Leibniz Supercomputing Centre

of the Bavarian Academy of Sciences and Humanities

Judgment Day for HPC: Will AI Take Control?

Prof. Dr. Dieter Kranzlmüller
Leibniz Supercomputing Centre (LRZ) &
Ludwig-Maximilians-Universität (LMU) Munich

With contributions from
Michael Ott, Axel Auweter,
Josef Weidendorfer,
Nicolay Hammer, Max Höb,
and many others …

JUDGMENT DAY FOR HPC: WILL AI TAKE CONTROL?

Leibniz Supercomputing Centre (LRZ)
of the Bavarian Academy of Sciences and Humanities

Partner for Digital Transformation of Science & Research

IT-Services for Science since 1962

Munich Universities, Bavaria, Germany, Europe, …

The World‘s Largest Turbulence Simulation

Astro physics: How stars are born

Bavarian-Canadian Project ClimEx

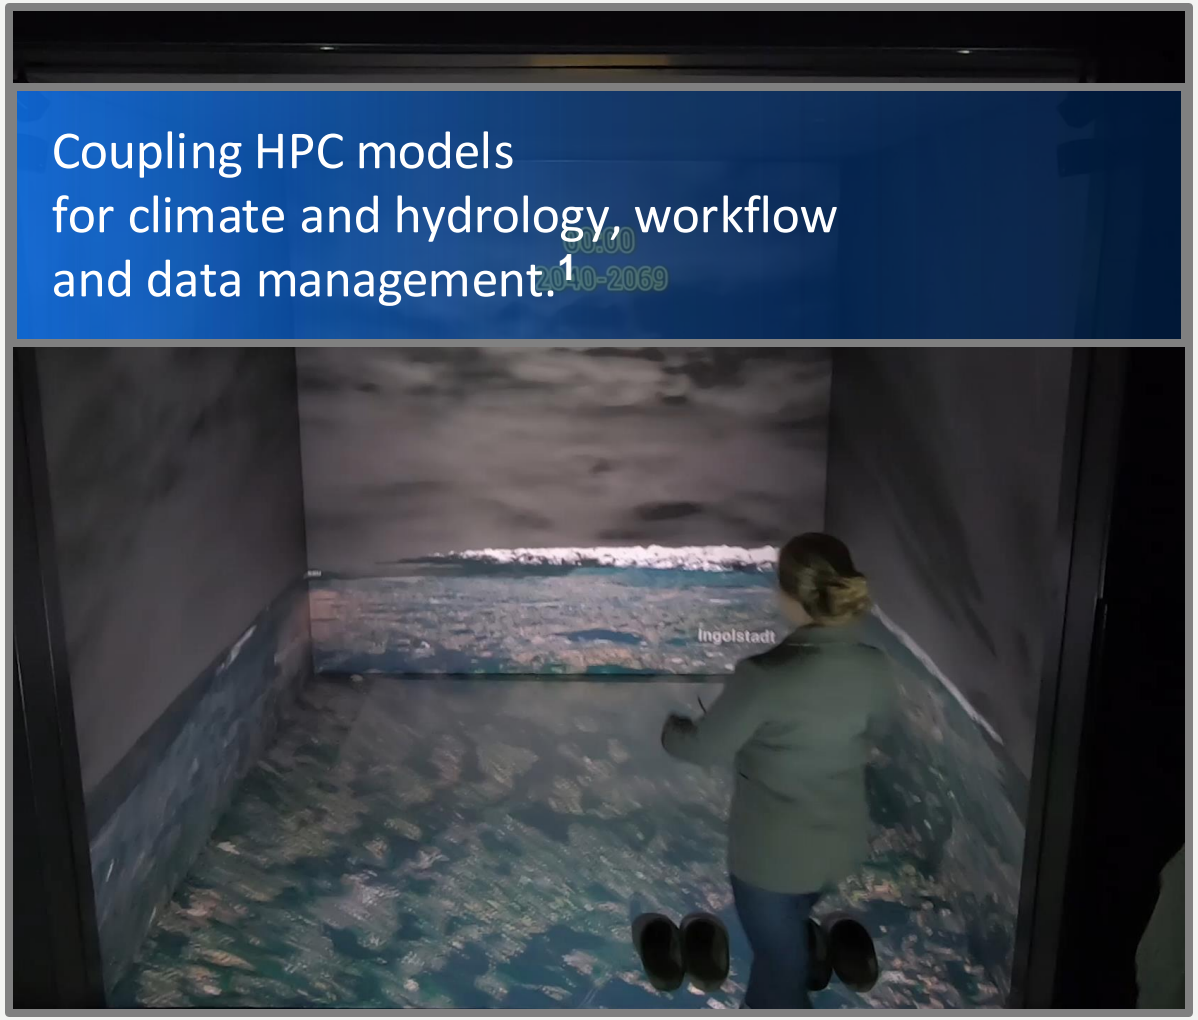

Coupling HPC models
for climate and hydrology, workflow
and data management.[1]

Impact of the climate crisis on extreme hydro-
logical events such as floods and droughts.[2]

www.climex-project.org; **1** © LRZ / Ouranos; **2** Danube Flooding 2013 © sueddeutsche.de

ClimEx – Scales of application

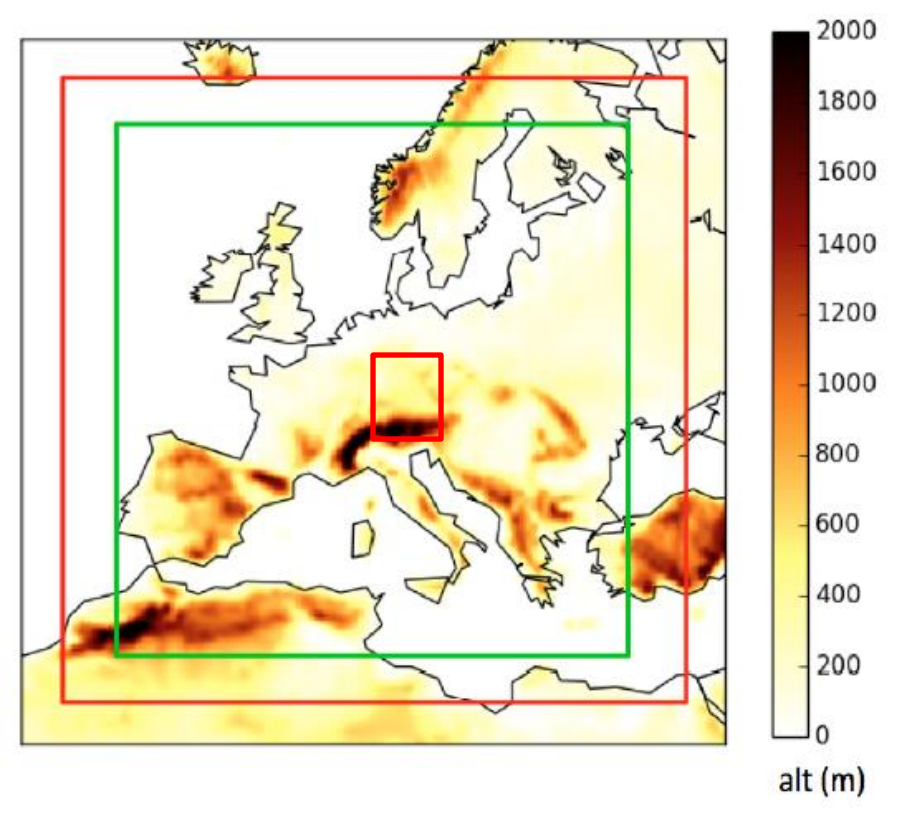

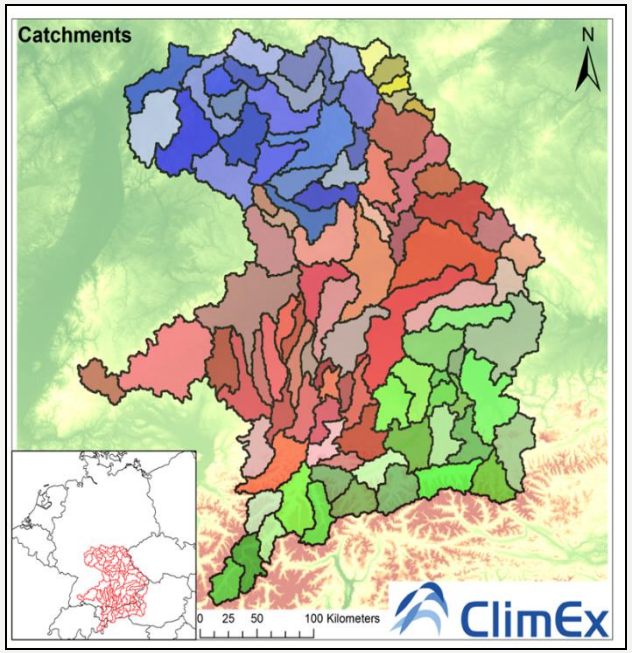

Scope: Assess the CC related dynamics of extremes events from the continental to the river basin scale

Goal: a) Improve process understanding of non-linear hydro-meteorological extreme events
b) Provide adaptation and management options for stakeholders to reduce related risks

Will extreme flows/floods (e.g. HQ100) be more severe?

Changes in Frequency and Intensity of HQ100 in Bavarian river basins

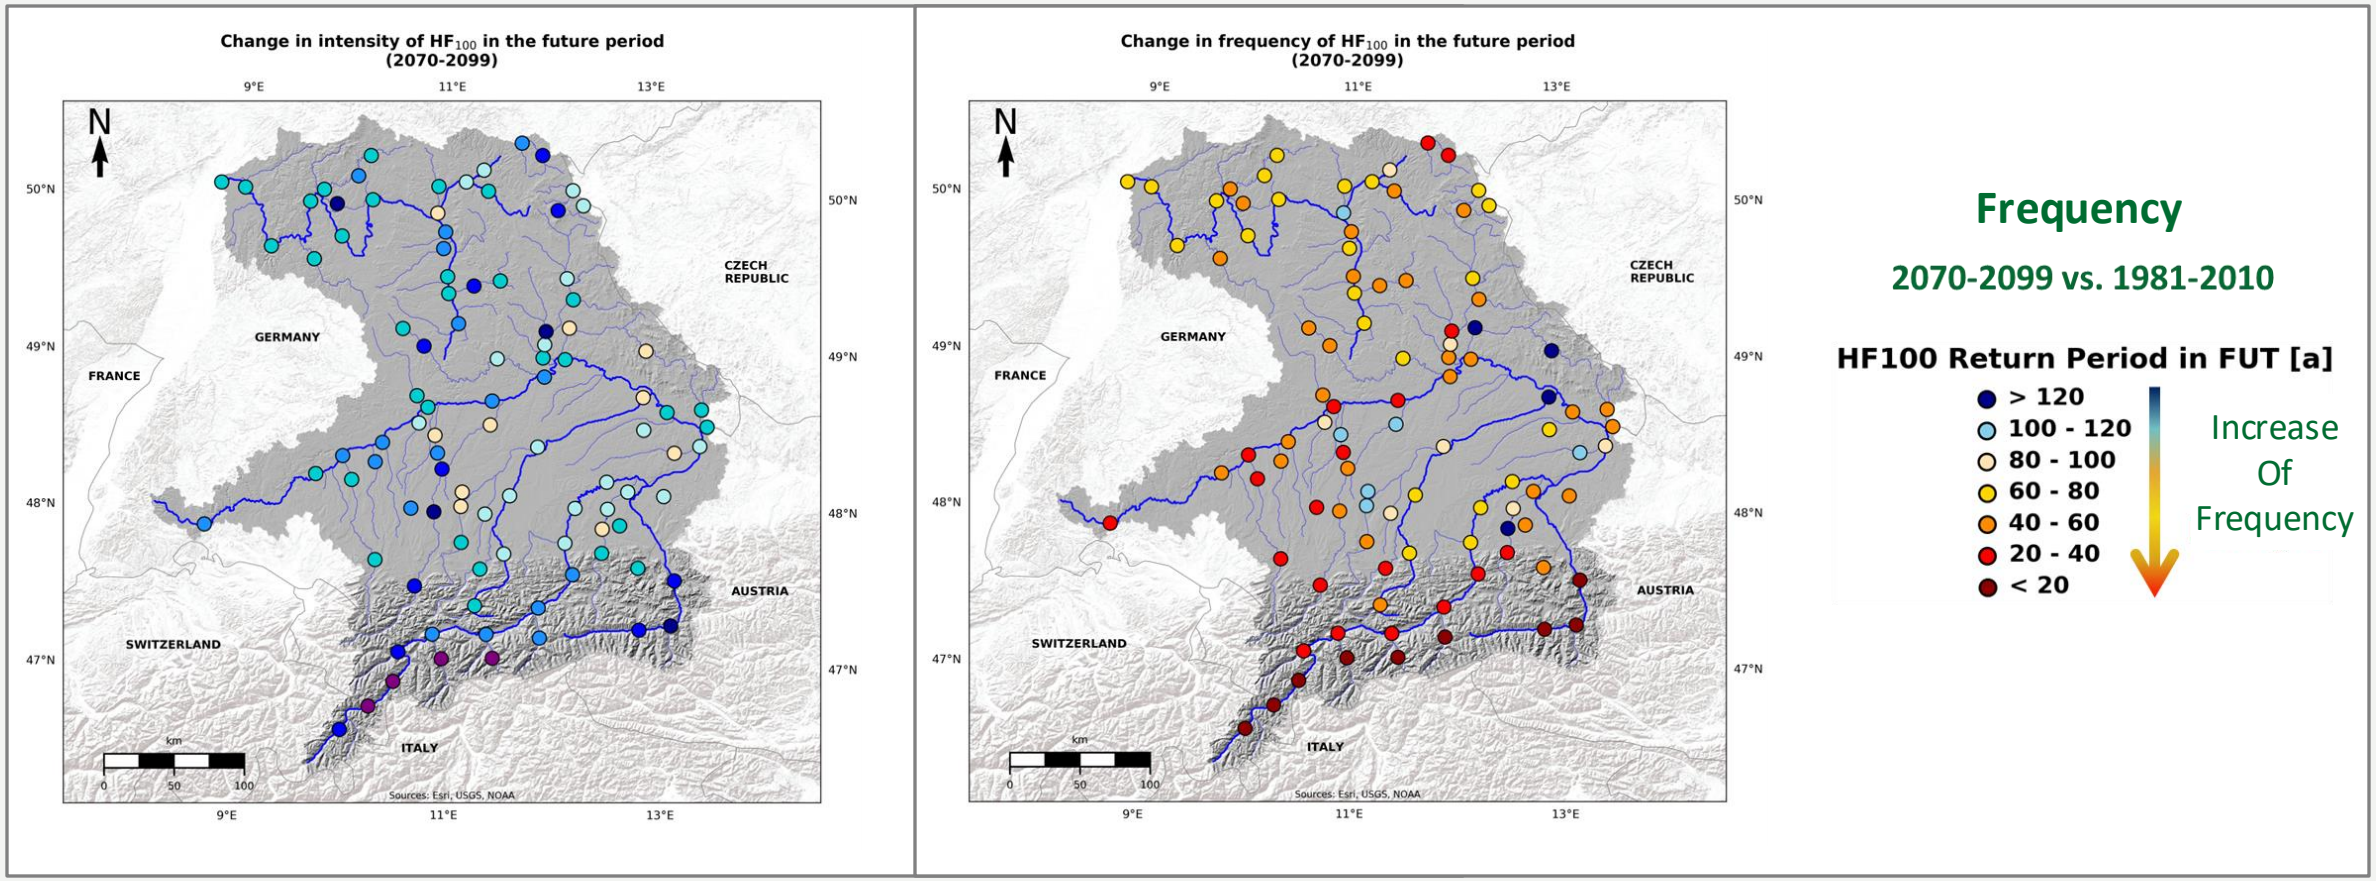

(Willkofer, Ludwig et al., in prep.)

