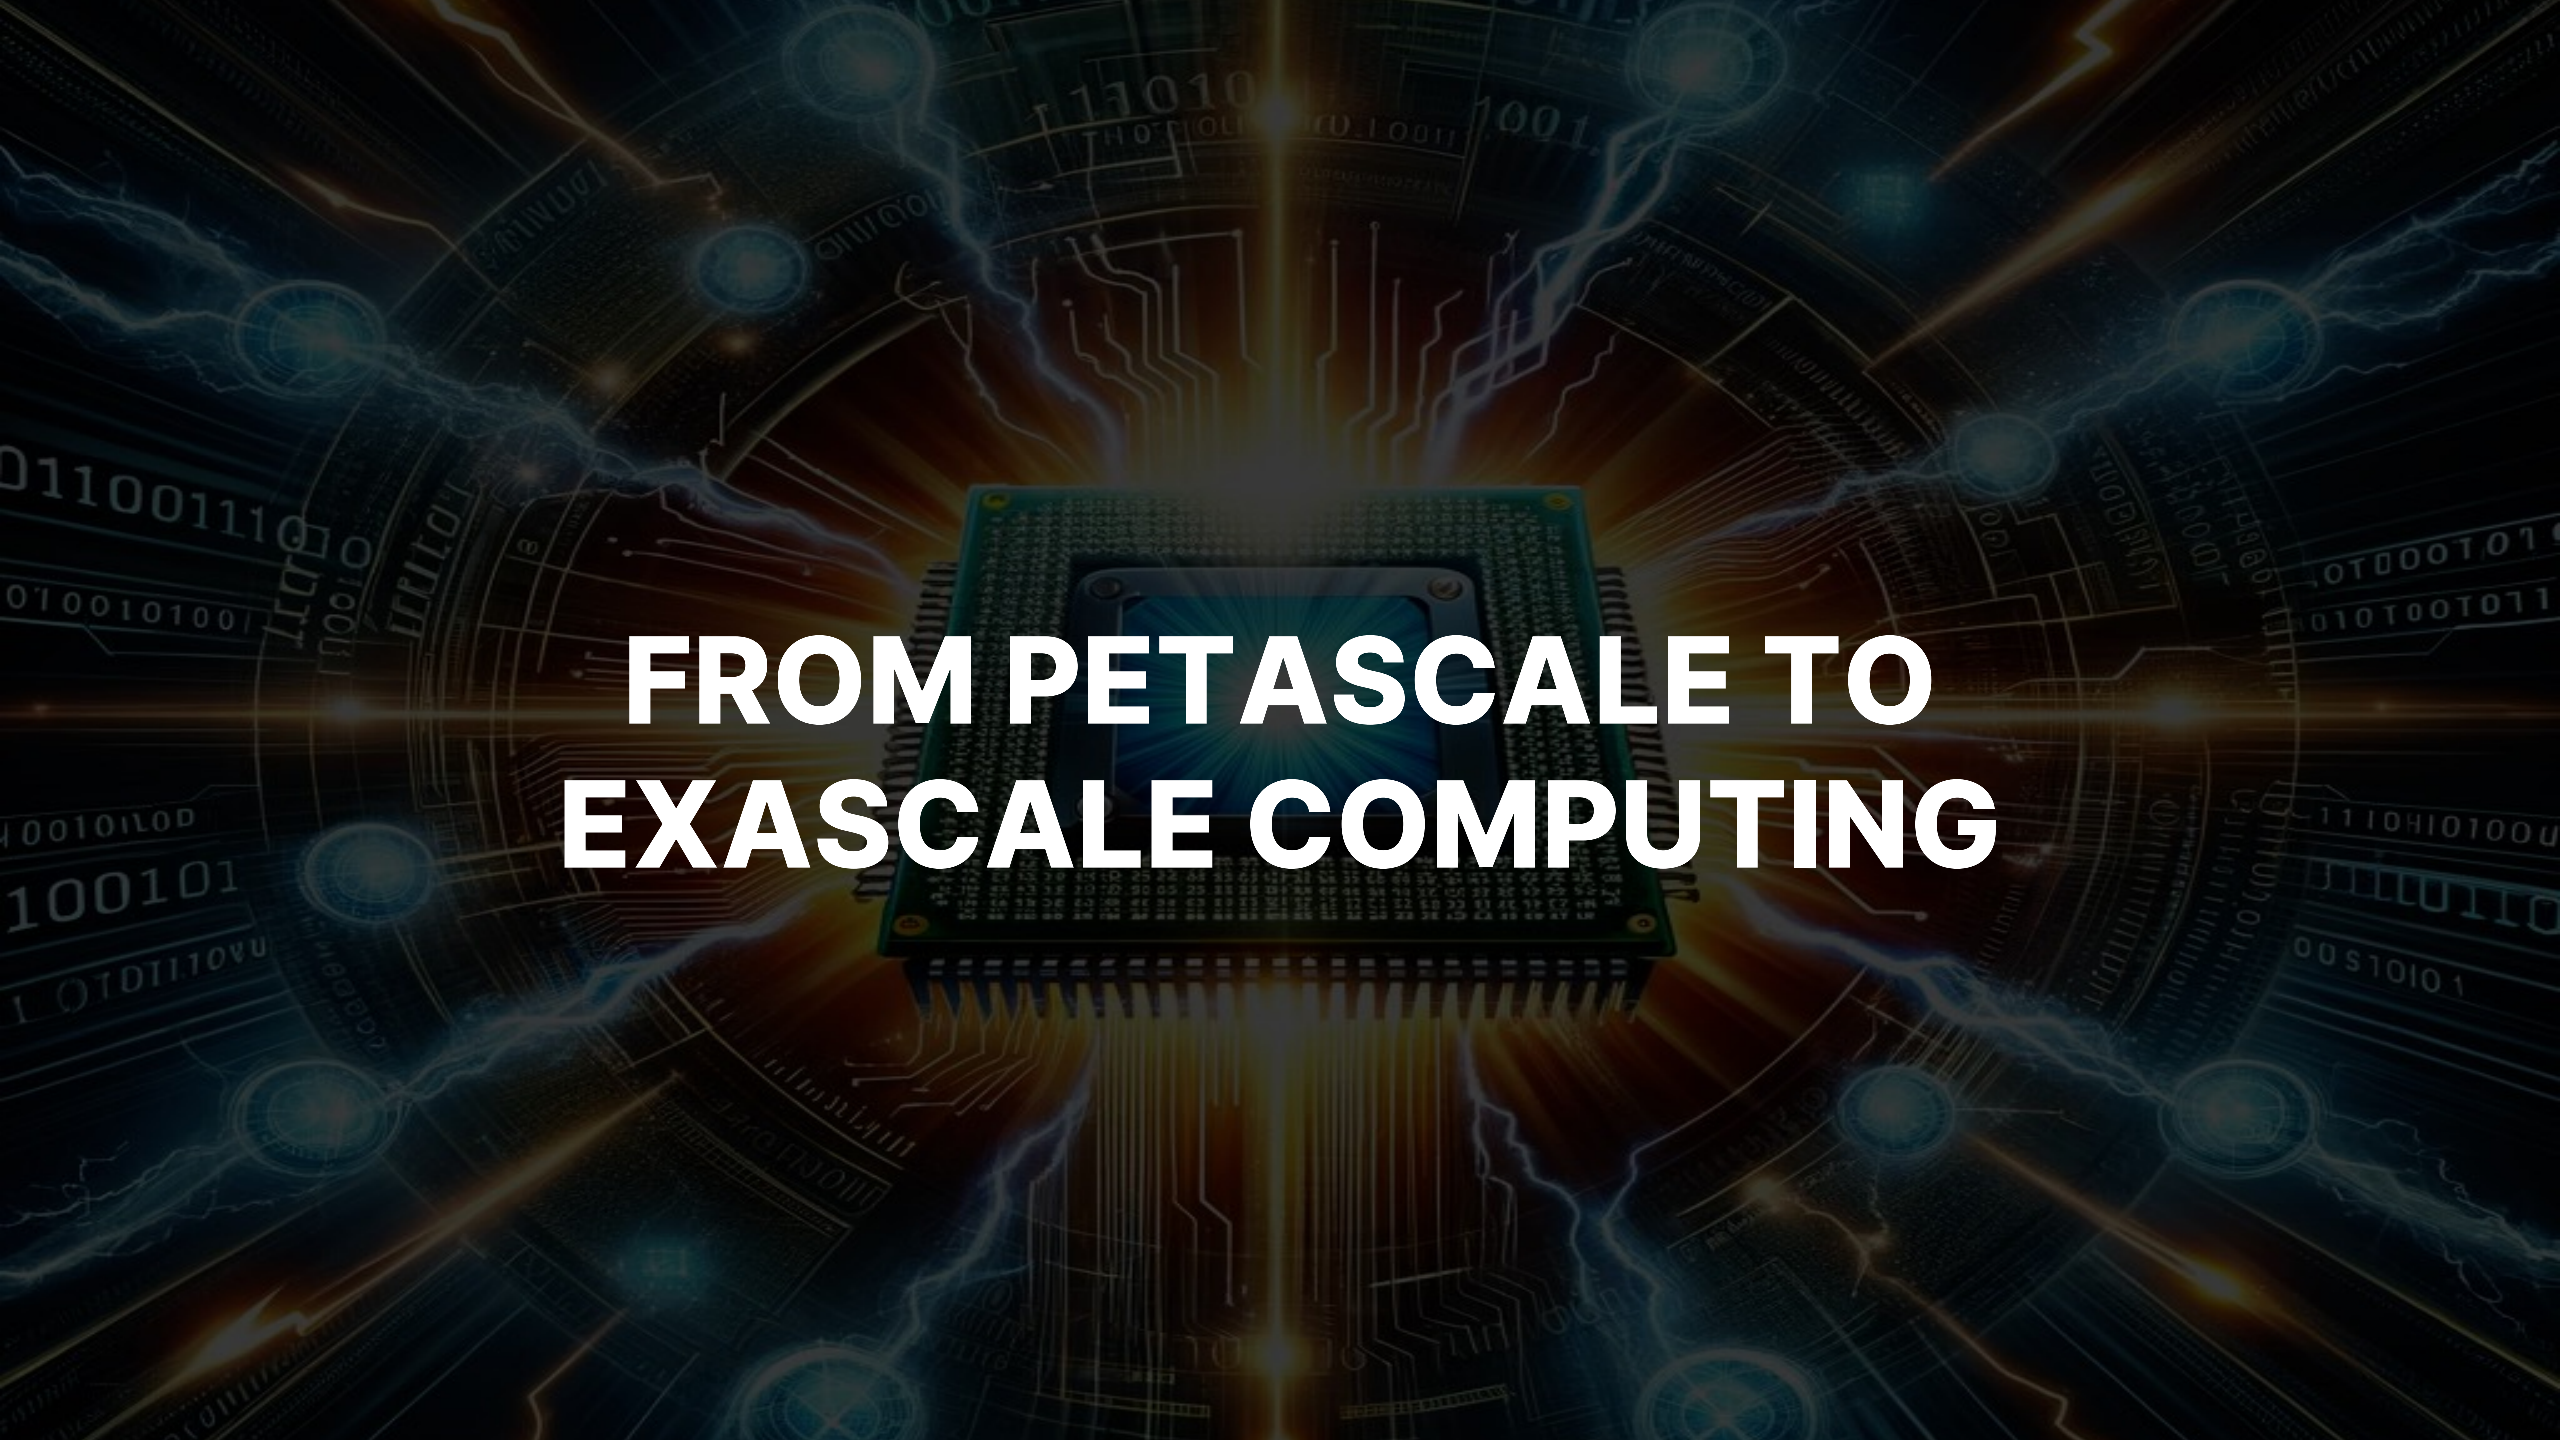

EXASCALE ERA

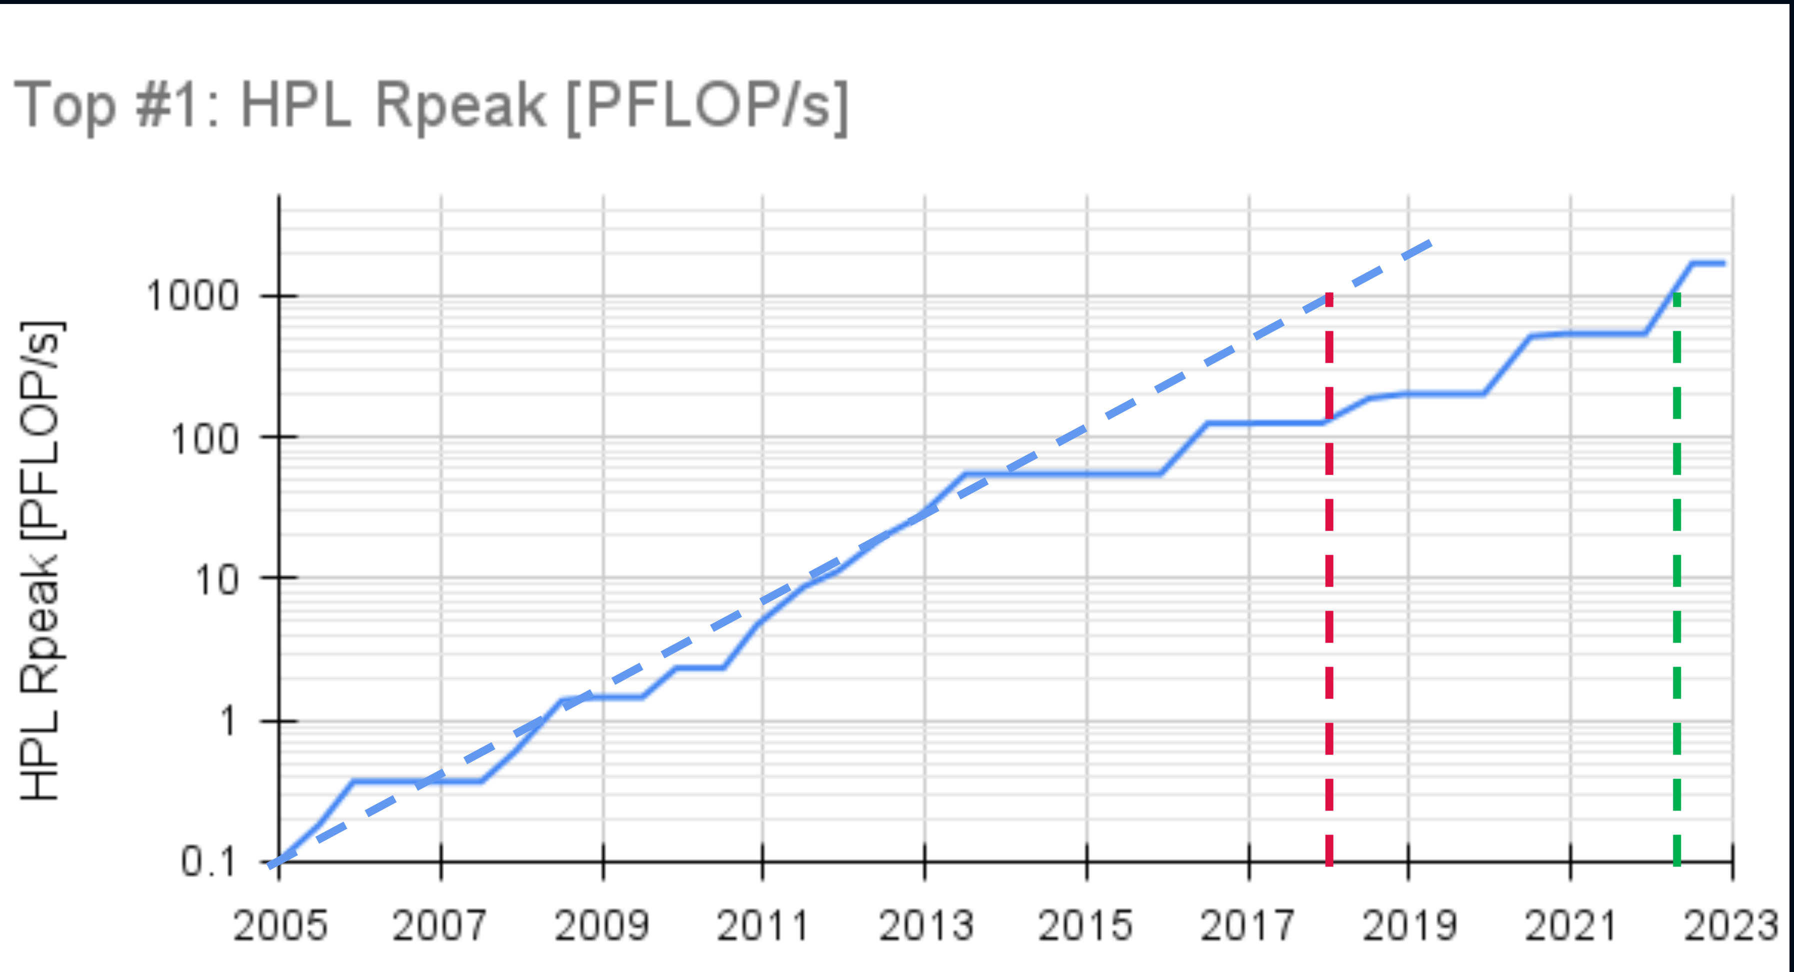

- **1997:** First **1 TFLOP/s** computer: (*Intel ASCI Read/9152*)
- **2008:** First **1 PFLOP/s** computer: (*Roadrunner*)
- So…. First **1 EFLOP/s** computer: **2018 !!**
  - Well… not really
- It took 4 more years… **2022**

Credit: Estela Suarez (Jülich Supercomputing Centre)

https://www.top500.org/

FRONTIER
First Exascale System

https://www.ornl.gov/news/ornl-celebrates-launch-frontier-worlds-fastest-supercomputer

OUTLINE

1. Evolution of Computing
2. **Machine Learning drives computational demands**
3. Supercomputing and Parallel Programming
4. Jülich Supercomputing Centre

