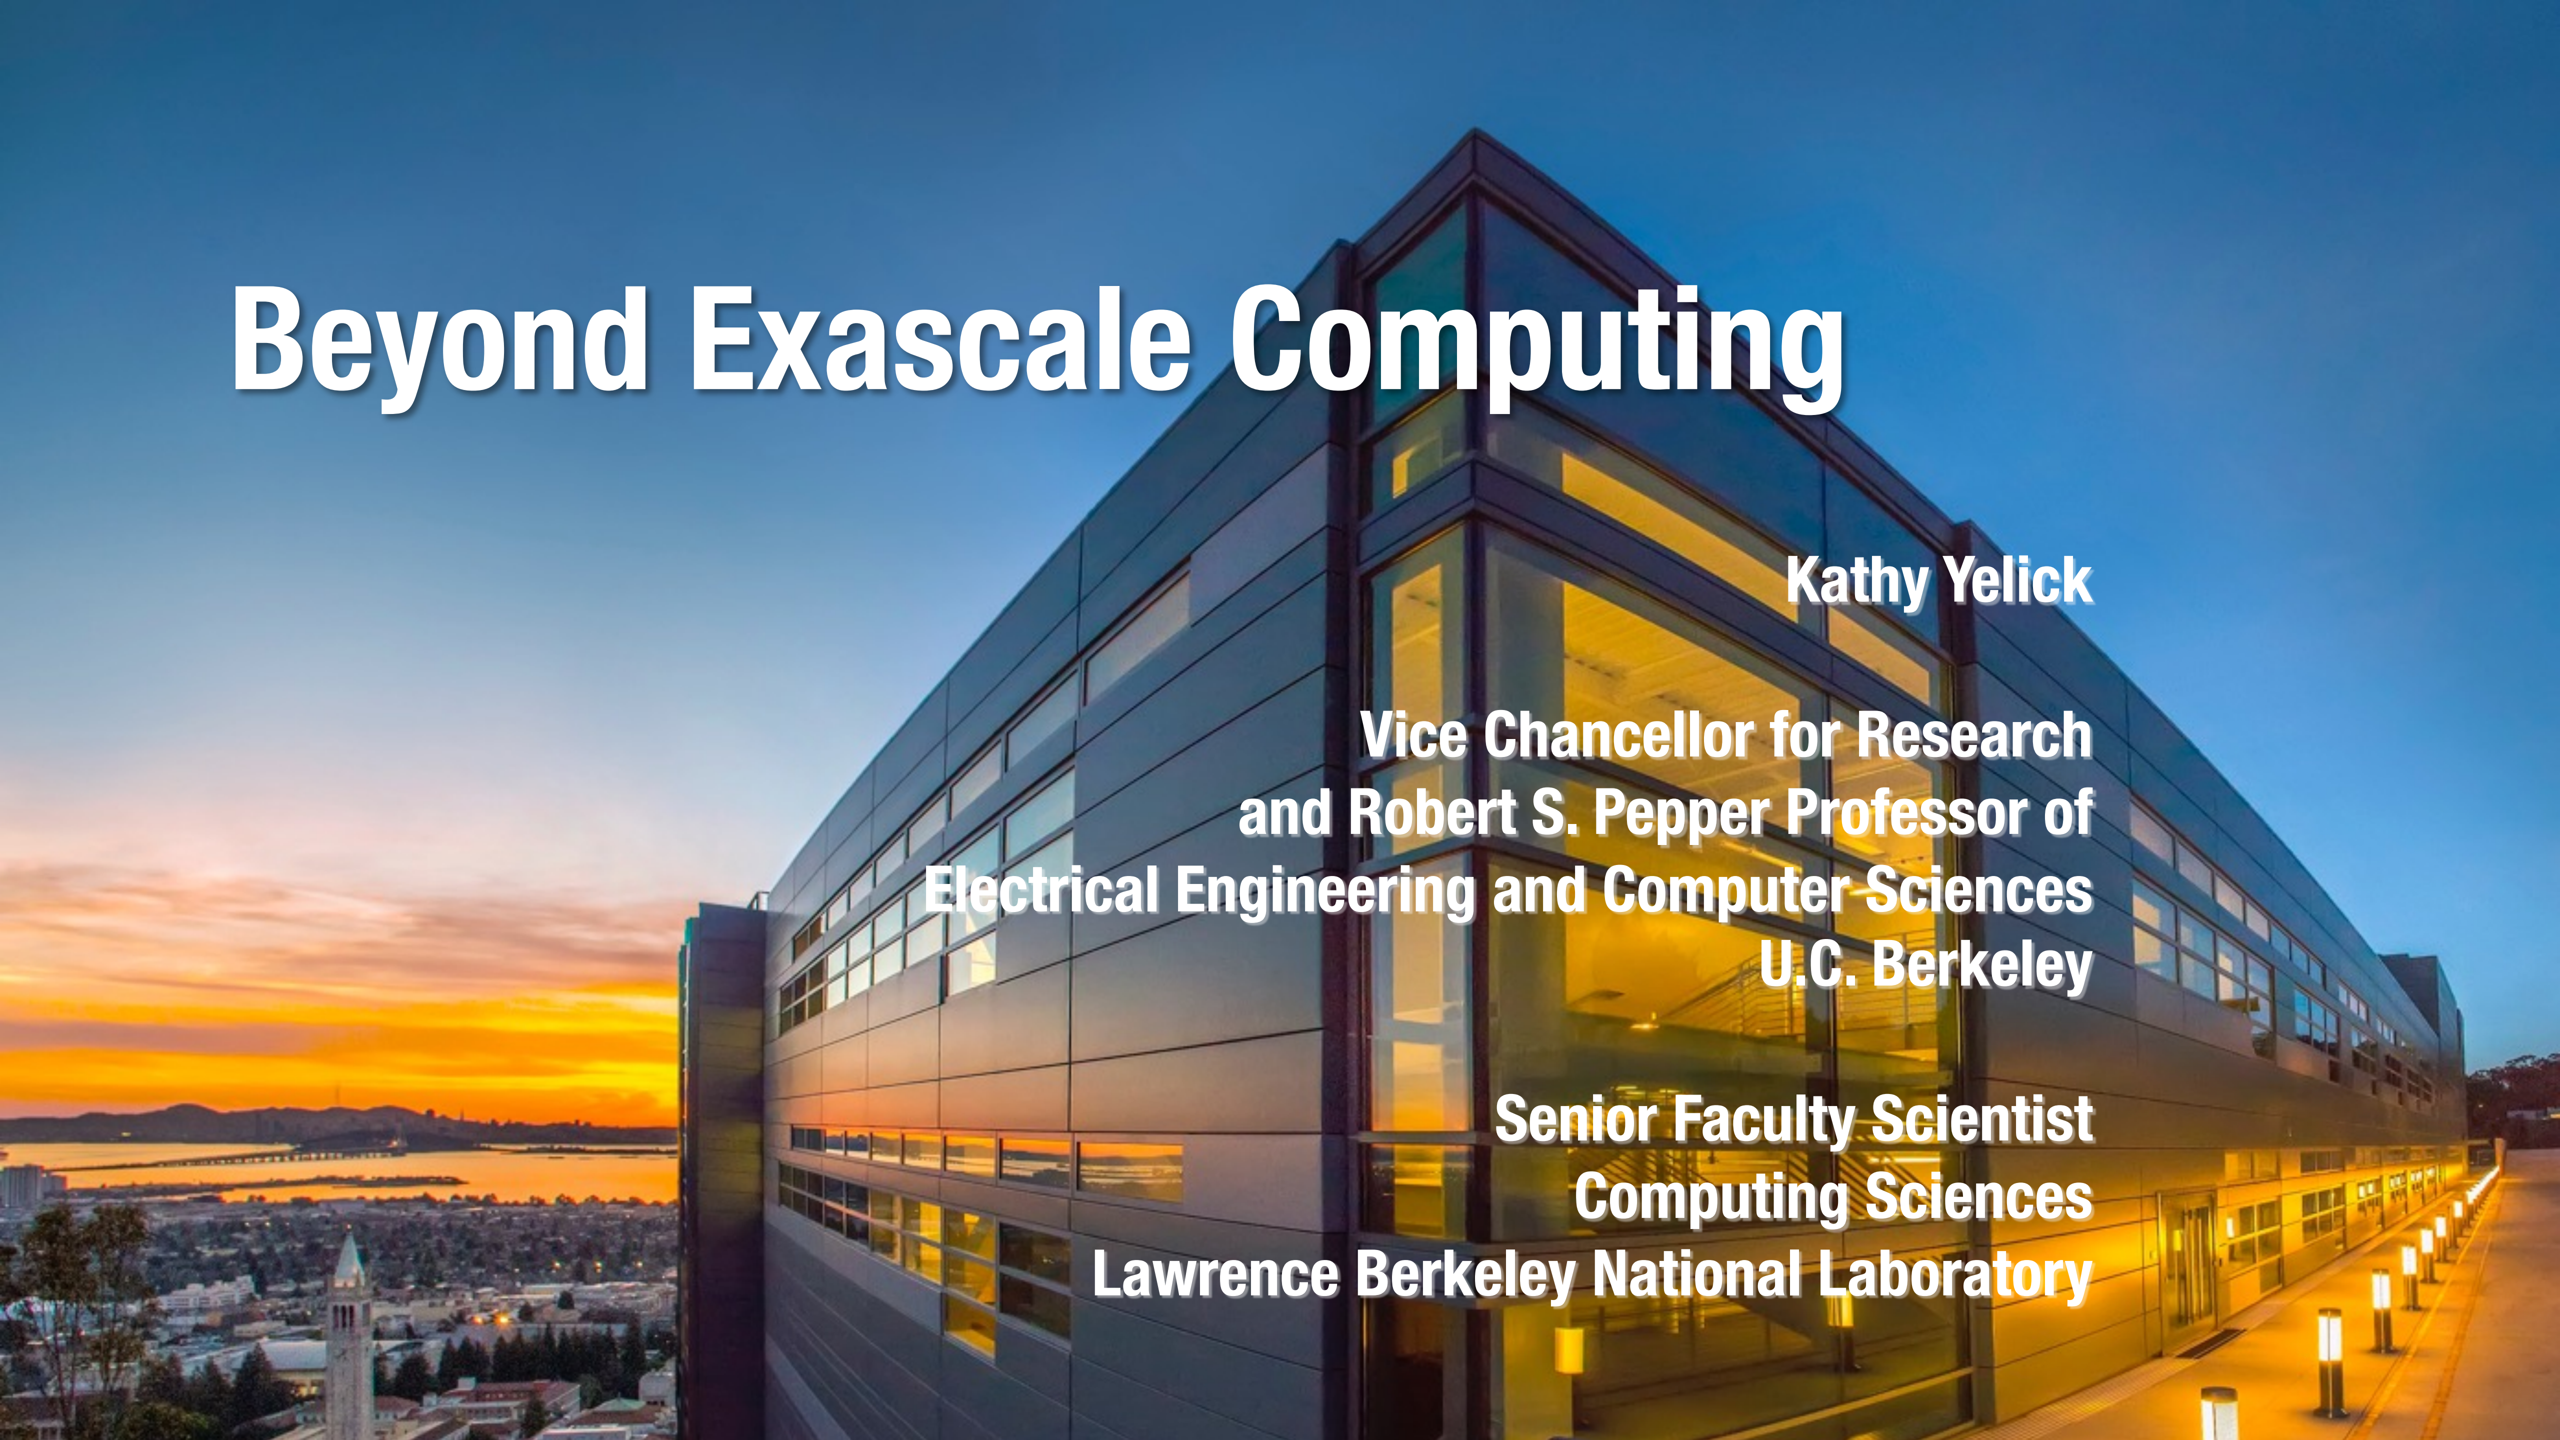

Beyond Exascale Computing

Kathy Yelick

Vice Chancellor for Research
and Robert S. Pepper Professor of
Electrical Engineering and Computer Sciences
U.C. Berkeley

Senior Faculty Scientist
Computing Sciences
Lawrence Berkeley National Laboratory

National Academies Study

Kathy Yelick (Chair)
John Bell
Bill Carlson
Fred Chong
Dona Crawford
Jack Dongarra
Mark Dean
Ian Foster
Charlie McMillan
Dan Meiron
Daniel Reed
Karen Willcox

(report at www.cstb.org)