Background…

From the 6th Assessment Report of the IPCC (AR6) (2021/2022):
Working Group I: The Physical Science Basis
Working Group II: Impacts, Adaptation and Vulnerability

WGI.B – Possible Climate Futures

B.2 Many changes in the climate system become larger in direct relation to increasing global warming. They include increases in the frequency and intensity of hot extremes, marine heatwaves, and heavy precipitation, agricultural and ecological droughts in some regions […]

WGII.C – Future Adaptation Options and their Feasibility

C.2 […] The effectiveness of adaptation to reduce climate risk is documented for specific contexts, sectors and regions (high confidence) and will decrease with increasing warming. Integrated, multi-sectoral solutions […] increase the feasibility and effectiveness of adaptation in multiple sectors.

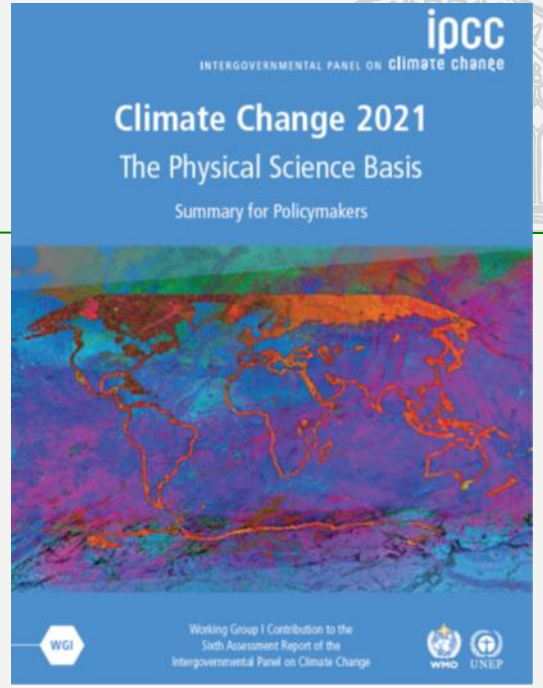

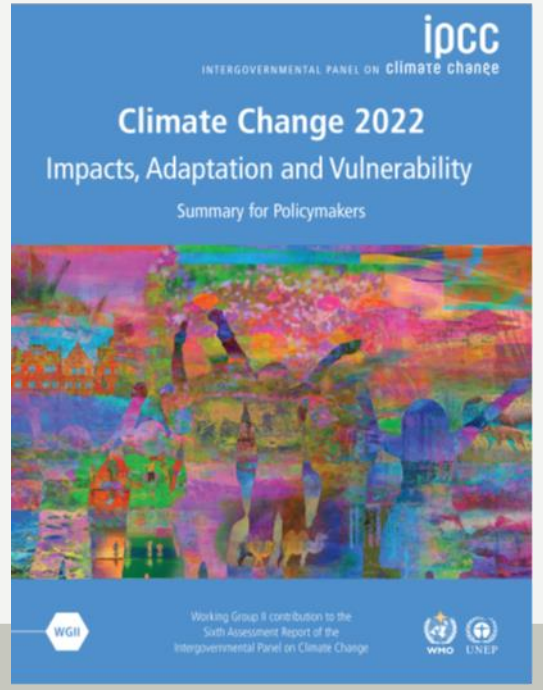

Slide courtesy Ralf Ludwig, LMU

Quelle: F. Löchner für LRZ

#	Site	Manufacturer	Computer	Country	Cores	Rmax [Pflops]	Power [MW]
1	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.07GHz, Mellanox EDR, NVIDIA GV100	USA	2,397,824	143.5	9.8
2	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA GV100	USA	1,572,480	94.6	7.4
3	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 260C 1.45GHz	China	10,649,600	93.0	15.4
4	National University of Defense Technology	NUDT	Tianhe-2A ANUDT TH-IVB-FEP, Xeon 12C 2.2GHz, Matrix-2000	China	4,981,760	61.4	18.5
5	Swiss National Supercomputing Centre (CSCS)	Cray	Piz Daint Cray XC50, Xeon E5 12C 2.6GHz, Aries, NVIDIA Tesla P100	Switzerland	387,872	21.23	2.38
6	Los Alamos NL / Sandia NL	Cray	Trinity Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries	USA	979,072	20.16	7.58
7	National Institute of Advanced Industrial Science and Technology	Fujitsu	AI Bridging Cloud Infrastructure (ABCI) PRIMERGY CX2550 M4, Xeon Gold 20C 2.4GHz, IB-EDR, NVIDIA V100	Japan	391,680	19.9	1.65
8	Leibniz Rechenzentrum	Lenovo	SuperMUC-NG ThinkSystem SD530, Xeon Platinum 8174 24C 3.1GHz, Intel Omni-Path	Germany	305,856	19.5	
9	Oak Ridge National Laboratory	Cray	Titan Cray XK7, [illegible]	USA	560,640	17.6	[illegible]
10	Lawrence Livermore National Laboratory	IBM	[illegible]	[illegible]	1,572,864	17.2	7.89

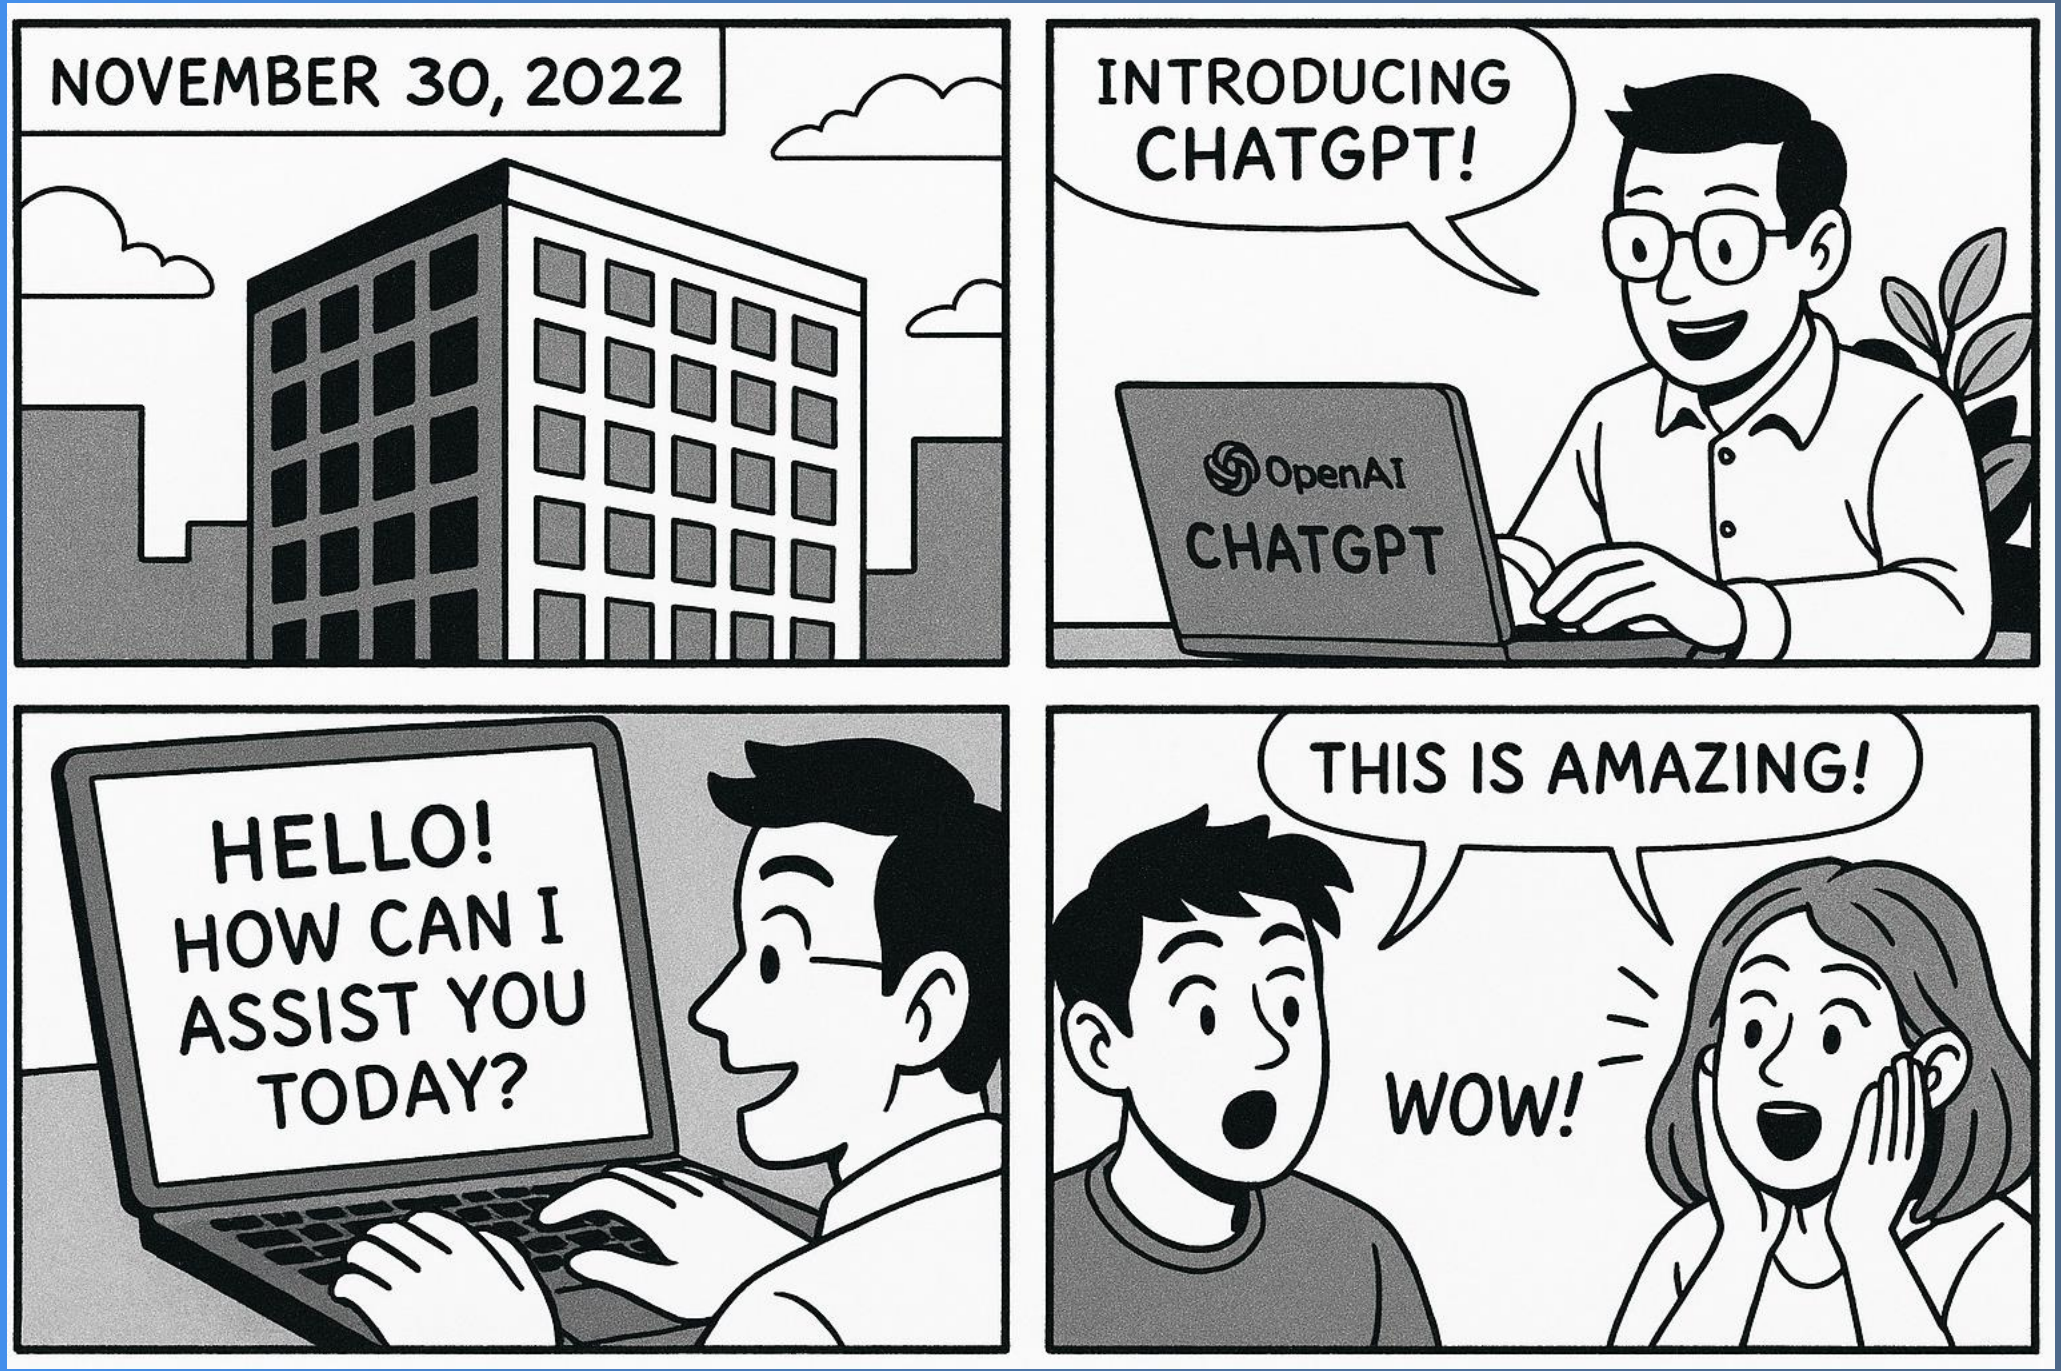
NOVEMBER 30, 2022
INTRODUCING CHATGPT!
OpenAI
CHATGPT
HELLO! HOW CAN I ASSIST YOU TODAY?
THIS IS AMAZING!
WOW!