THE ERA OF LARGE-SCALE DEEP LEARNING

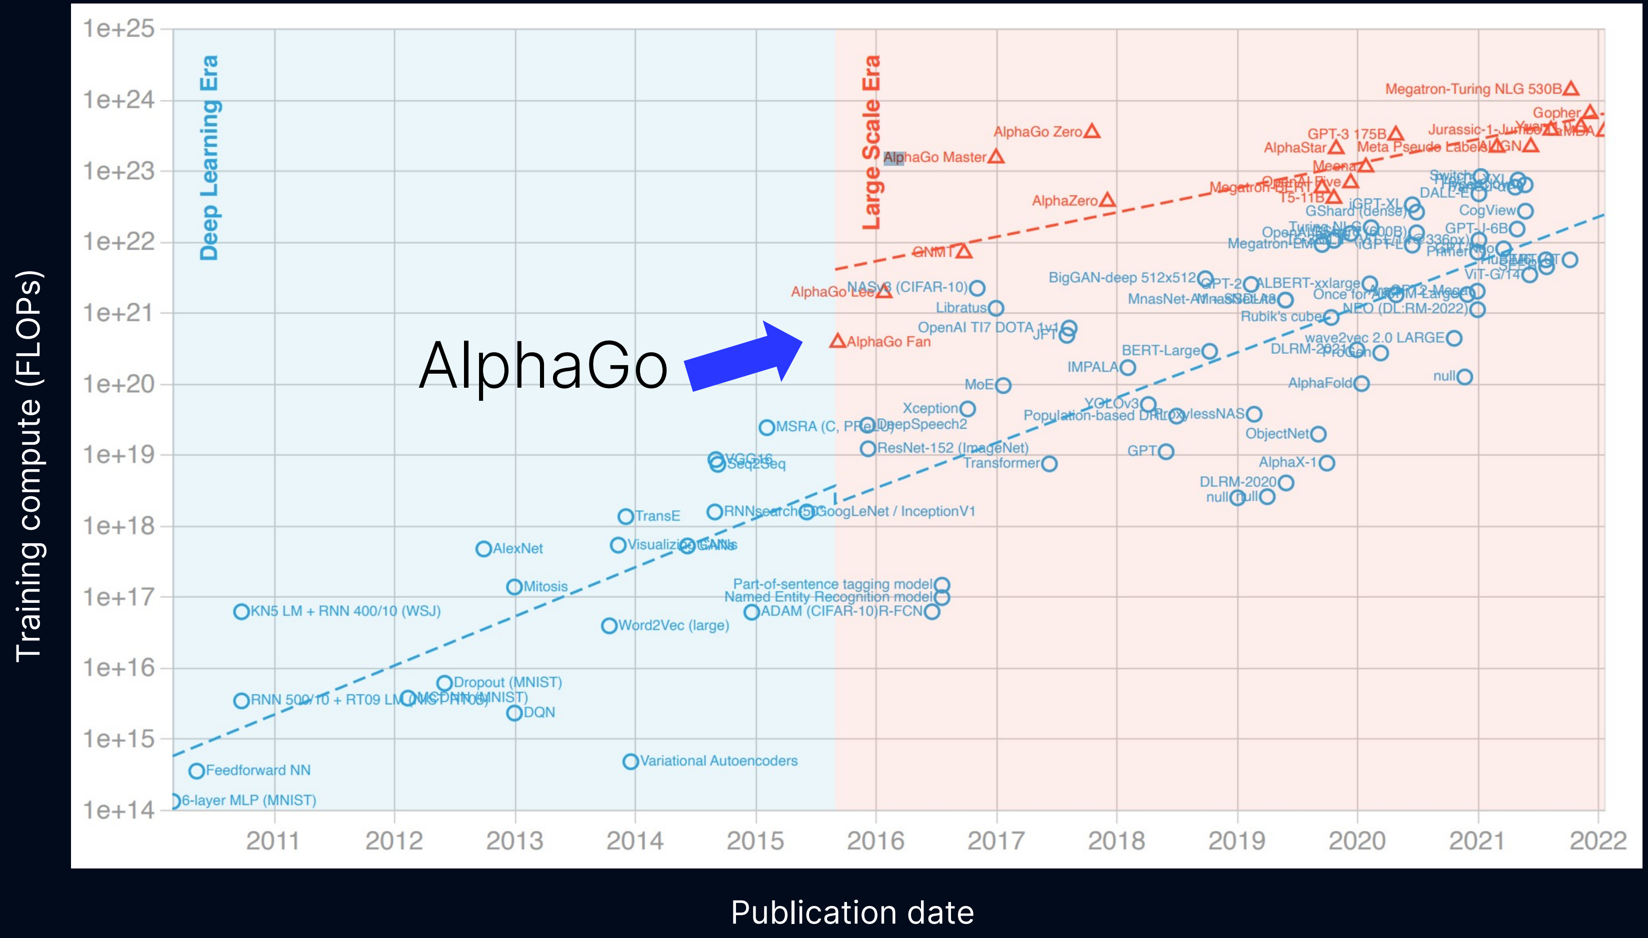

- **2015**: a new trend of large-scale models
- Computational capacity significantly higher (e.g., AlphaGo) than other models published in the same year
- OpenAI estimates 3.4-month doubling!

J. Sevilla, L. Heim, A. Ho, T. Besiroglu, M. Hobbhahn and P. Villalobos, "Compute Trends Across Three Eras of Machine Learning," 2022 International Joint Conference on Neural Networks (IJCNN), pp. 1-8, 2022, https://doi.org/10.1109/IJCNN55064.2022.9891914

NEW TRENDS IN AI FOUNDATION MODELS

Strong trend towards FMs trained on extensive domain-agnostic datasets, using:

- Unsupervised learning
- Self-supervised representation learning,
- Multimodal learning

Deliver more robust insights and decision-making, and bring advances in:

- Mainstream problems, e.g.: Natural Language Processing (NLP), Computer Vision
- But also to many scientific fields, e.g., Earth observation

Chat GTP-4 Could Pass the Bar Exam

How Our Technology Evolves FAST

Brian Lamacraft · Follow
Published in ILLUMINATION · 2 min read · Jan 5

AI chatbot's MBA exam pass poses test for business schools

ChatGPT earned a solid grade and outperformed some humans on a Wharton course

ChatGPT Passes Google Coding Interview for Level 3 Engineer With $183K Salary

'Amazingly, ChatGPT gets hired at L3 when interviewed for a coding position,' reads a Google document, but ChatGPT itself says it can't replicate human creativity and problem-solving skills.

By Emily Dreibelbis February 1, 2023

AI Passes U.S. Medical Licensing Exam

— Two papers show that large language models, including ChatGPT, can pass the USMLE

by Michael DePeau-Wilson, Enterprise & Investigative Writer, MedPage Today January 19, 2023

CURRENT POPULARITY OF SUPERCOMPUTERS

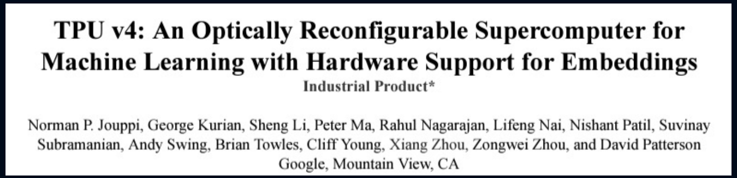
TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings
Industrial Product*

Norman P. Jouppi, George Kurian, Sheng Li, Peter Ma, Rahul Nagarajan, Lifeng Nai, Nishant Patil, Suvinay Subramanian, Andy Swing, Brian Towles, Cliff Young, Xiang Zhou, Zongwei Zhou, and David Patterson
Google, Mountain View, CA


Tech > Science

BABY STEPS Google artificial intelligence supercomputer creates its own 'AI child' that can outperform its human-made rivals

The NASNet system was created by a neural network called AutoML earlier this year

Mark Hodge
Published: 15:22, 5 Dec 2017 | Updated: 11:27, 6 Dec 2017


FORBES > INNOVATION > SUSTAINABILITY

Tesla's Biggest News At AI Day Was The Dojo Supercomputer, Not The Optimus Robot

James Morris Contributor
I write about the rapidly growing world of electric vehicles

Follow

0
Oct 6, 2022, 07:23am EDT

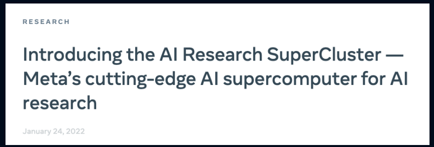
RESEARCH

Introducing the AI Research SuperCluster — Meta's cutting-edge AI supercomputer for AI research

January 24, 2022

- Big tech companies deploy their AI supercomputers
- Supercomputing now goes far beyond traditional scientific computing, which was driven by large governments
- Major industries building highly specialized supercomputers are taking the lead

N. P. Jouppi, G. Kurian, et al., "TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings", 2023, https://doi.org/10.48550/arXiv.2304.01433

https://www.forbes.com/sites/jamesmorris/2022/10/06/teslas-biggest-news-at-ai-day-was-the-dojo-supercomputer-not-the-optimus-robot/?sh=22ba4ab780bd

https://ai.facebook.com/blog/ai-rsc/

https://www.thesun.co.uk/tech/5072741/google-nasnet-ai-child-reinforcement-learning/

Torsten Hoefler, "Efficient AI: From supercomputers to smartphones", Scalable Parallel Computing Lab @ ETH Zurich, https://youtu.be/xxwT45ljG4o

HOW TO CREATE A FOUNDATION MODEL?

1) Gather data at scale

2) Train foundation model one time and evaluate

3) Fine-tune model for multiple downstream tasks

4) Inference (operational)

HOW TO CREATE A FOUNDATION MODEL?

Data for specific events

Big Data

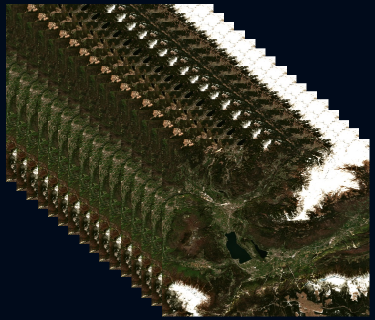

Satellite Image Time Series

MAIN TRAINING

Trained model

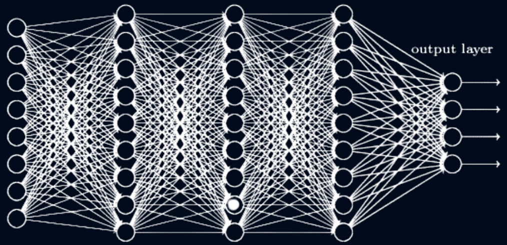

E.g. transformer

FINE TUNING

Fine-tuned model

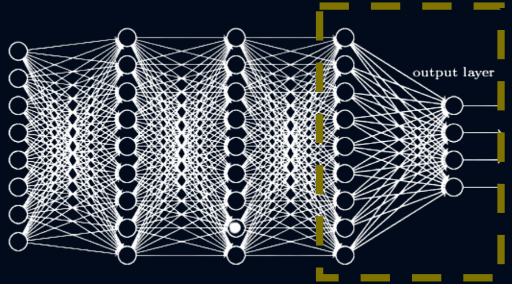

High Performance Computing

Cloud Computing, Computing workstations

TIME REQUIRED FOR A FULL TRAINING OF GPT-3 MODEL

175 billion weight parameters

Ranking in November 2020
(TOP500 (7 World, 1 Europe), Green500 (1 in TOP100) TOP10 AI (4)

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE
ParTec CLUSTER COMPETENCE CENTER
Mellanox TECHNOLOGIES
AtoS Bull
NVIDIA.

- 1× Nvidia Ampere100 ≈ 90 years
- 1× Nvidia Hopper100 ≈ 15-30 years
- 2,000× Nvidia Ampere100 ≈ 16 days
  (if scaled well on JUWELS Booster)

Julich Supercomputing Centre, "JUWELS Cluster and Booster: Exascale Pathfinder with Modular Supercomputing Architecture at Julich Supercomputing Centre," Journal of large-scale research facilities, vol. 7, no. A138, 2021.

S. Kesselheim, A. Herten, K. Krajsek, J. Ebert, J. Jitsev, M. Cherti, M. Langguth, B. Gong, S. Stadtler, A. Mozaffari, G. Cavallaro, R. Sedona, A. Schug, A. Strube, R. Kamath, M. G. Schultz, M. Riedel, and T. Lippert, "JUWELS Booster – A Supercomputer for Large-Scale AI Research," in High Performance Computing (H. Jagode, H. Anzt, H. Ltaief, and P. Luszczek, eds.), (Cham), pp. 453–468, Springer International Publishing, 2021.

Jenia Jitsev, Towards Scalable Deep Learning, Scalable Learning & Multi-Purpose AI Lab, Helmholtz AI, LAION @ JSC

TIME REQUIRED FOR A FULL TRAINING OF GPT-4 MODEL

1,8 Trillion weight parameters

Ranking in November 2020
(TOP500 (7 World, 1 Europe), Green500 (1 in TOP100) TOP10 AI (4)

- 1× Nvidia Ampere100 ≈ 1,200 years
- 1× Nvidia Hopper100 ≈ 200-600 years
- 2,000× Nvidia Ampere100 ≈ 900 days

Julich Supercomputing Centre, "JUWELS Cluster and Booster: Exascale Pathfinder with Modular Supercomputing Architecture at Julich Supercomputing Centre," Journal of large-scale research facilities, vol. 7, no. A138, 2021.