Which of the following statements is true?

- Instead of HPC systems, we should only buy AI systems. They can do everything.
- Exascale is what delivers > 1 Exaflop/s in 64 Bit Linpack
- AI is today, where HPC was 30 years ago

Fotos:
https://www.stern.de/politik/ausland/mit-donald-trump-oder-ohne--wie-republikaner-2022-wieder-an-die-macht-wollen-31455198.html
https://www.spiegel.de/ausland/fox-news-soll-donald-trump-vorab-fragen-eines-tv-duells-zugespielt-haben-a-a2e43dc6-c5fe-4305-9ef1-bba97a8d976f

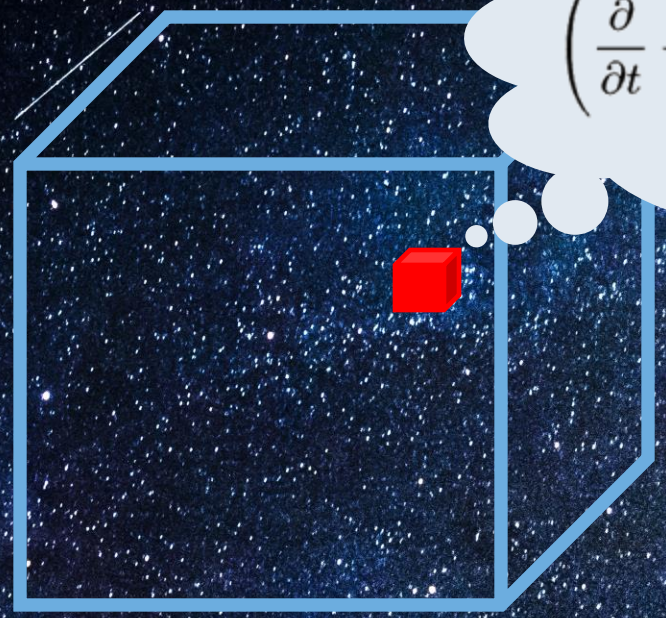

$$\left(\frac{\partial}{\partial t} + \vec{v} \cdot \nabla_{\vec{x}} + \frac{\vec{F}}{m} \cdot \nabla_{\vec{v}}\right) f(\vec{x}, \vec{v}, t) = \left.\frac{\partial f}{\partial t}\right|_{\text{Stoss}}$$

https://de.wikipedia.org/wiki/Boltzmann-Gleichung

Photo by Federico Beccari on Unsplash

Simple Task: Train a neural network for sine of x

- Given.: y = sin(x)
- Find: Neural network for sin(x)

- Let us ask ChatGPT:
- „Approximate the sine function using a neural network with pytorch using tanh as activation function“

- ==> Jupyter Notebook

```python
import torch
import torch.nn as nn
import torch.optim as optim
import numpy as np
import matplotlib.pyplot as plt

# Step 1: Create the neural network model
class SineNN(nn.Module):
    def __init__(self):
        super(SineNN, self).__init__()
        # A simple neural network with two hidden layers
        self.hidden1 = nn.Linear(1, 64)  # First hidden layer with 64 units
        self.hidden2 = nn.Linear(64, 64) # Second hidden layer with 64 units
        self.output = nn.Linear(64, 1)   # Output layer with 1 unit (for sine value)

    def forward(self, x):
        # Apply tanh activation in the hidden layers
        x = torch.tanh(self.hidden1(x))  # First hidden layer
        x = torch.tanh(self.hidden2(x))  # Second hidden layer
        x = self.output(x)  # Output layer
        return x

# Step 2: Prepare the dataset
# Generate input data (x values)
x_values = np.linspace(-2 * np.pi, 2 * np.pi, 100)
y_values = np.sin(x_values)

# Convert to PyTorch tensors
x_train = torch.tensor(x_values, dtype=torch.float32).view(-1, 1)  # Shape (1000, 1)
y_train = torch.tensor(y_values, dtype=torch.float32).view(-1, 1)  # Shape (1000, 1)

# Step 3: Initialize the neural network, loss function, and optimizer
model = SineNN()
criterion = nn.MSELoss()        # Mean Squared Error loss
optimizer = optim.Adam(model.parameters(), lr=0.001)  # Adam optimizer
```

Result: neural network for sine of x

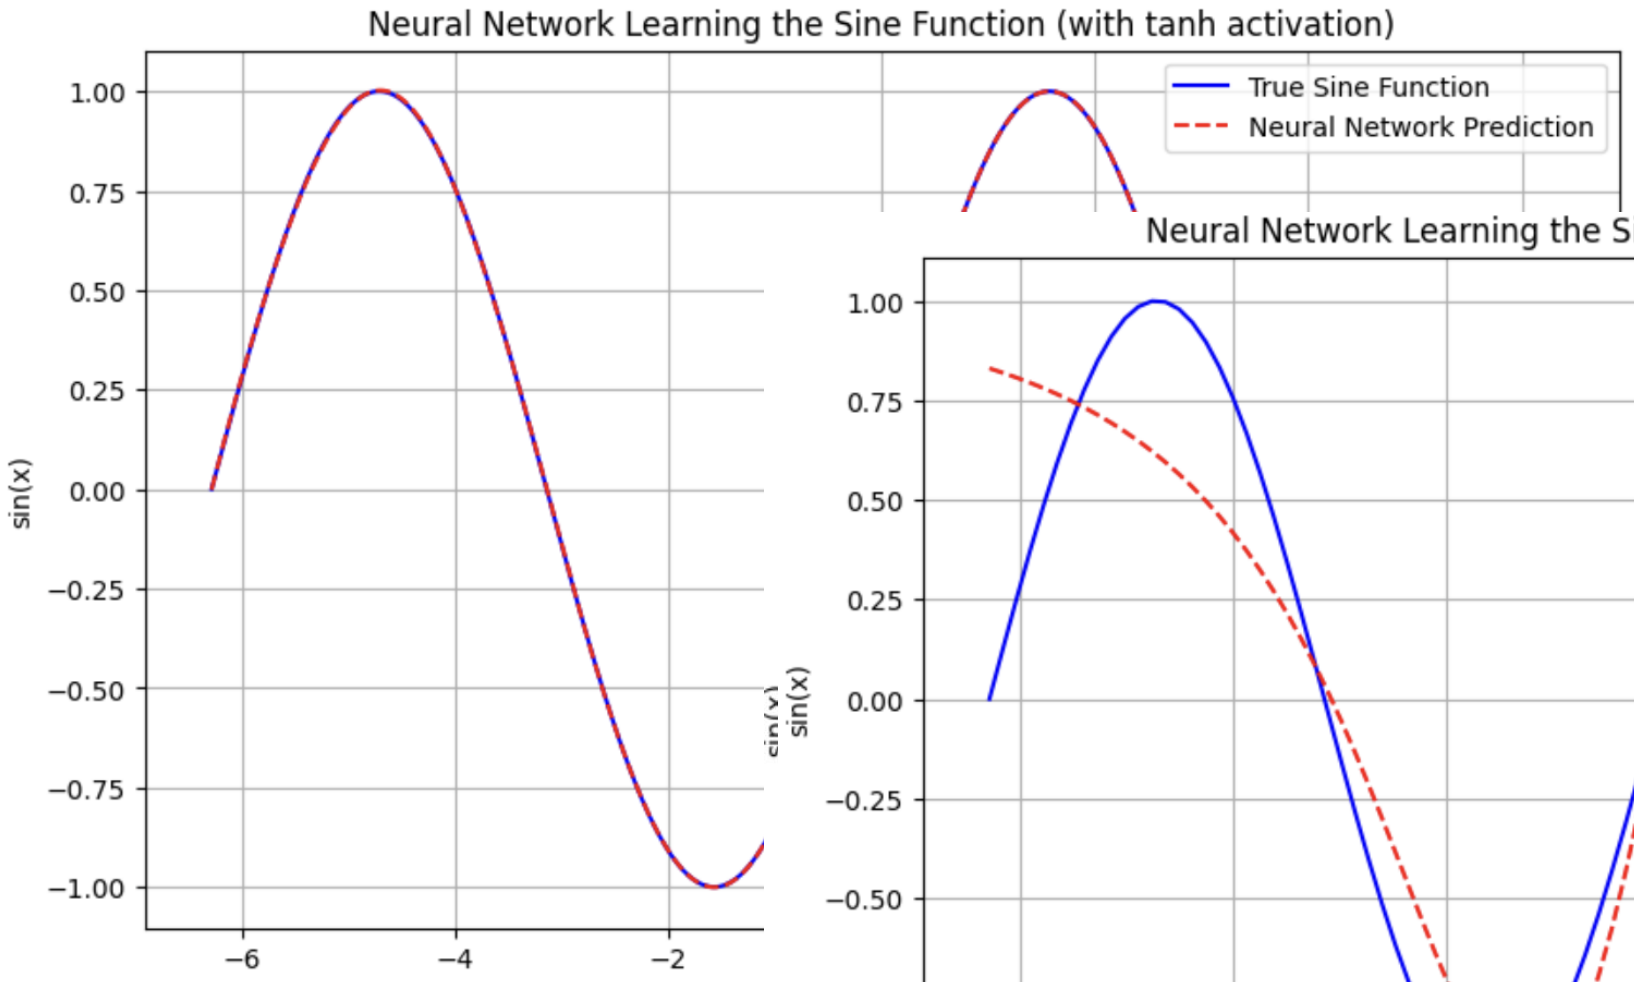

100 epochs
with 100 data points

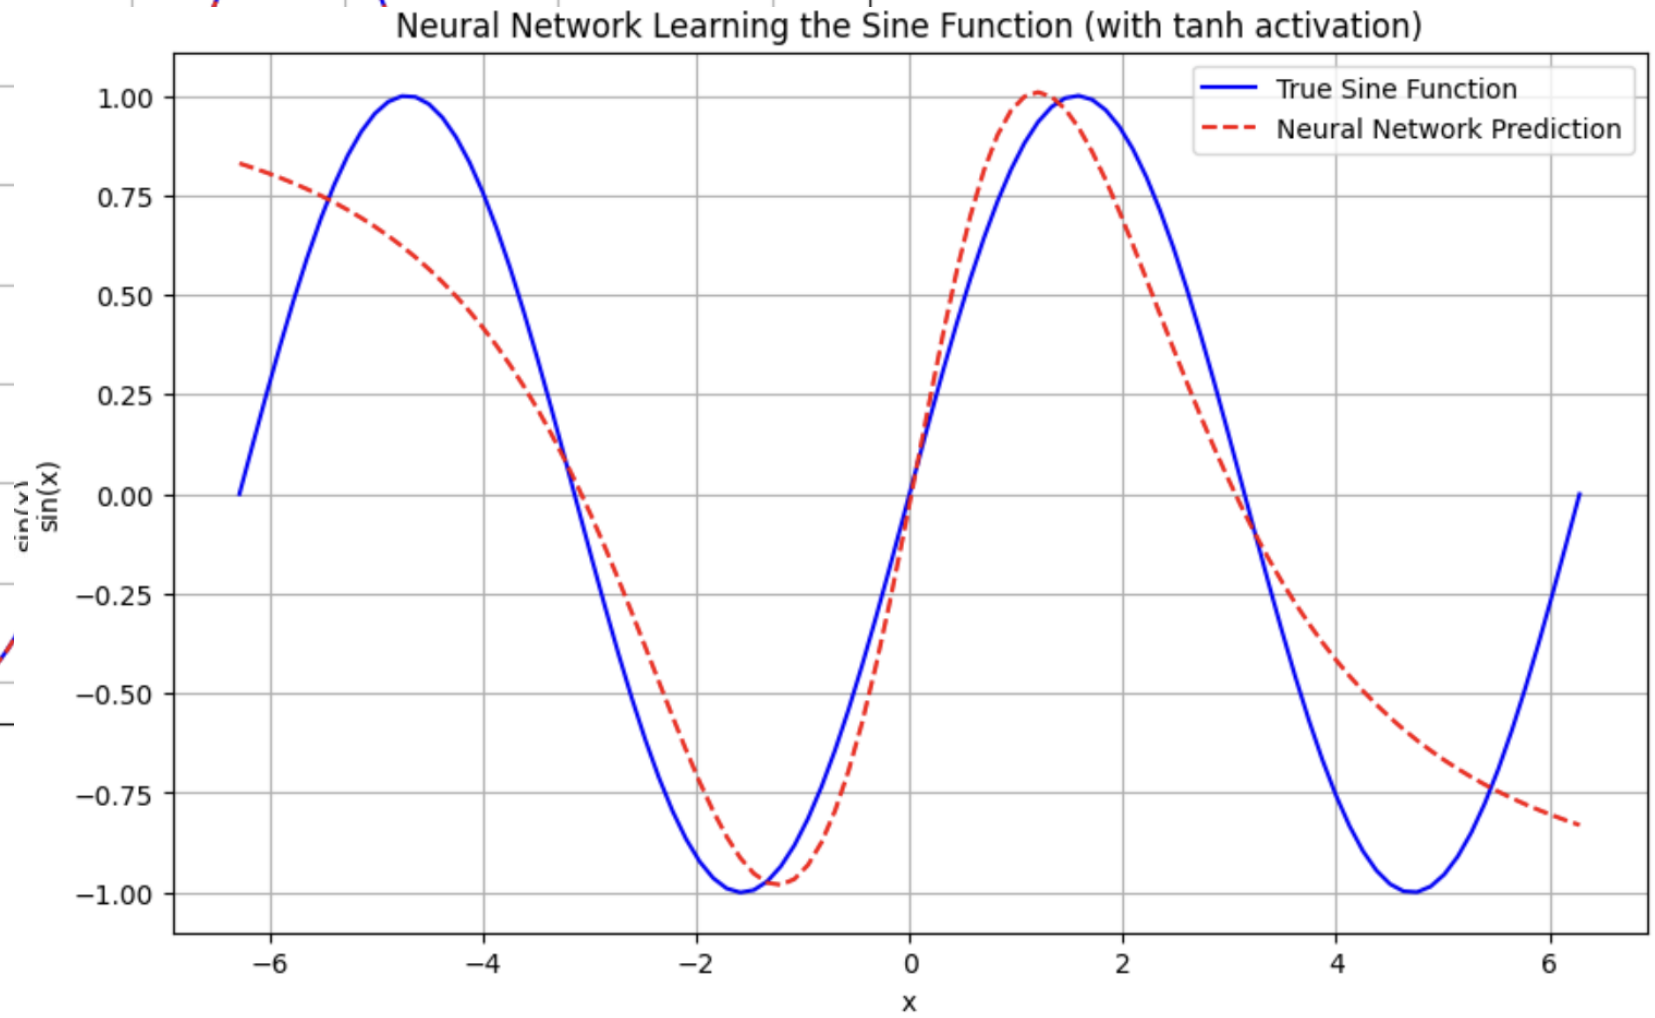

5000 epochs
with 1000 data points

SuperMUC-NG Phase 2

Lenovo Intel

240
Direct hot water cooled compute nodes (CPU:Sapphire Rapids + GPU:Ponte Vecchio)

SD650-I v3
Lenovo platform

1 PB
DAOS storage system mit Intel Xeon Scalable CPUs and 3rd Gen Intel® Optane™ persistent memory

ELECTRICITY CONSUMPTION

How much electrical power do we need for the digital transformation?

- **Smartphone**: iPhone 16 (Standard)
  - Battery Capacity 3.561 mAh
  - Per charge (with 0,30 € per kWh) → approx. 0,005 € per day
  - With one charge per day → approx. 1,950 € per year

Created by Thays Malcher from Noun Project

- Personal computer:
  - Current Intel/AMD multicore processor with dedicated GPU
  - Per day (with 4 hours running) → 200 kWh per year
    → approx. 60 € per year
- Electric vehicle :
  - BMW i3, 12.9 kWh/100 km → 3,87 € per 100 km
  - Annual mileage approx. 15.000 km → 580,50 € per year

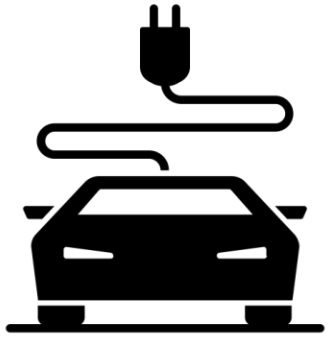

Created by Symbolon from Noun Project

Leibniz Supercomputing Centre (LRZ)
of the Bavarian Academy of Sciences and Humanities

Partner for Digital Transformation of Science & Research

IT-Services for Science since 1962

Munich Universities, Bavaria, Germany, Europe, …

Average Power Consumption Top 10/50/500 Systems

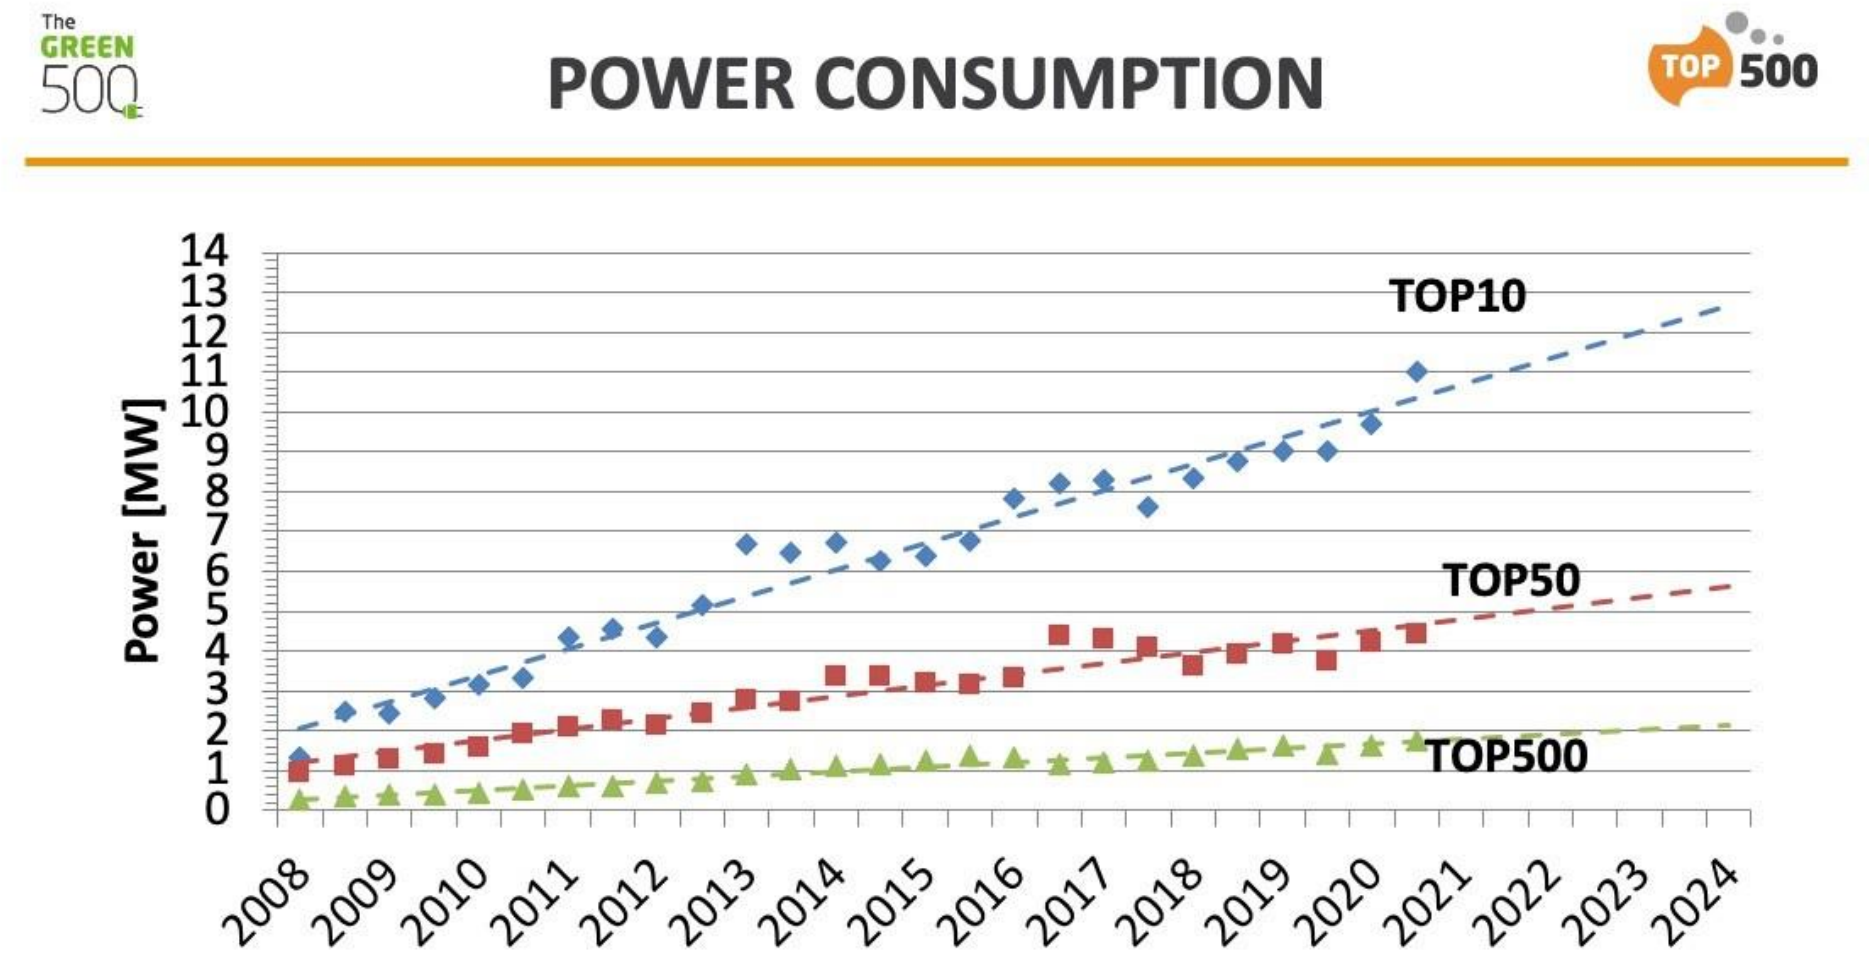

Source: https://twitter.com/top500supercomp/status/1329583763844599808

LRZ HPC/AI Systems and Energy Efficiency

SuperMUC

CoolMUC-3

SuperMUC-NG

2012

2017

2018

2025

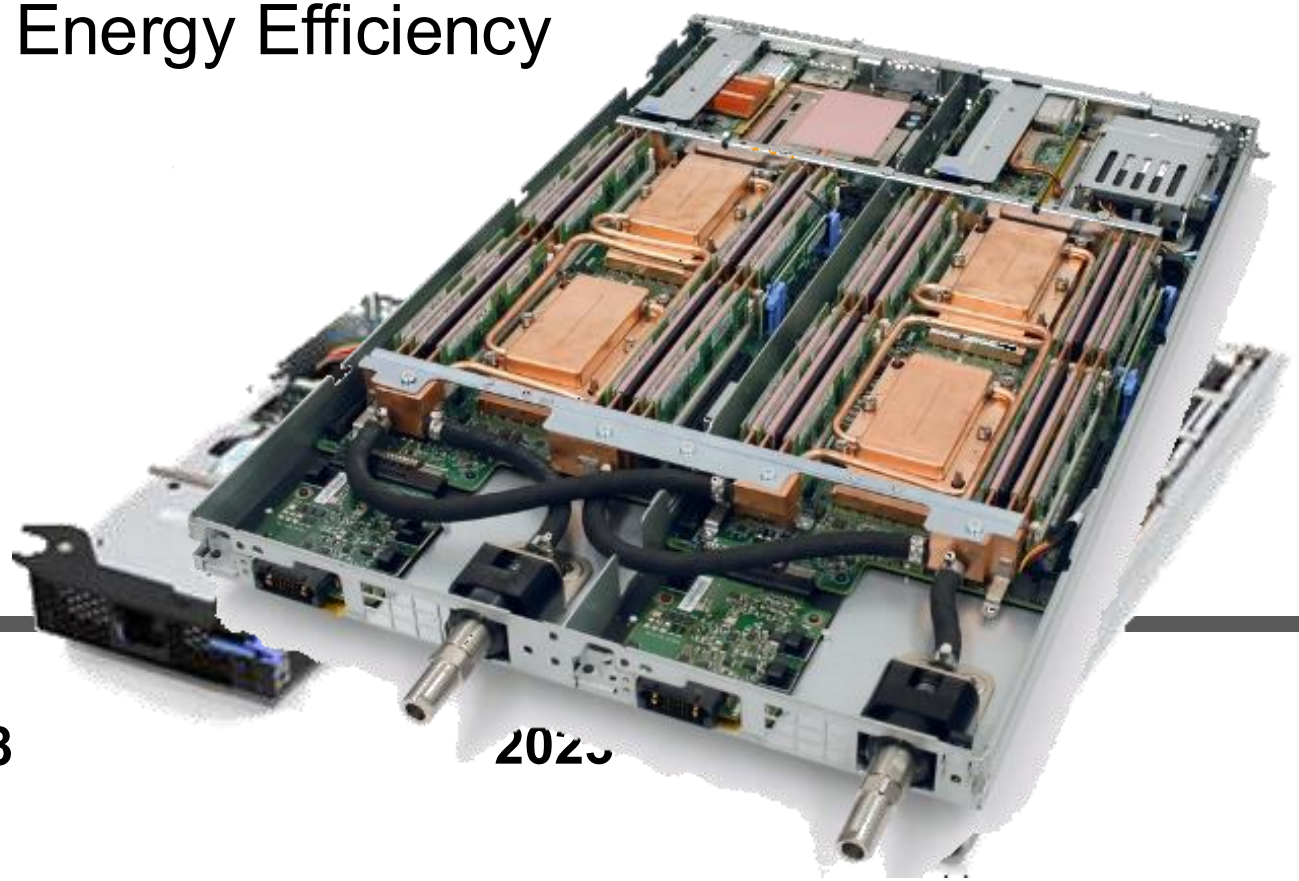

2024: SuperMUC-NG Phase 2

- **Direct hot-water-cooled compute nodes**
- **Intel CPUs & Intel PVC GPUs**
- **PUE = 1,03**