S. Kesselheim, A. Herten, K. Krajsek, J. Ebert, J. Jitsev, M. Cherti, M. Langguth, B. Gong, S. Stadtler, A. Mozaffari, G. Cavallaro, R. Sedona, A. Schug, A. Strube, R. Kamath, M. G. Schultz, M. Riedel, and T. Lippert, "JUWELS Booster – A Supercomputer for Large-Scale AI Research," in High Performance Computing (H. Jagode, H. Anzt, H. Ltaief, and P. Luszczek, eds.), (Cham), pp. 453–468, Springer International Publishing, 2021.

Jenia Jitsev, Towards Scalable Deep Learning, Scalable Learning & Multi-Purpose AI Lab, Helmholtz AI, LAION @ JSC

OUTLINE

1. Evolution of Computing
2. Machine Learning drives computational demands
3. **Supercomputing and Parallel Programming**
4. Jülich Supercomputing Centre

WHAT ACTUALLY IS A SUPERCOMPUTER?

WHY USE PARALLEL COMPUTERS?

- Parallel computers can be the only way to achieve specific computational goals in a given time
- Sequential system is too "slow"
  - Calculation takes days, weeks, months, years, ...
  - **Use more than one processor to get calculation faster**
- Sequential system is too "small"
  - Data does not fit into the memory
  - **Use parallel system to get access to more memory**

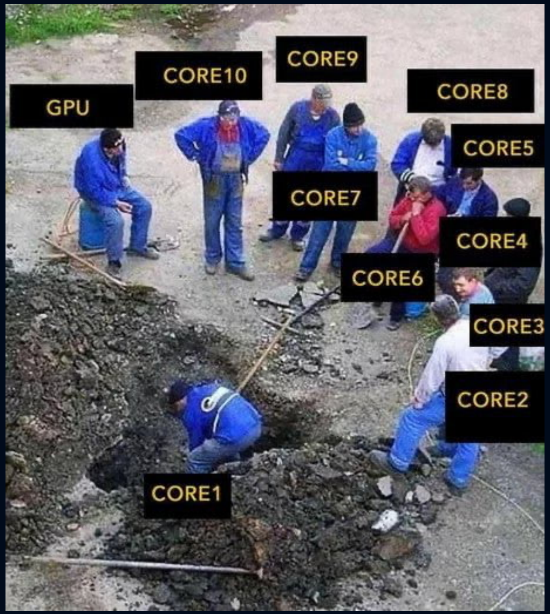

https://9gag.com/gag/av5vmzd

- You realize you have a parallel system (**multicore**) and you want to make use of its special features

Credit: Bernd Mohr (Jülich Supercomputing Centre)

HIGH-PERFORMANCE COMPUTING SYSTEMS

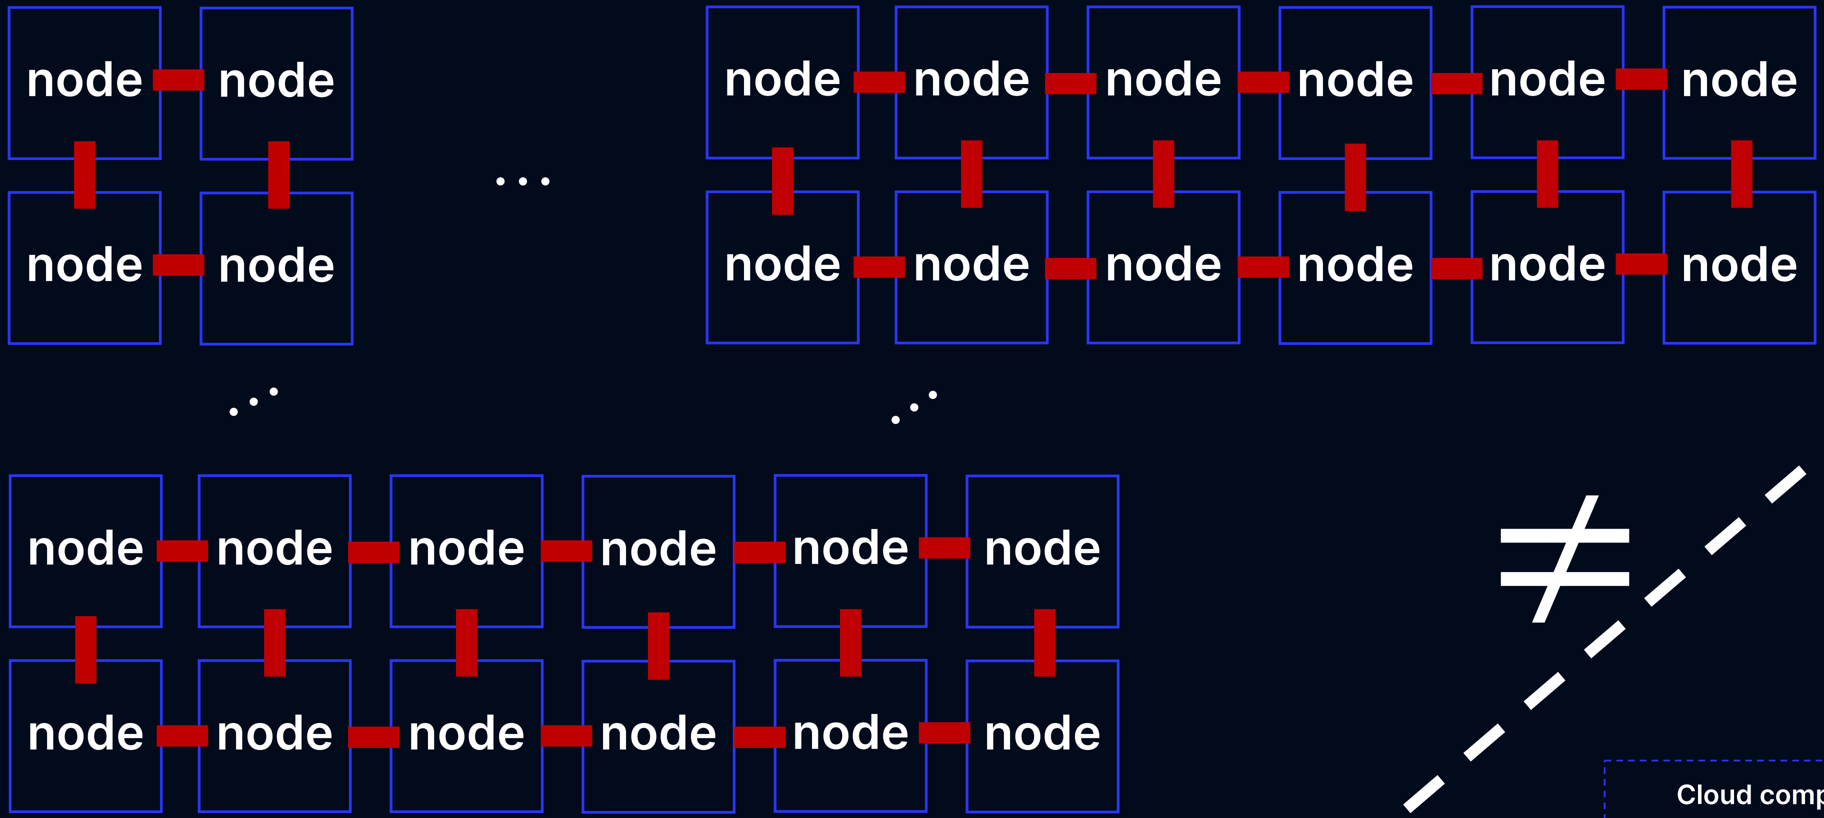

- High number of **compute nodes**
- Vast amounts of memory
- **High-speed interconnects**

≠

High Throughput Computing

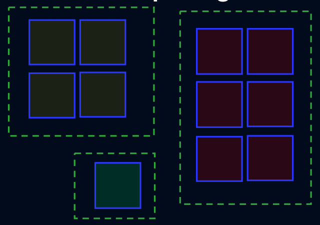

Many independent computations

Cloud computing

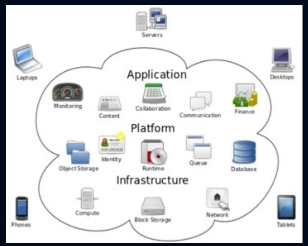

Virtual machines in compute center(s)

Computer cluster

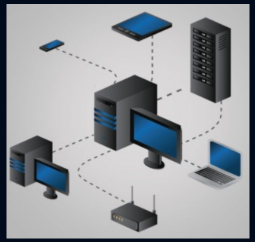

Many computers bound together locally

HPC BUILDING BLOCKS

- Hardware
  - Login and compute nodes (CN)
  - Network
  - Storage
- Software
  - Operating System (OS)
  - Compilers
  - Libraries
  - Scheduler

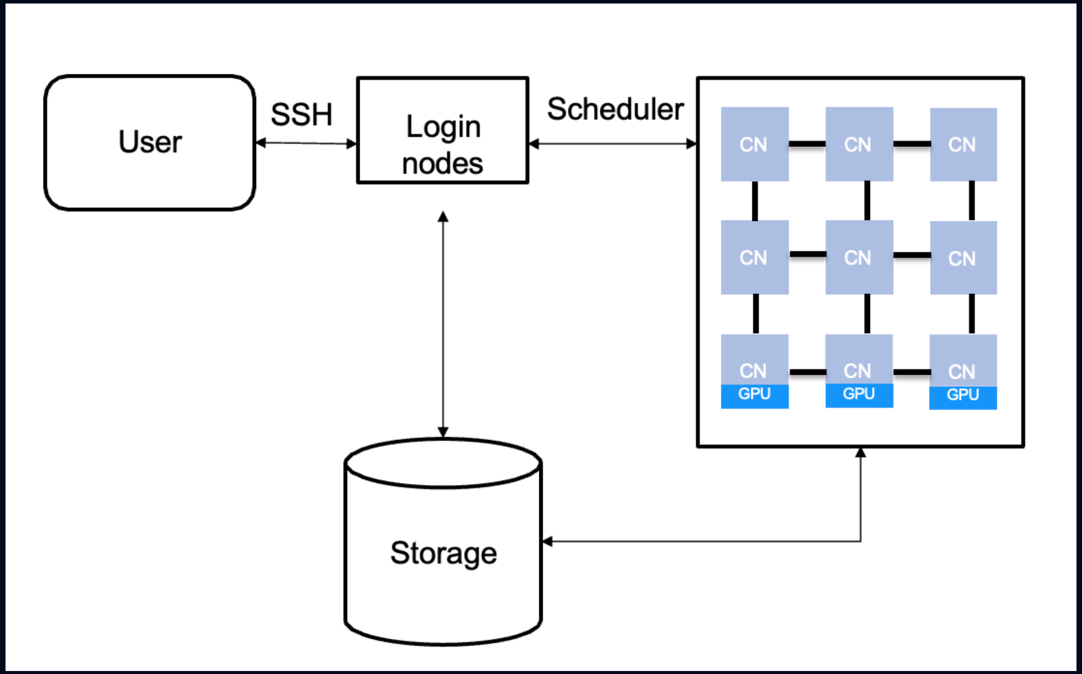

Credit: Bernd Mohr (Jülich Supercomputing Centre)

USING THE SUPERCOMPUTER MEANS SUBMITTING A JOB TO A BATCH SYSTEM

Job scheduling according to priorities. The jobs with the highest priorities will be scheduled next.