LRZ HPC/AI Systems and Energy Efficiency

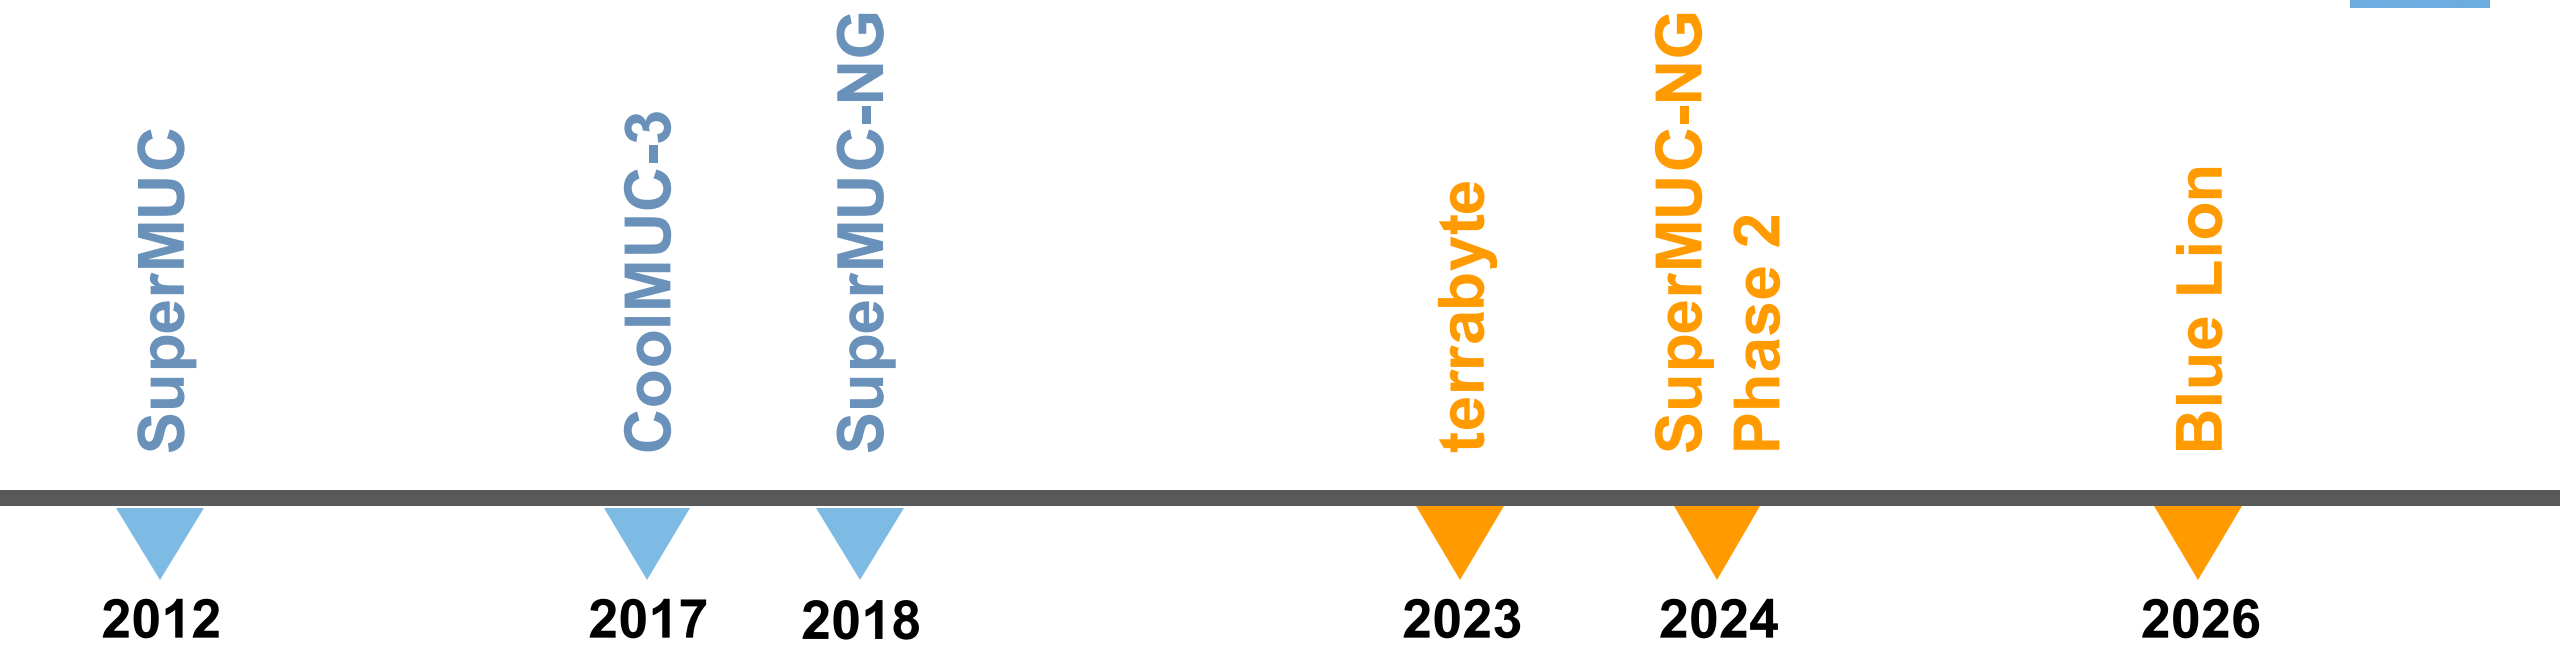

2026: Blue Lion

- **Upcoming Supercomputing System**

The Four Pillar Model

Neighboring Buildings

Data Center

Improve PUE (Power Usage Effectiveness)	Reduce Hardware Power Consumption	Optimize Resource Usage, Tune System	Optimize Application Performance

Holistic Optimization Strategy

Building Infrastructures	**HPC System Hardware**	**HPC System Software**	**HPC Applications**

Torsten Wilde, Axel Auweter, Hayk Shoukourian
„The 4 Pillar Framework for energy efficient HPC data centers“,
Comput Sci Res Dev (2014) 29: 241. https://doi.org/10.1007/s00450-013-0244-6

Lessons learned towards Sustainability

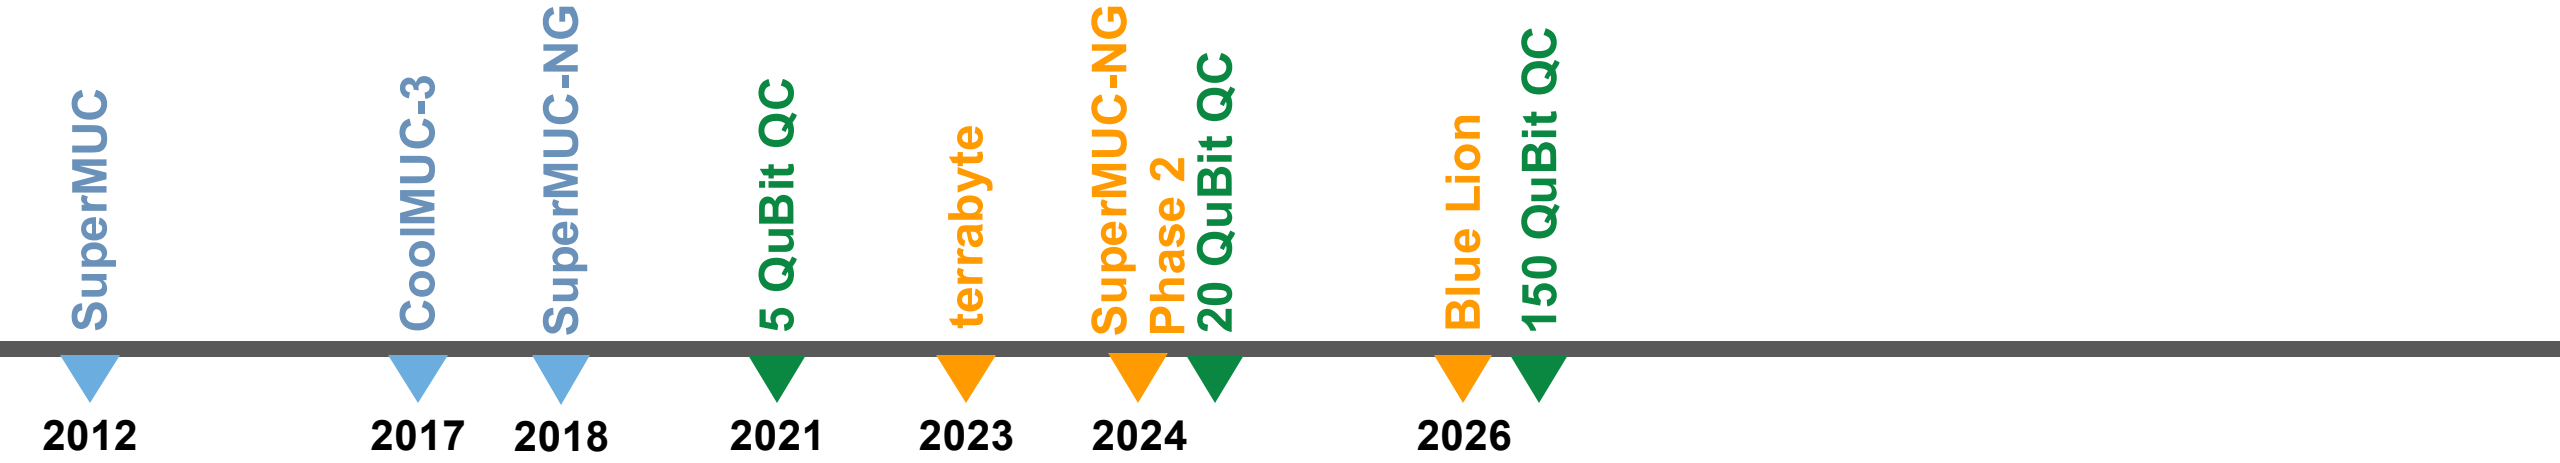

- DVFS - Dynamic Voltage Frequency Scaling (on average: 68% of nominal frequency)
- Control of power capping / sleep modes
- 100% renewable energies by contract
- Cooling infrastructure optimizations - Glycol free cooling
- Waste heat usage
- Adsorption cooling machines
- Include job-power usage in job protocols
- Future computing technologies – Accelerators, Quantum, …

What about ChatGPT: https://chat.openai.com/

Question: **Is ChatGPT energy efficient?**

Answer:

- *ChatGPT's energy efficiency depends on factors like **hardware**, **model size**, and **deployment**.*
- *Large models can be **energy-intensive**.*
- *Researchers are working on techniques to improve efficiency.*
- *The actual efficiency varies based on **implementation** and **infrastructure**.*
- *Improvements are expected over time.*

Question: **What is the energy consumption for training ChatGPT?**

Answer:

- *To provide a rough estimate, training GPT-3, which has 175 billion parameters, could consume hundreds of megawatt-hours (MWh) of electricity during the training process.*

NVIDIA Corporation HQ, Santa Clara, California

www.top500.org

Top 500 Supercomputers

- List: June 2024
- Vendor: NVIDIA
- Top 5 out of 22 entries

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
10	**Eos NVIDIA DGX SuperPOD** - NVIDIA DGX H100, Xeon Platinum 8480C 56C 3.8GHz, NVIDIA H100, Infiniband NDR400, Nvidia NVIDIA Corporation United States	485,888	121.40	188.65	
15	**Selene** - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	555,520	63.46	79.22	2,646
25	**Sejong** - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Infiniband HDR, Nvidia NAVER Corp South Korea	277,760	32.97	40.77	
34	**Yep1** - NVIDIA DGX H100, Xeon Platinum 8480C 56C 3.8GHz, NVIDIA H100 SXM5 80 GB, Infiniband NDR400, Nvidia Ahrefs US DC United States	73,584	23.32	34.58	
35	**BioHive-2** - NVIDIA DGX H100, Xeon Platinum 8480C 56C 2GHz, NVIDIA H100, Infiniband NDR400, Nvidia Recursion United States	40,320	23.32	34.22	

EOS NVIDIA DGX SUPERPOD

https://blogs.nvidia.com/blog/eos/

4608 H100 GPUs

Top 500 Supercomputers for HPC and Artificial Intelligence

NVIDIA Technology:

- DGX A100
  https://www.nvidia.com/en-us/data-center/dgx-a100/
- NVIDIA DGX H100
  https://www.nvidia.com/en-us/data-center/dgx-h100/
- NVIDIA DGX SuperPod
  https://www.nvidia.com/en-us/data-center/dgx-superpod/

- All current systems are air cooled
- Water cooling is only a supplement
- Overhead for air cooling:
  - +30% Fans
  - +30% Cooling Infrastructure

NVIDIA H100 Tensor-Core-GPU

- Max. energy consumption (without cooling):
  700 W per year: **3741 kWh**
  (at 61% utilization)

- Compare: average household in Germany per year (2020):
  approx. **3200 kWh**

- NVIDIA estimates sales for 2024 of approx. 2 Mio. H100
- Total power: 1.4 GW (without cooling)

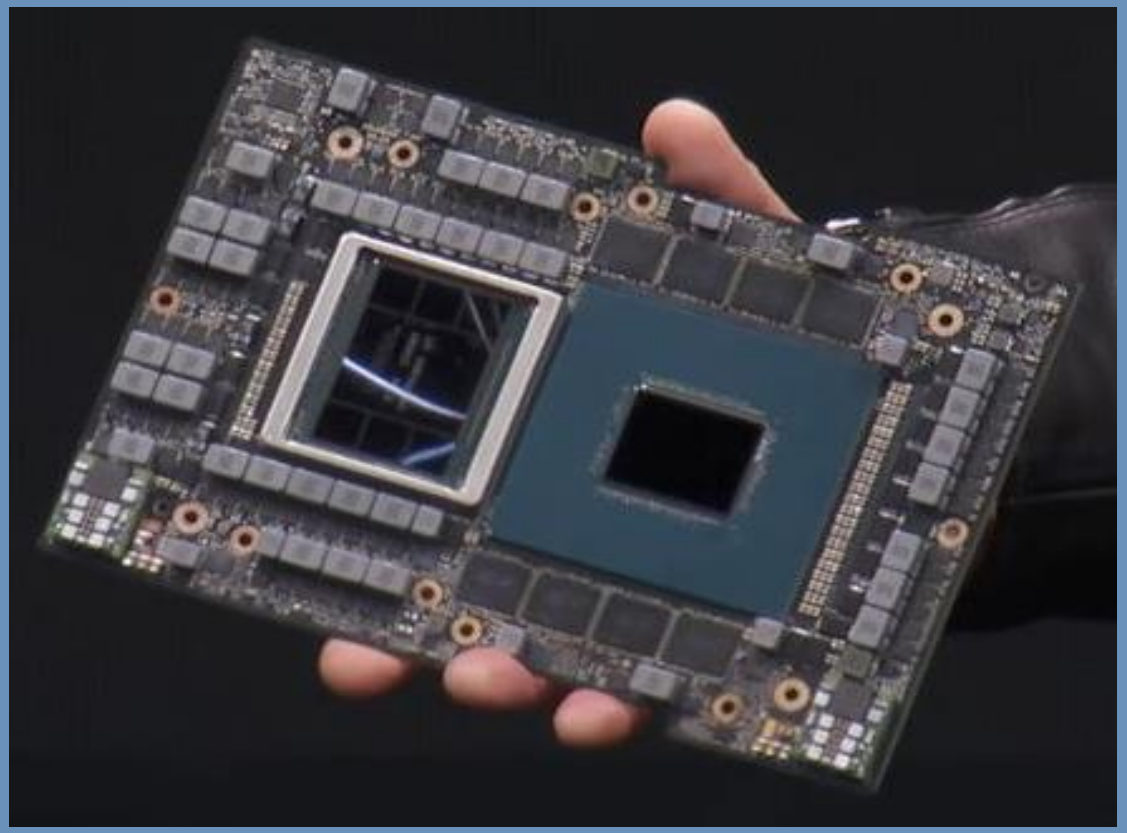

NVIDIA Grace Hopper Chip
https://www.anandtech.com/show/18877/nvidia-grace-hopper-has-entered-full-production-announcing-dgx-gh200-ai-supercomputer

Example AI Infrastructure

The Munich Center for Machine Learning (MCML)

Left:

- 2-3 NVIDIA DGX A100
- =16-24 GPUs per Rack
- Air-cooled
  (in a water-cooled rack)
- **PUE**: 1,65-1,80

Power Usage Effectiveness:

the **total amount of power**
entering a data center
divided by
the power used to run
the **IT equipment**

Example:

PUE 1,65 = 1 für IT
+ 65% für
Infrastructure

Air vs. Water

Name	Unit	Air	Water	Factor
Thermal Conductivity	W/(mK)	0.026	0.598	23x
Heat Capacity	J/(gK)	1.006	4.185	4x
Volumetric Heat Capacity	kJ/(m^3K)	1.196	4178	3493x
Thermal Inertia	J/($m^2Ks^{1/2}$)	5.563	1581	284x

Hot water cooling in the compute room

SuperMUC-NG Node – with hot water cooling

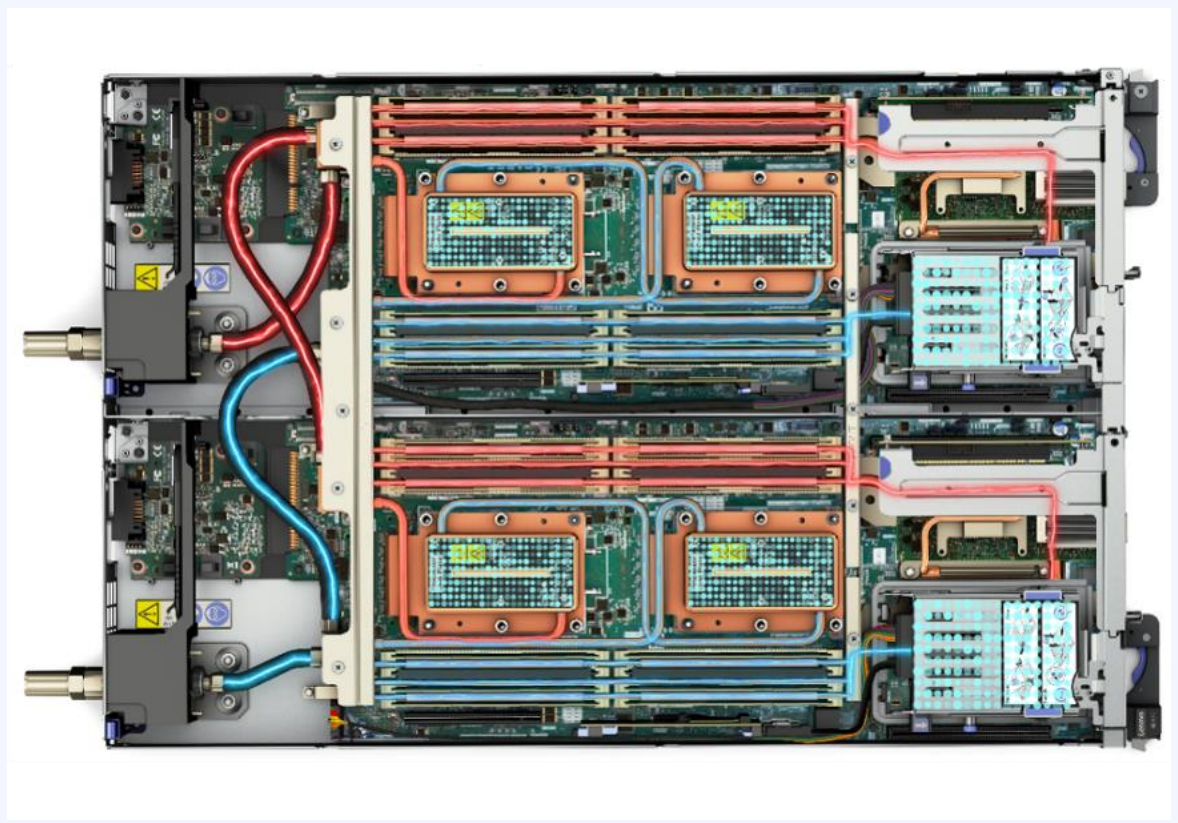

Source: Photos Lenovo – Produkt: https://www.lenovo.com/de/de/data-center/servers/high-density/ThinkSystem-SD650/p/77XX7DSSD65

SuperMUC-NG Cool Manager

3.000 l/h pro Rack

Reduced server power consumption

- Lower processor power consumption (~ 5%)
- No fan per node (~ 4%)

Reduce cooling power consumption

- At 40°C free cooling all year long (> 15%)

Energy Aware Scheduling

- Only CPU bound jobs get max frequency (~ 5%)

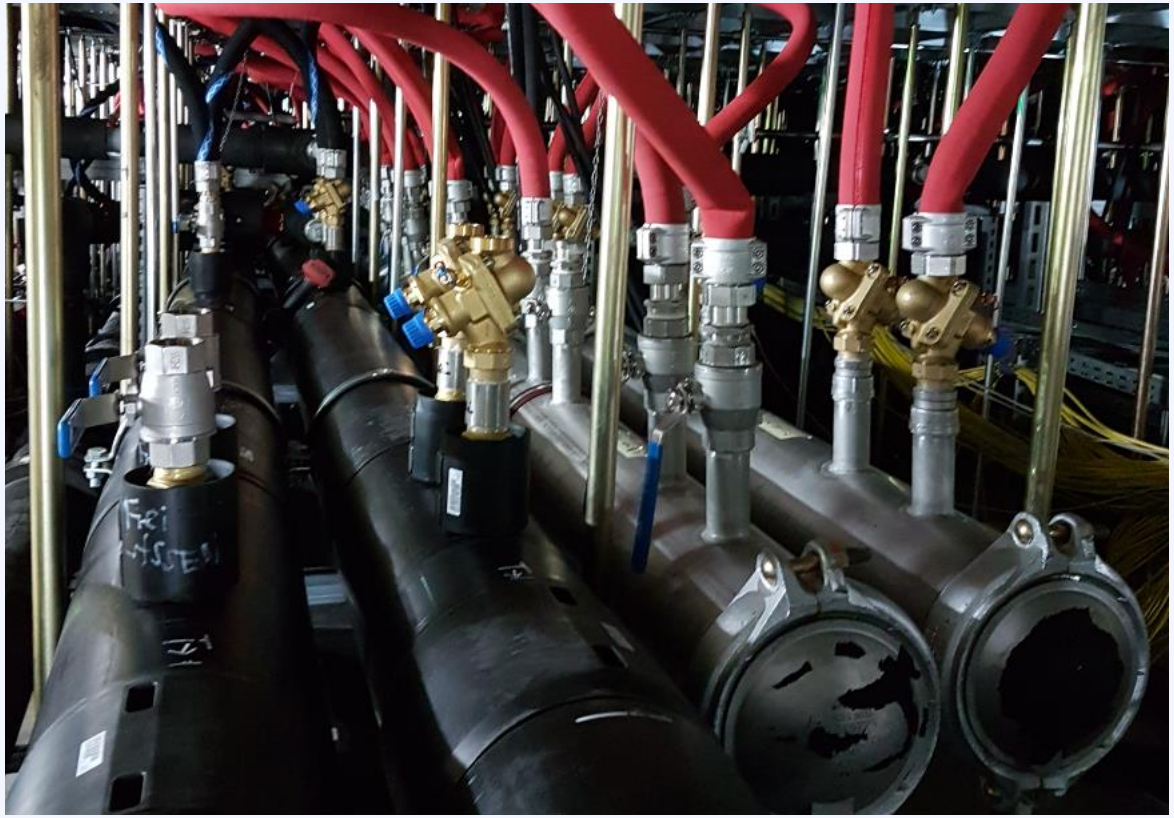

Example AI Infrastructure

The Munich Center for Machine Learning (MCML)

Left:

- 2-3 NVIDIA DGX A100
- =16-24 GPUs per Rack
- Air-cooled
  (in a water-cooled rack)
- PUE: 1,65-1,80

Right:

- Up to 144 GPUs per Rack
- Lenovo HGX-based nodes
  (4x A100)
- Direct hot-water cooled
  (allows free cooling)
- PUE: 1,03-1,05

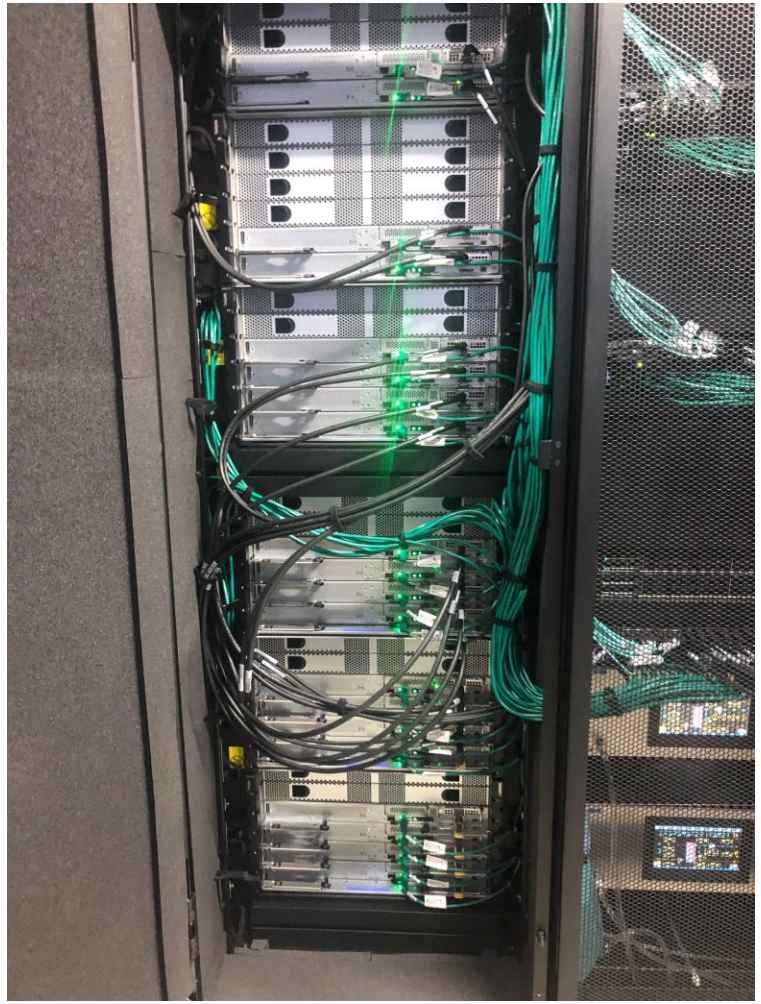

HELMUT SCHMIDT UNIVERSITÄT
Universität der Bundeswehr Hamburg
hpc.bw
HSUper
576 MEGWARE Compute Nodes
5 MEGWARE GPU Nodes
Direct-to-chip liquid cooling
1162 Intel® Xeon Scalable 8360Y
41,832 CPU cores
149 TB DDR4 main memory
3.21 PFlop/s RPEAK (CPU-only)
NVIDIA® HDR InfiniBand Network
1 PB net BeeGFS Storage System
1 PB net Ceph Storage System
MEGWARE
MEGWARE
MEGWARE

Savings potentials through direct water cooling

	Infrastructure	IT
Investment	Fewer chillers requires	IT slightly more expensive (depending on quantity)
Operations costs	Power consumption for refrigeration is significantly reduced	Power consumption of water-cooled IT is lower (fewer fans)

Foto: Connie Zhou/Google/Zuma https://www.nature.com/articles/d41586-018-06610-y

Example

Google‘s Water Demand in Luxembourg

- New Google data center in Luxembourg - €1.2 billion
- 12% of Luxembourg‘s electricity demand
- Cooling water demand: 10 million l/day
- Google refuses to provide information on the grounds that it "*could reveal details of the DC and the agreement with the Luxembourg government to competitors*"