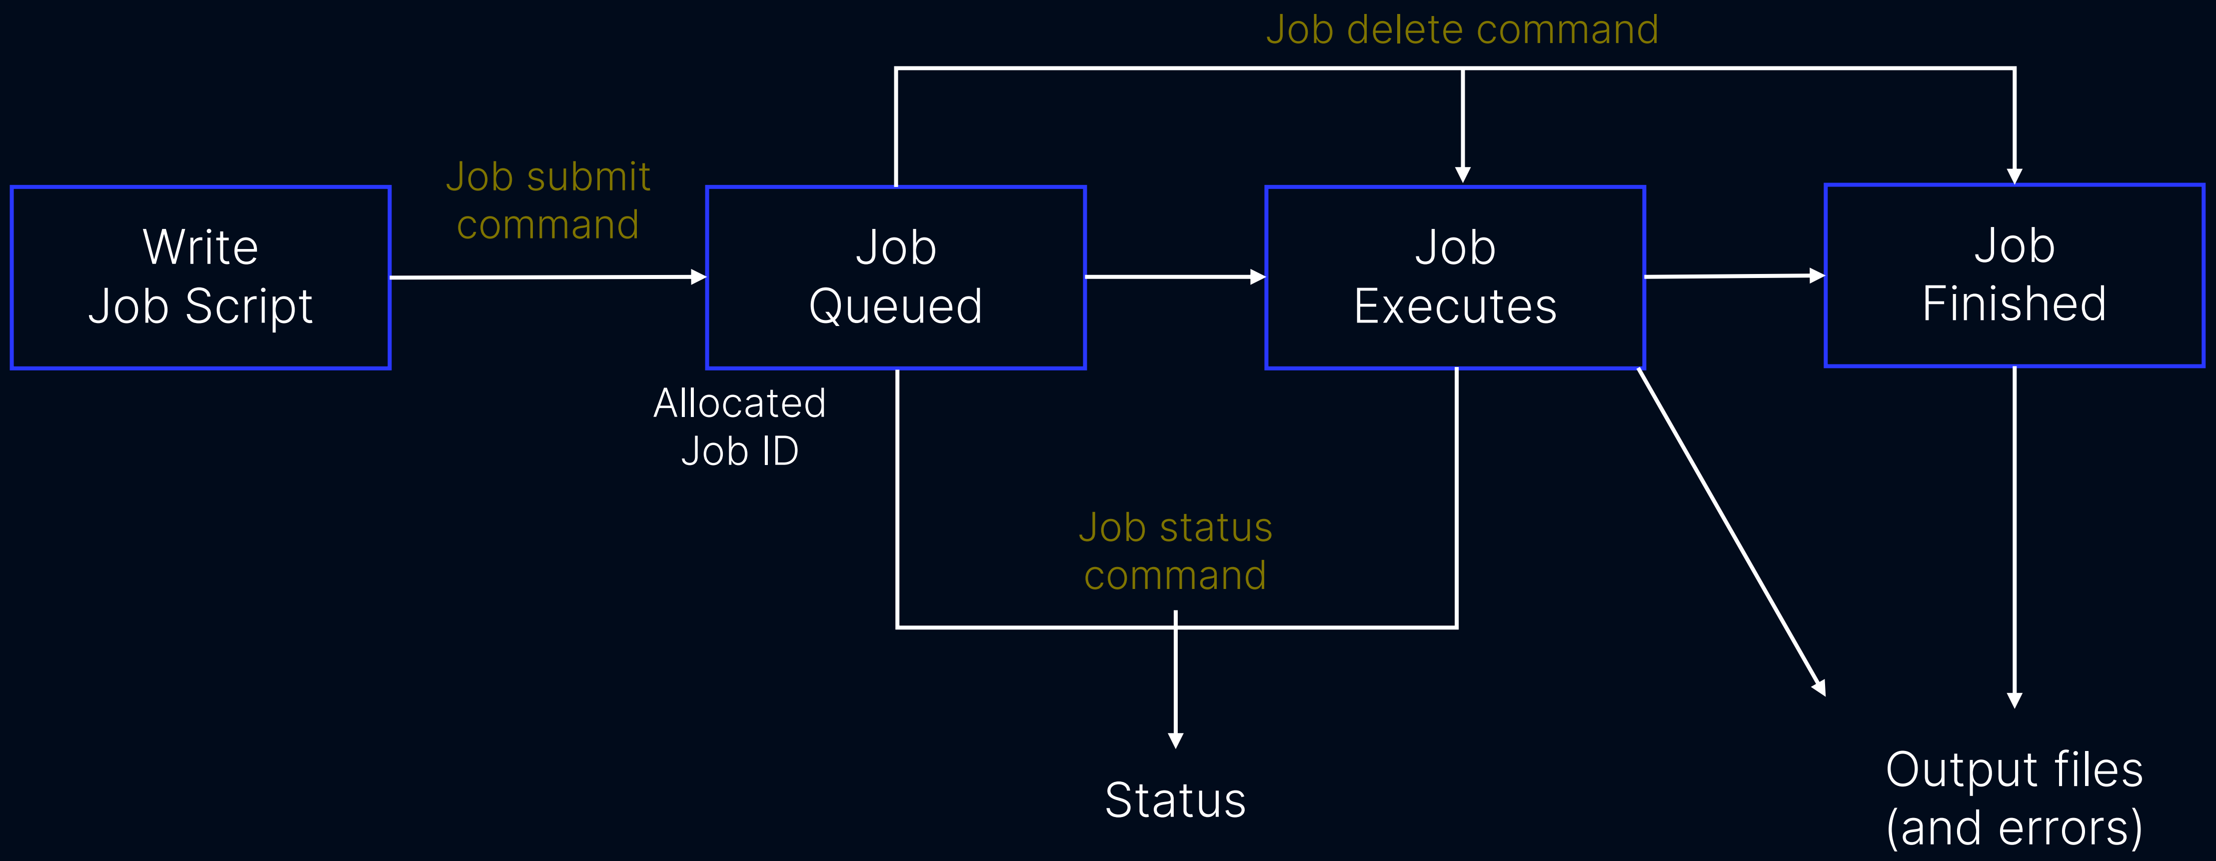

JUPYTER-JSC WEBSERVICE

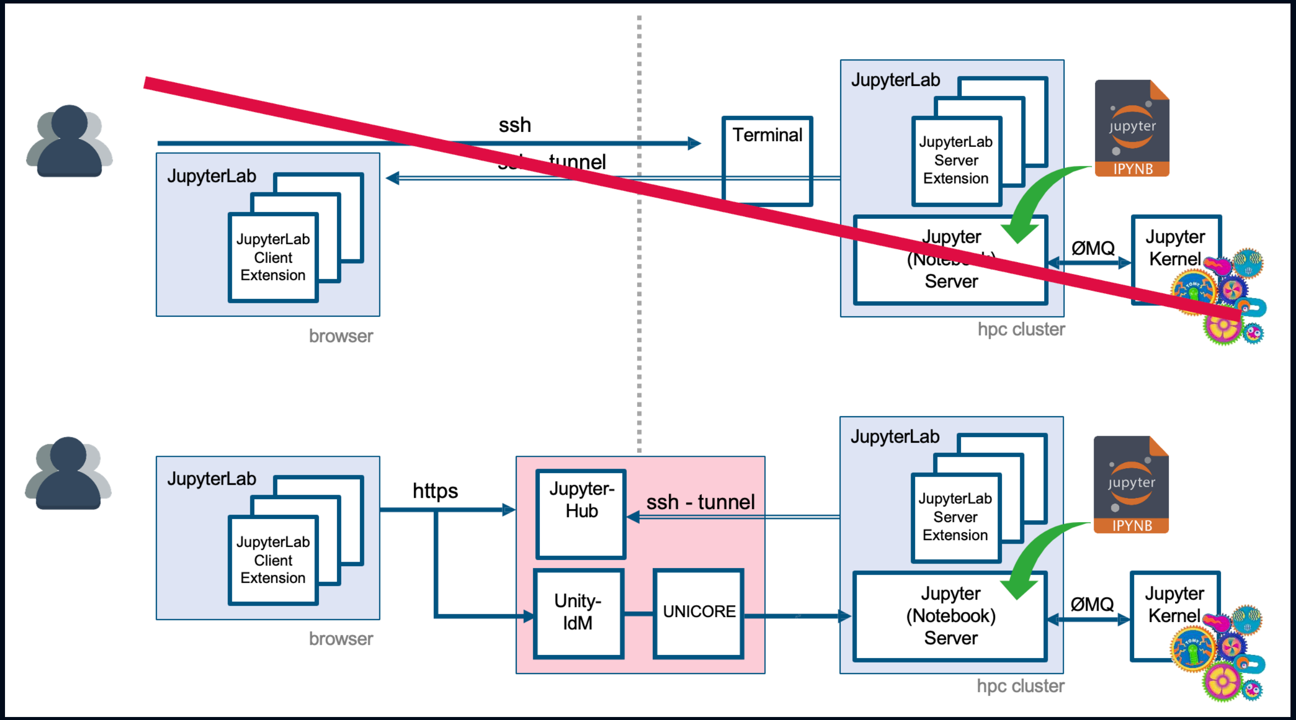

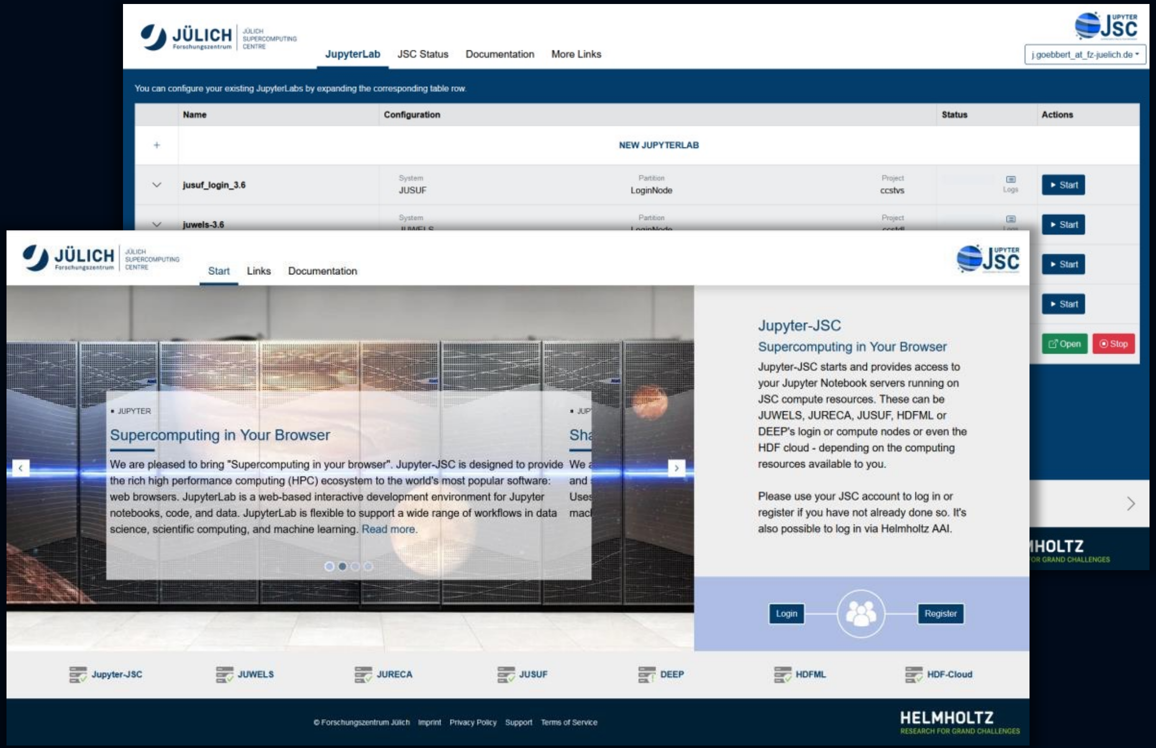

Credit: Jens H. Göbbert Tim Kreuzer, (Jülich Supercomputing Centre)
https://jupyter.jsc.fz-juelich.de/

DISTRIBUTED DEEP LEARNING

Paradigms of Parallelism

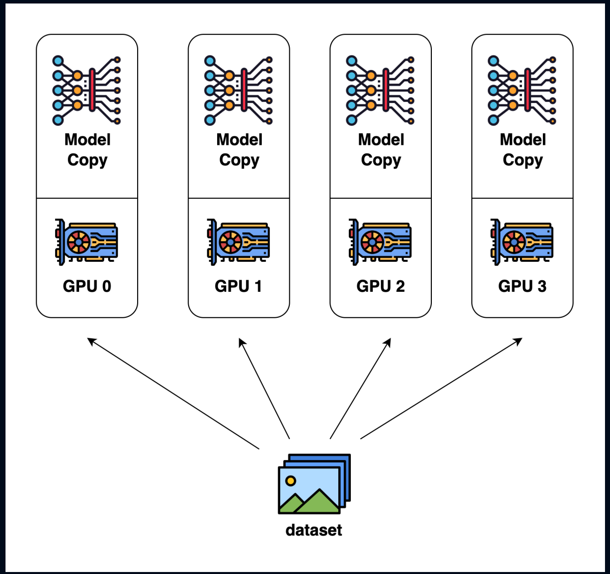

Data Parallel

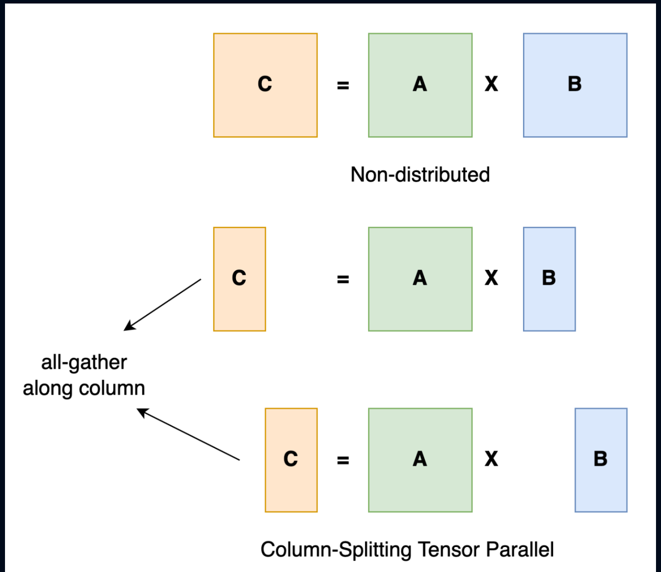