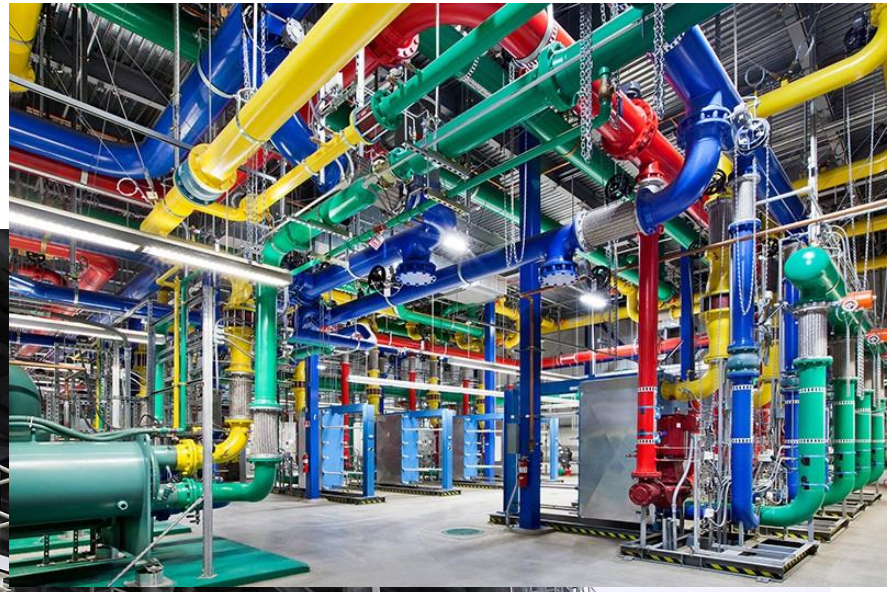

https://orf.at/stories/3184229/

Foto: Reuters/Yves Herman

Nature News Feature - https://www.nature.com/articles/d41586-018-06610-y

How to stop data centres from gobbling up the world's electricity

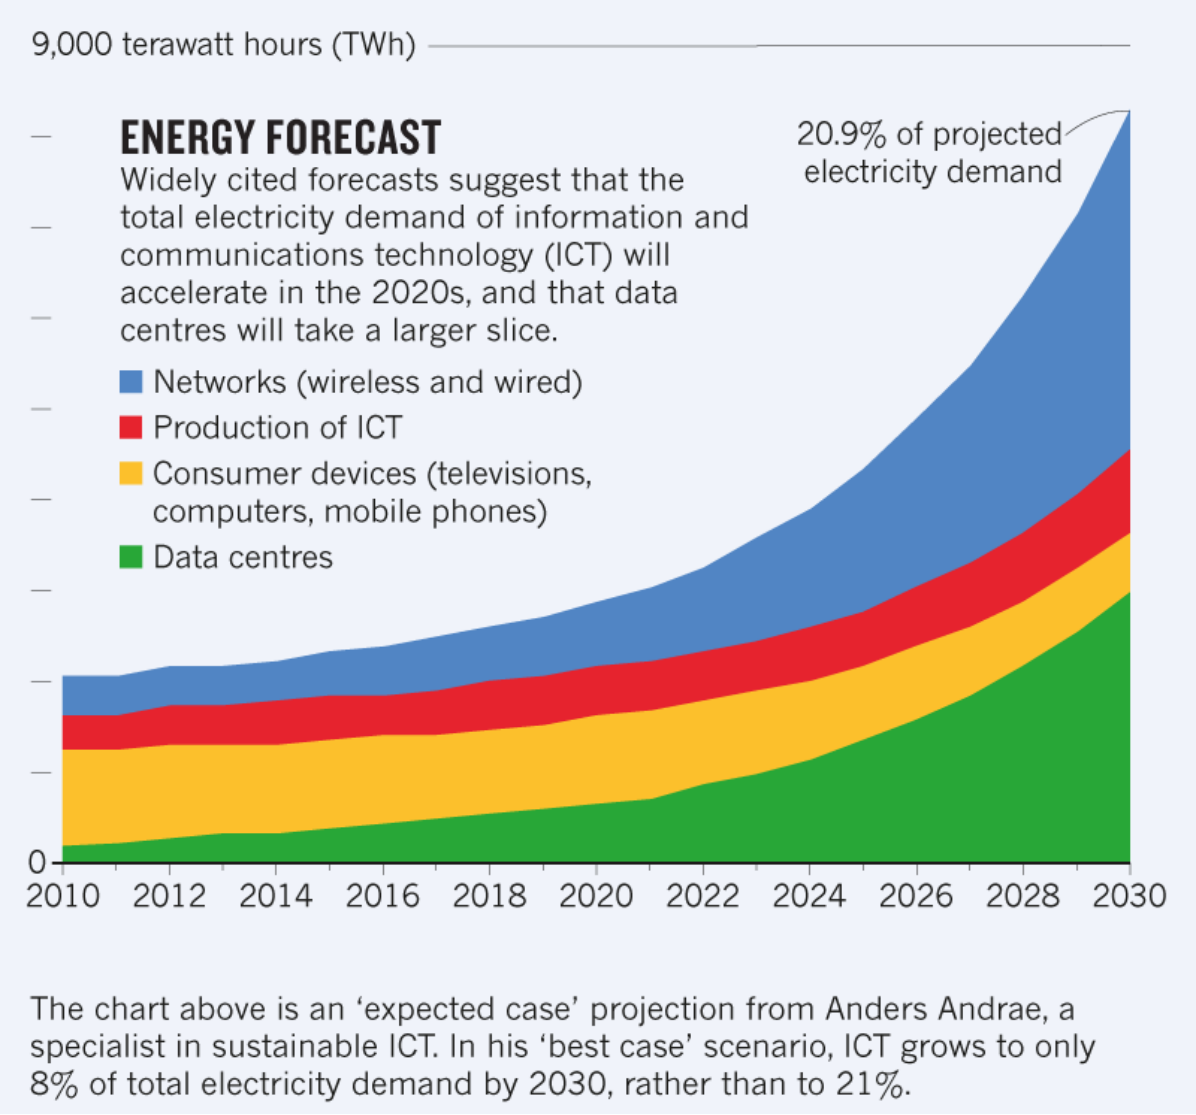

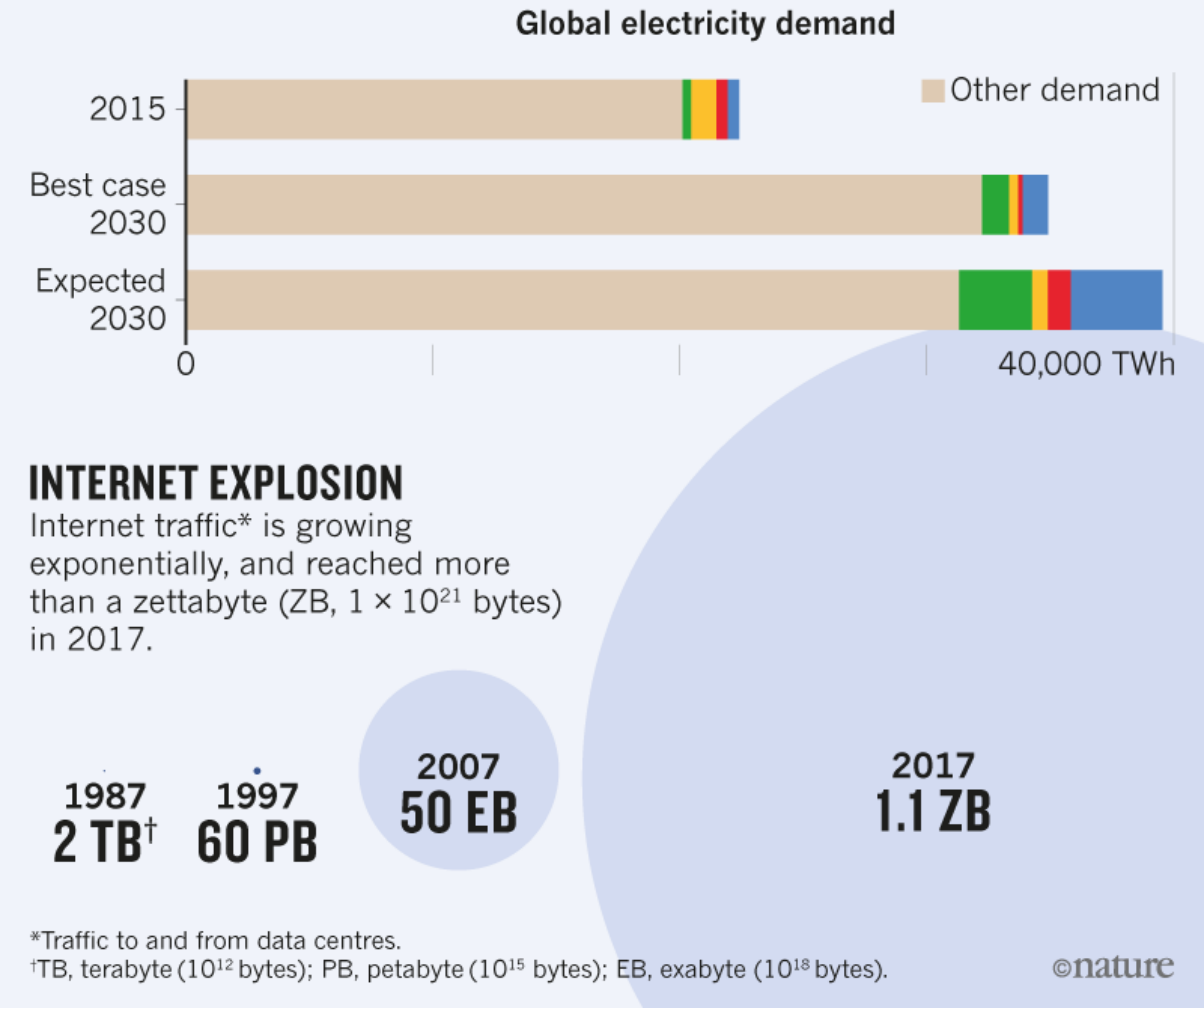

https://www.sueddeutsche.de/wirtschaft/microsoft-stargate-ki-openai-1.6509900

Project „Stargate“

- 5 Steps Project to build AI Data Centers
- 2026 First Data Center Microsoft jointly with OpenAI
- In Mt. Pleasant, Wisconsin
- Costs: 10 Billion USD

- 2030 Completion of Supercomputer „Stargate“
- Total: 115 Billion USD
- Electricity: 5 GWatt

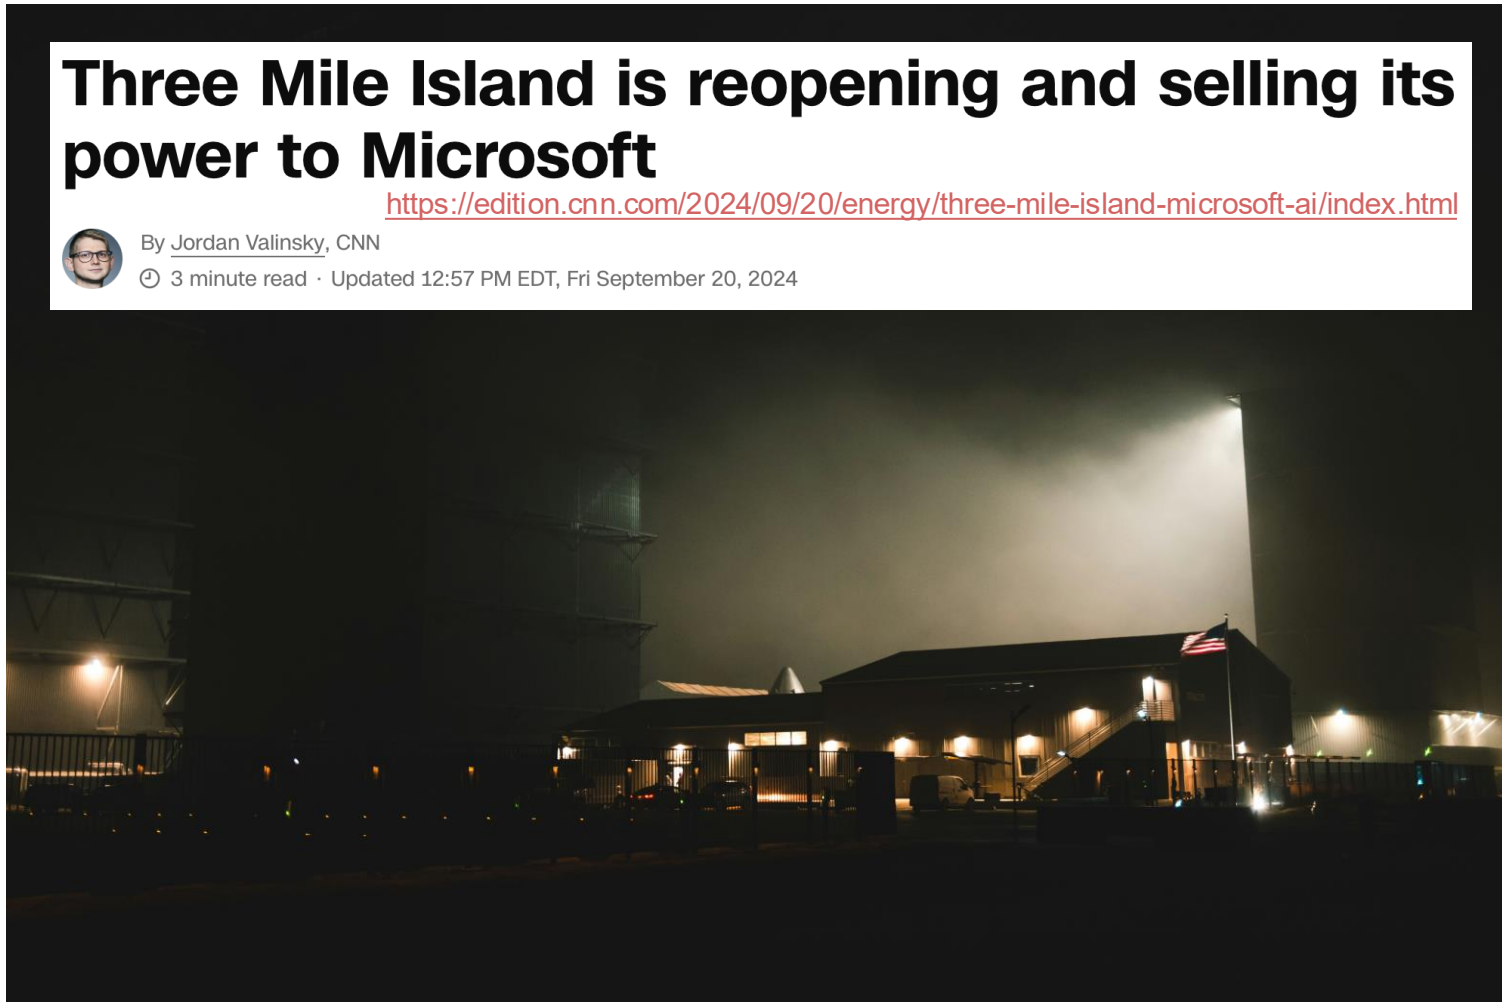

Foto von Jeswin Thomas auf Unsplash

https://www.hugendubel.de/de/buch_kartoniert/kate_crawford-atlas_of_ai-41511599-produkt-details.html

How sustainable is AI? Recommendation: Kate Crawford, „*Atlas of AI*“

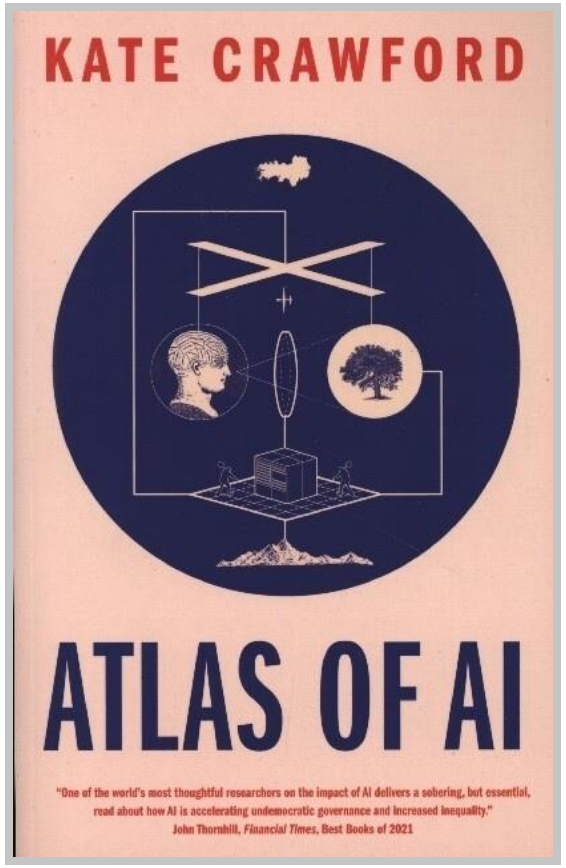

- In addition to the electricity requirement, valuable raw materials are also needed:
  - Rare earths for the production of the chips
  - Water for cooling the data centers
  - …

Judgment Day for HPC: Will AI Take Control?