Model Parallel

...

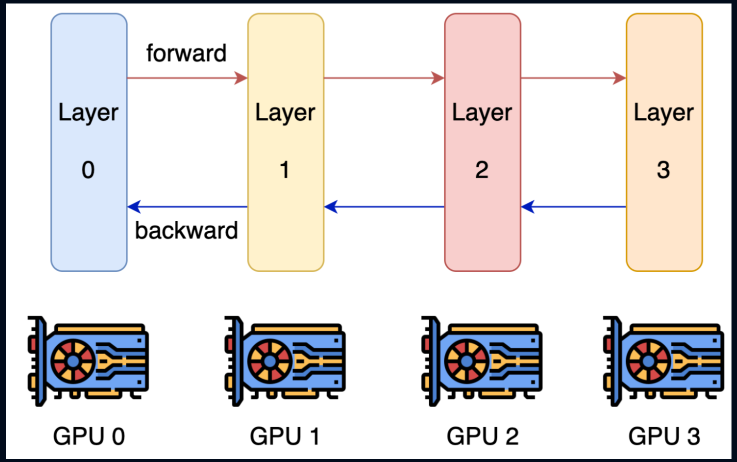

Pipeline Parallel

Credit: Colossal-AI, https://colossalai.org/docs/concepts/paradigms_of_parallelism/

OUTLINE

1. Evolution of Computing

2. Machine Learning drives computational demands

3. Supercomputing and Parallel Programming

4. **Jülich Supercomputing Centre**

JÜLICH SUPERCOMPUTING CENTRE (JSC)

Multi-system facility - Tier-0/1 HPC resources of the highest performance class

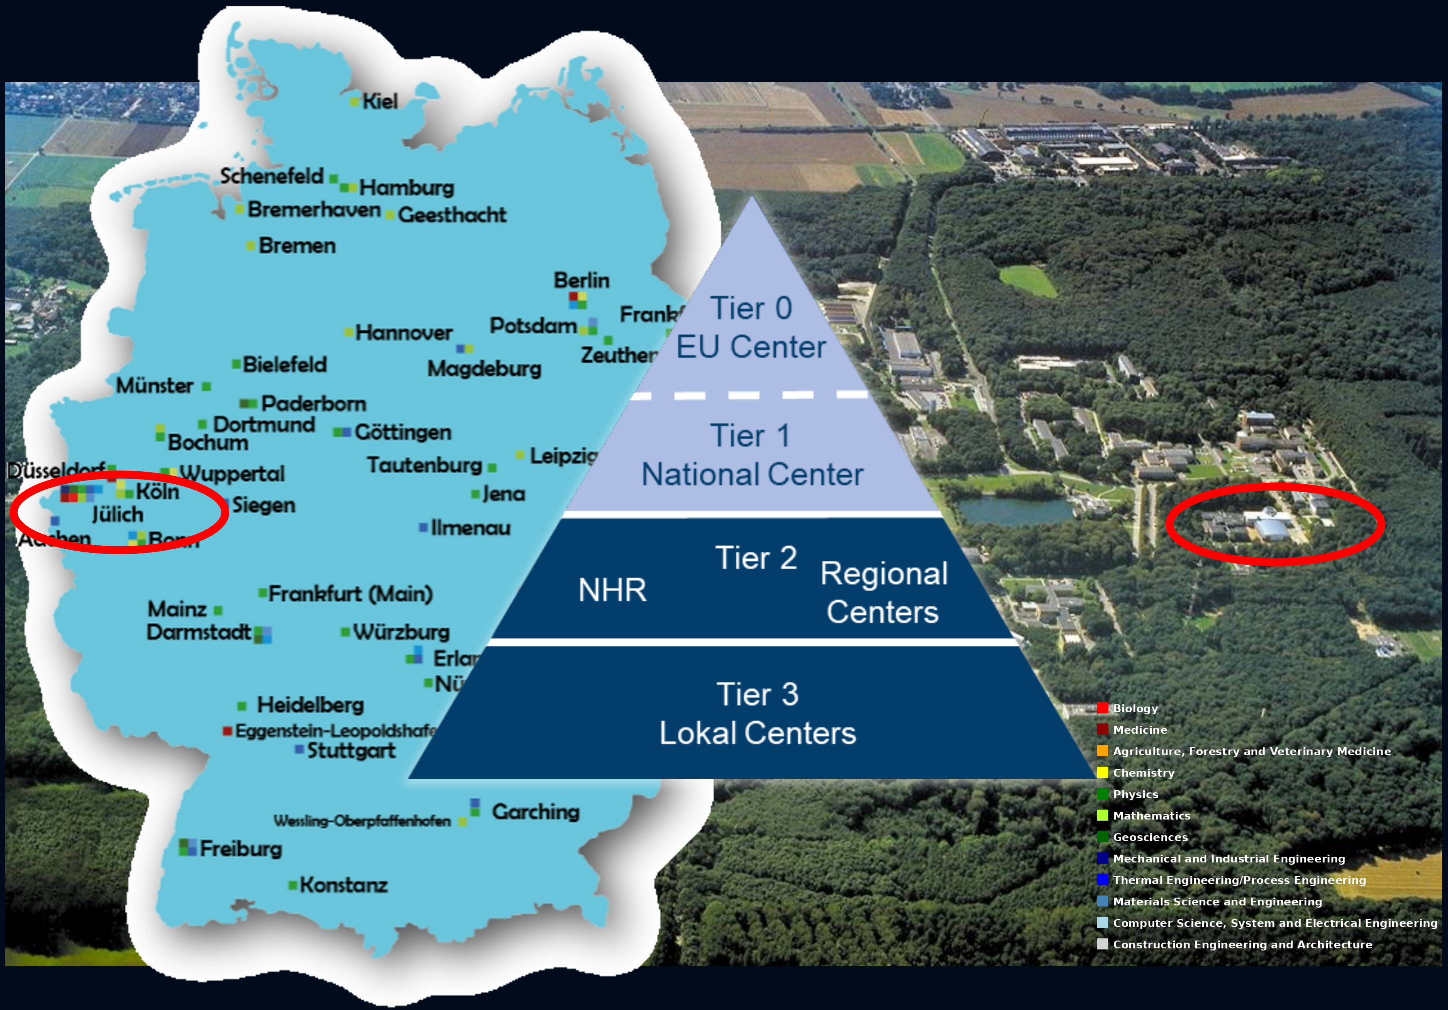

Jülich Supercomputing Centre (JSC), https://www.fz-juelich.de/en/ias/jsc

(DUAL) HARDWARE STRATEGY AT JSC

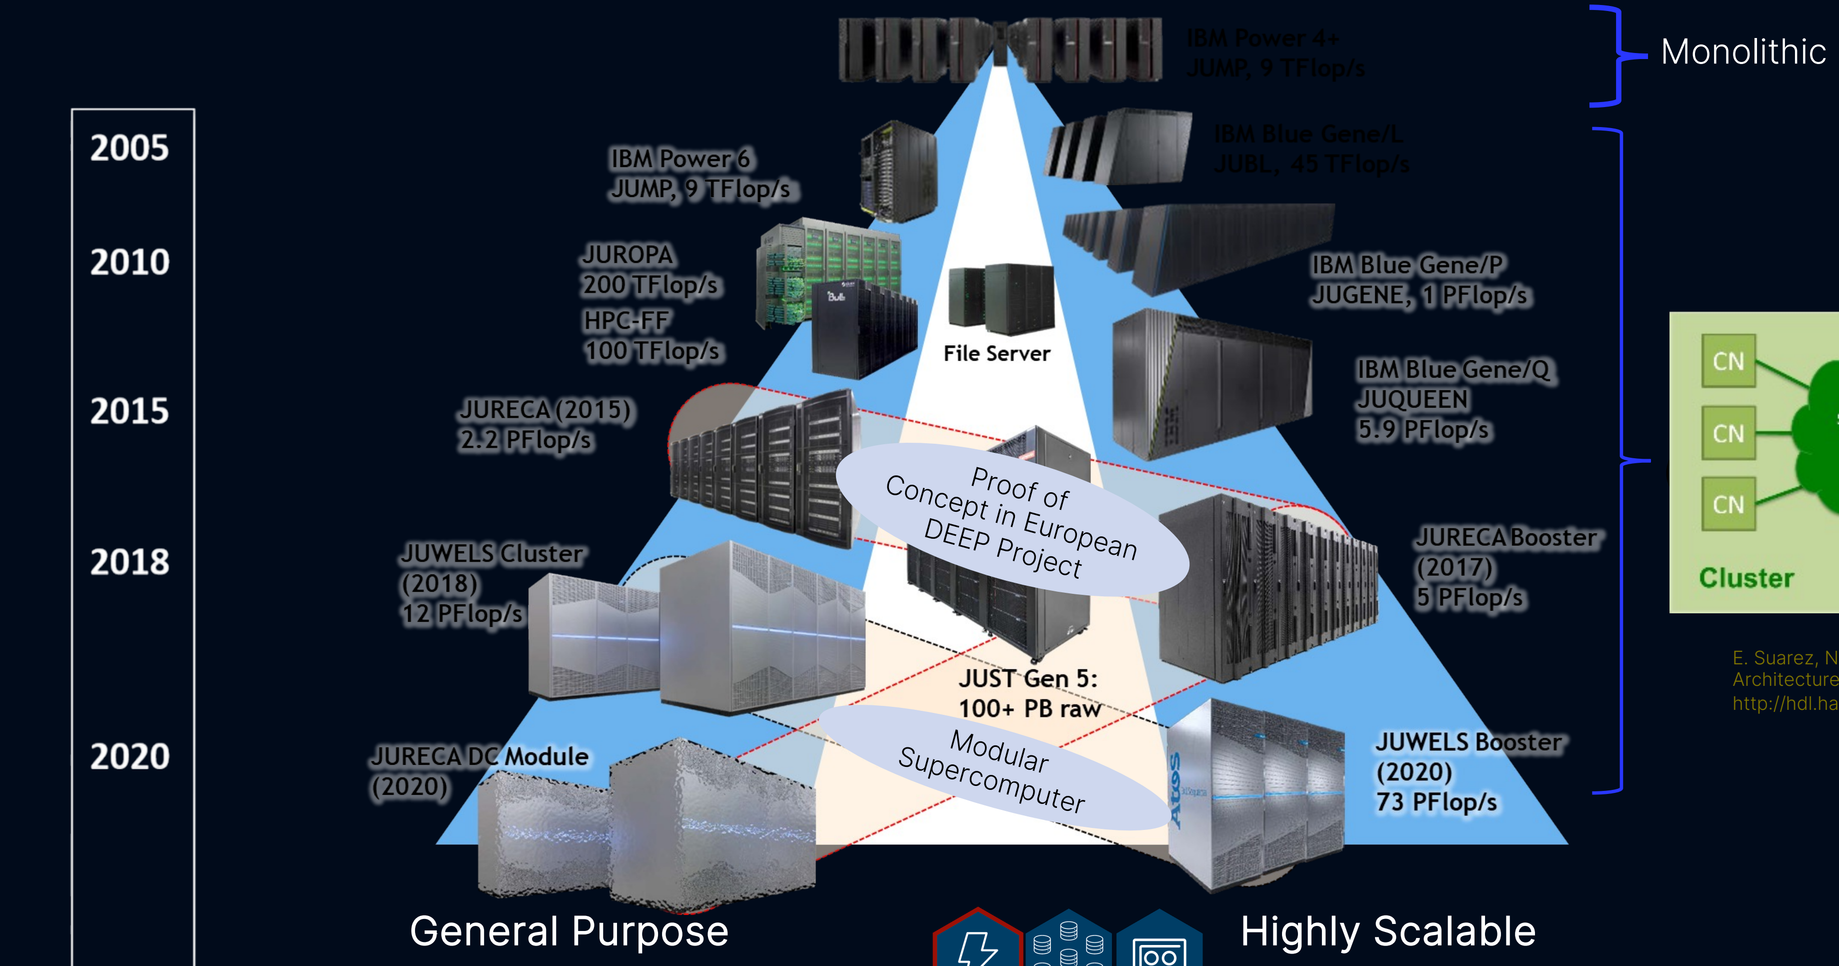

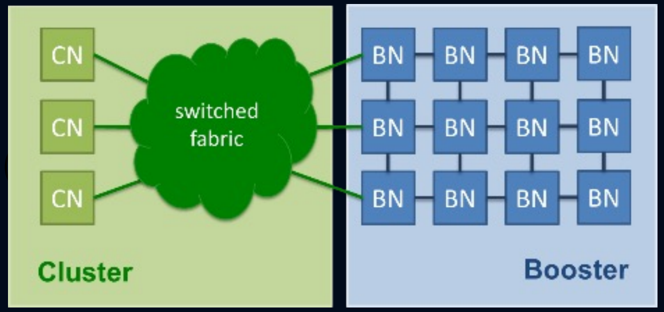

E. Suarez, N. Eicker, and T. Lippert, "Modular Supercomputing Architecture: from Idea to Production", 2019
http://hdl.handle.net/2128/22212