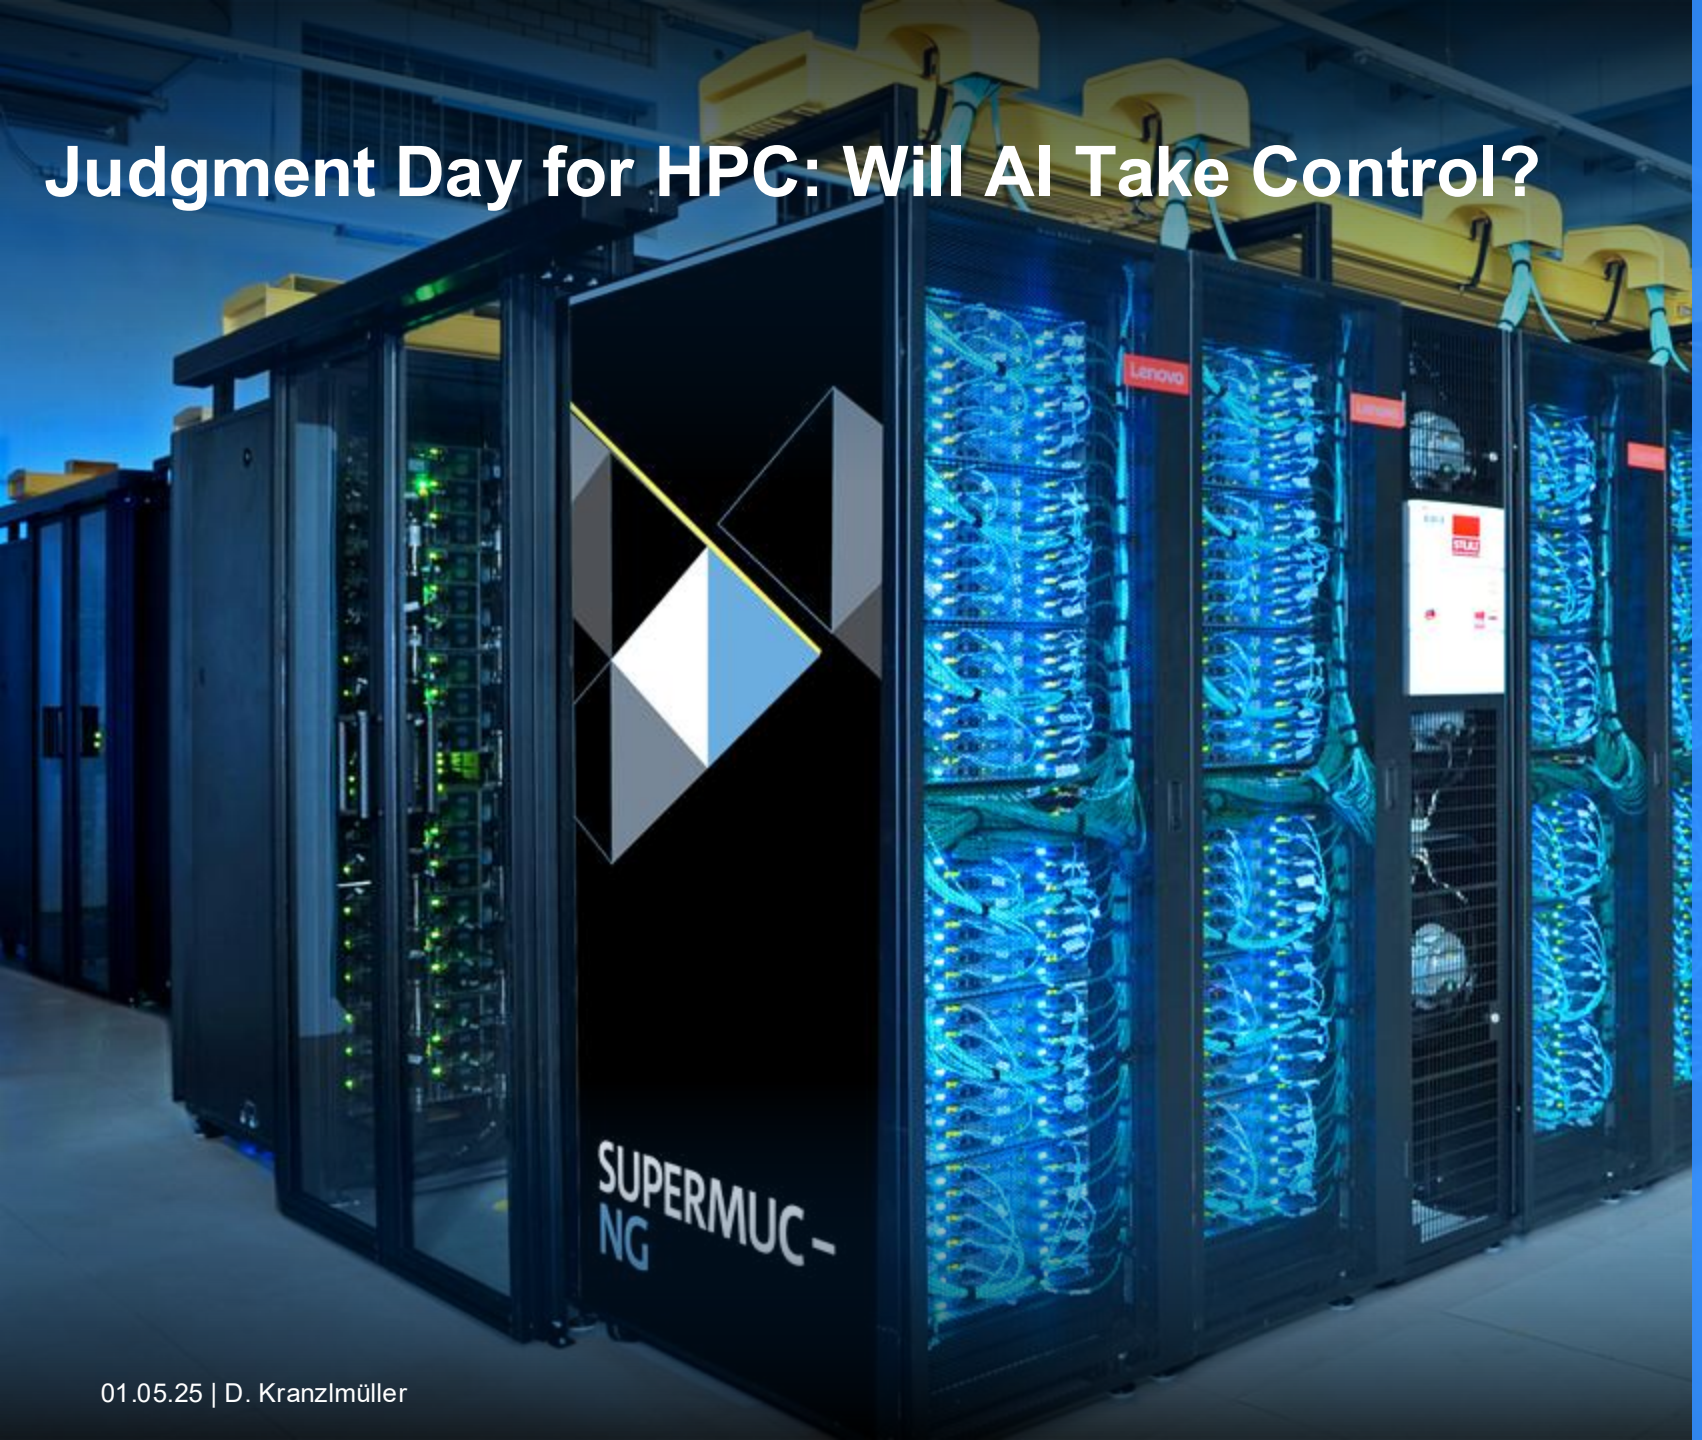

- **AI** has a role in HPC as a complementary feature, but not as a replacement
- We need an **integrated (heterogeneous) approach**
- AI can still learn a lot from traditional **HPC**

Contact: Prof. Dr. Dieter Kranzlmüller
kranzlmueller@lrz.de | @Kranzlmueller

https://ec.europa.eu/commission/presscorner/detail/en/ip_25_1013

Commission sets course for Europe's AI leadership

AI Continent Action Plan:

1. Building a large-scale AI data and computing infrastructure
   - Network of AI Factories
   - AI Gigafactories (with approximately 100,000 state-of-the-art AI ch
2. Increasing access to large and high-quality data
3. Developing algorithms and fostering AI adoption in strategic EU se
4. Strengthening AI skills and talents
5. Regulatory simplification

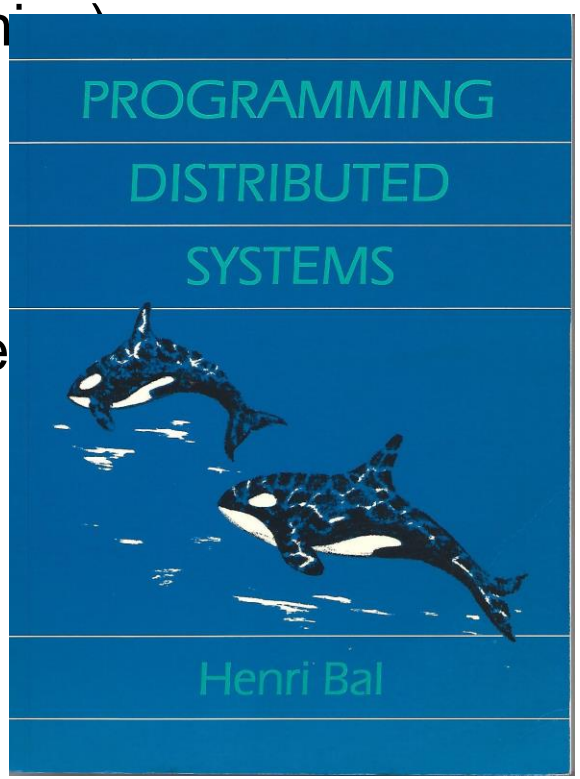

Published 1 January 1990

Ask ChatGPT

What did Henri Bal say?

Henri Bal, **a prominent figure in parallel and distributed computing**, has emphasized the **critical importance of adapting programming practices to leverage the full capabilities of modern computing systems**. In a 2015 interview, he highlighted the necessity for investment in research and education to effectively utilize the available computational power:

"*Investments in **research on programming** new massively-parallel machines are required to gain knowledge. Thus it must be examined how programs should be written for parallel computing methods and what extent of parallel calculations can be performed automatically." ("Soon we will use only one thousandth of available computer ...)*

This statement underscores **Bal's advocacy for advancing programming methodologies** to keep pace with the evolving landscape of high-performance computing.

From Large Language Models to **„Smart“ Language Models**

- *Sepp Hochreiter*, JKU Linz, founder of NXAI,
  „*Wants to sweep Open AI off the market* “
  https://industriemagazin.at/news/sepp-hochreiter-ich-moechte-open-ai-vom-markt-fegen/

- *Stephan Günnemann*, Director MDSI at TUM, founder of Pruna AI
  „*Smash your AI model*“
  https://www.pruna.ai/

- *Björn Ommer*, LMU and Stable Diffusion,
  „*Fighting the Industry‘s Megalomania* “
  https://www.sueddeutsche.de/wirtschaft/kuenstliche-intelligenz-ki-neurips-konferenz-bjoern-ommer-1.6318009

The next LRZ supercomputer

Blue Lion (formely known as ExaMUC)

The HPC Innovation/Integration Circle

HPC

Modeling & Simulation (M&S)
Natural World Hypothesis ▸
Equations ▸ Algorithms ▸
Computing ▸ Data ▸ Analysis

AI & Machine Learning

Data ▸ Algorithms
▸ Computing ▸ Pattern Recognition

Big Data

The HPC Innovation/Integration Circle

HPC

Modeling & Simulation (M&S)
Natural World ▸ Hypothesis ▸
Equations ▸ Algorithms ▸
Computing ▸ Data ▸ Analysis

AI & Machine Learning

Data ▸ Algorithms
▸ Computing ▸ Pattern Recognition

Quantum Computing

Big Data

LRZ Quantum Computer Q-Exa – Germanys 1. Quantum Demonstrator

HASTA LA VISTA, BABY