DEEP Projects
https://www.deep-projects.eu/

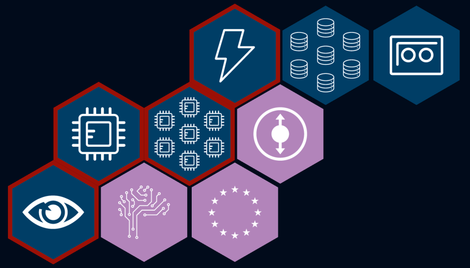

Exascale Modular Supercomputer

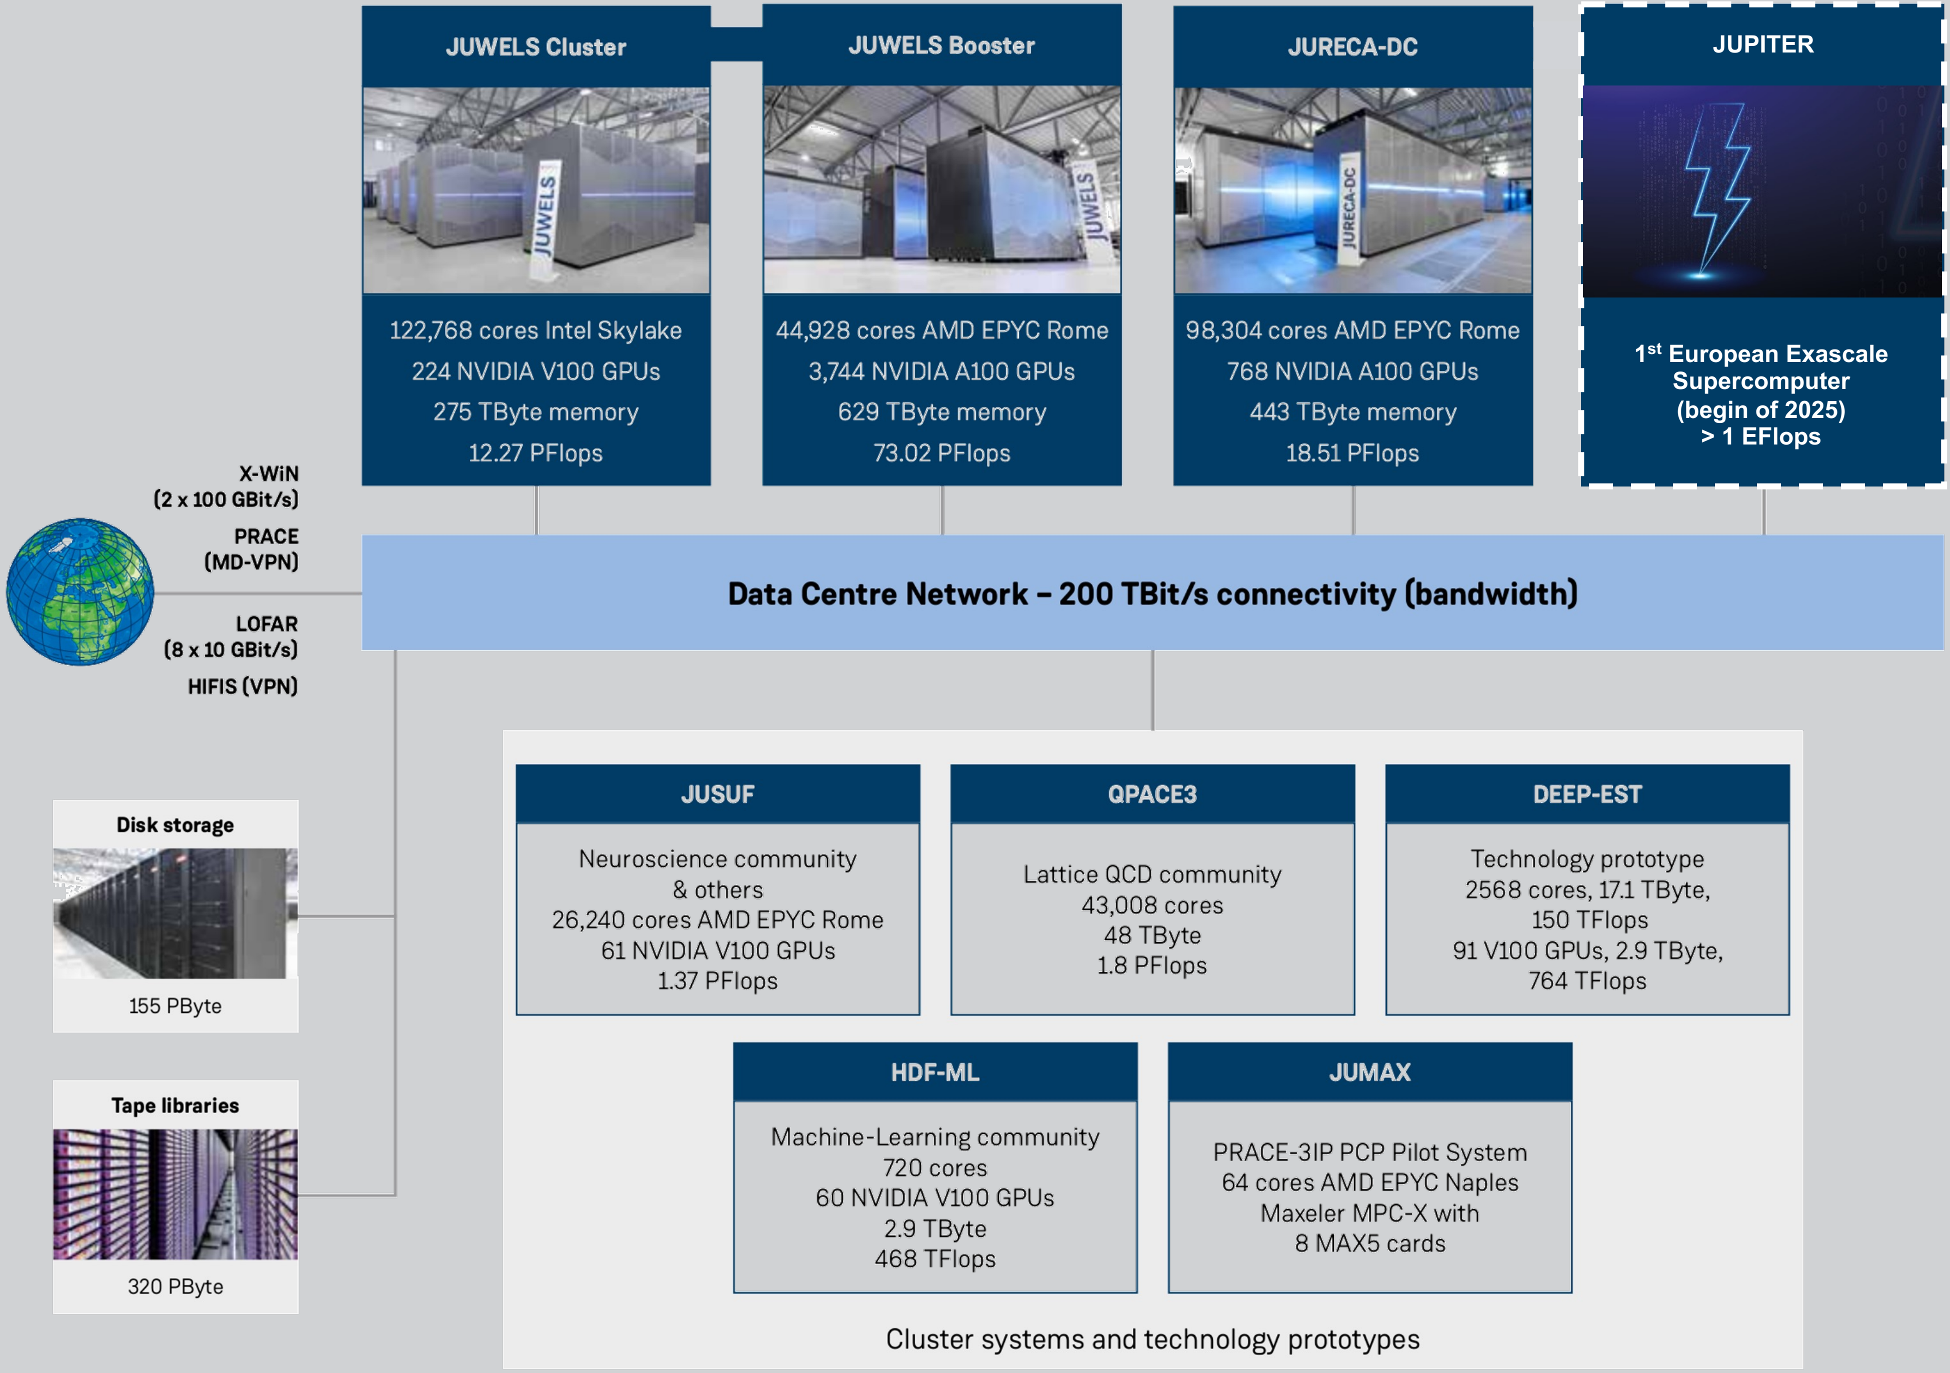

Credit: Jens Henrik Göbbert (Jülich Supercomputing Centre)

JURECA-DC

DC = Data Centric

- Intended for mixed capacity and capability workloads
- Designed with big-data science needs in mind

JURECA-DC, Jülich Supercomputing Centre (JSC), https://www.fz-juelich.de/en/ias/jsc/systems/supercomputers/jureca

JURECA-DC CPU NODES

- 576 compute nodes Atos
  - 2× 64-core AMD Epyc 7742 Rome CPUs AMD&
    - 2x 8 memory channels
    - 2x 256 GB DDR4 @ 3.2 GHz
      - 96 nodes with 2x 512 GB DDR4 @ 3.2 GHz
  - 2x 4 NUMA domains
  - PCIe Gen4 • 1x HDR100 InfiniBand adapter (100Gbps)

Credit: D. Alvarez, S. Achilles, Jülich Supercomputing Centre (JSC)
https://www.fz-juelich.de/en/ias/jsc/systems/supercomputers/jureca

JURECA-DC GPU NODES

- 192 compute nodes Atos
  - 2× 64-core AMD Epyc 7742 Rome CPUs AMD
    - 2x 8 memory channels
    - 2x 256 GB DDR4 @ 3.2GHz
    - 96 PCIe Gen4 lanes
  - 512 GB DDR memory
  - 4x Nvidia A100 GPUs NVIDIA
    - 9.7 / 19.5 TF/s peak
    - 40 GB HBM2
    - 1.5 TB/s memory performance
    - NVLink3 full mesh
      - 4 links (200GB/s) between GPU pairs
  - PCIe Gen4 x32 (64 GB/s)
- 2x HDR200 InfiniBand adapter (1 per GPU) NVIDIA

Credit: D. Alvarez, S. Achilles, Jülich Supercomputing Centre (JSC)
https://www.fz-juelich.de/en/ias/jsc/systems/supercomputers/jureca

JURECA-DC LOGIN NODES

- 12 login nodes
  - 2x 64-core AMD Epyc 7742 Rome CPUs
  - 1024 GB DDR4 @ 3.2 GHz
  - 100 GigE external network
  - 2x Nvidia RTX8000 GPUs
    - Different compute capabilities than in compute nodes!
  - Used for:
    - Compile/submit jobs
      - Careful with make -j !
    - Small pre- and post-processing/visualization
  - Shared nodes!

Credit: D. Alvarez, S. Achilles, Jülich Supercomputing Centre (JSC)
https://www.fz-juelich.de/en/ias/jsc/systems/supercomputers/jureca

JUPITER

The Arrival of Exascale in Europe

- JUPITER: EuroHPC JU system hosted at JSC
- Launched with focus on applications
- ~6000 nodes,
  24 000 H100 GPUs, 1 728 000 Arm cores, 24 000 NDR200 endpoints
- Landing in Modular Data Center
- Preparation is ongoing:
  - JUREAP
  - GH200 test systems

→ jupiter.fz-juelich.de

The GREEN 500 CERTIFICATE

JEDI - BullSequana XH3000, Grace Hopper Superchip 72C 3GHz, NVIDIA GH200 Superchip, Quad-Rail NVIDIA InfiniBand NDR200

EuroHPC/FZJ, Germany

is ranked

No. 1

among the World's TOP500 Supercomputers

with 72.733 GFlops/watts Performance

in the Green500 List published at the ISC24

Conference on June 01, 2024.

Congratulations from the Green500 Editors

Wu-chun Feng
Virginia Tech

Kirk Cameron
Virginia Tech

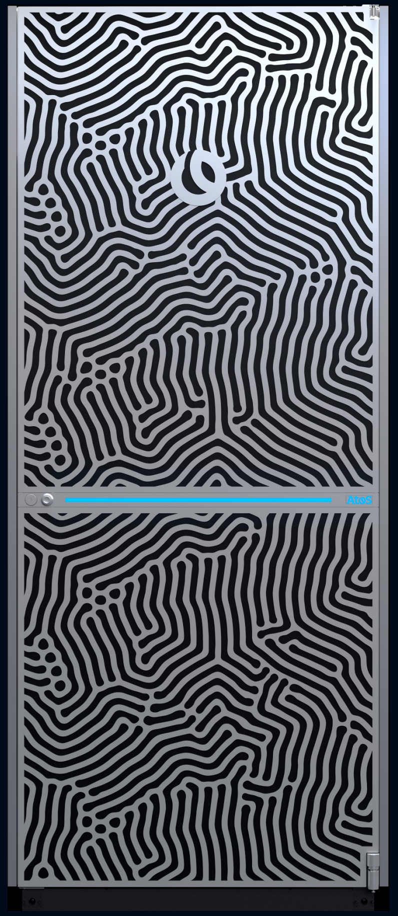

Funding Agencies

EuroHPC Joint Undertaking

Federal Ministry of Education and Research

Ministerium für Innovation, Wissenschaft und Forschung des Landes Nordrhein-Westfalen

Thank you for your attention