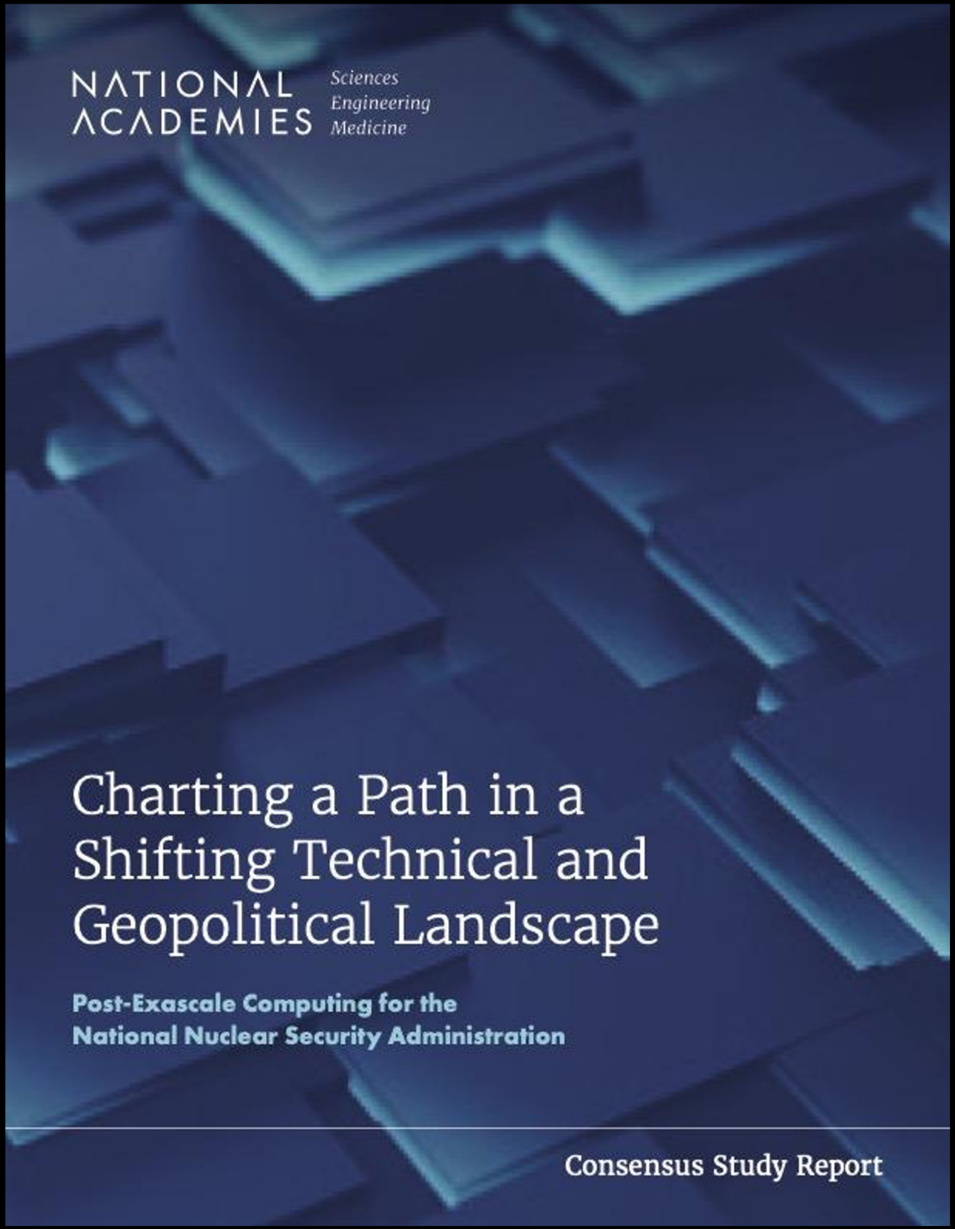

DOE ASCR Report

Jack Dongarra (Chair)
Ewa Deelman (Vice Chair)
Tony Hey
Satoshi Matsuoka
Vivek Sarkar
Greg Bell
Ian Foster
David Keyes
Dieter Kranzlmeuller
Bob Lucas
Lynne Parker
John Shalf
Dan Satanzione
Rick Stevens
Katherine Yelick

(report at www.osti.gov/)

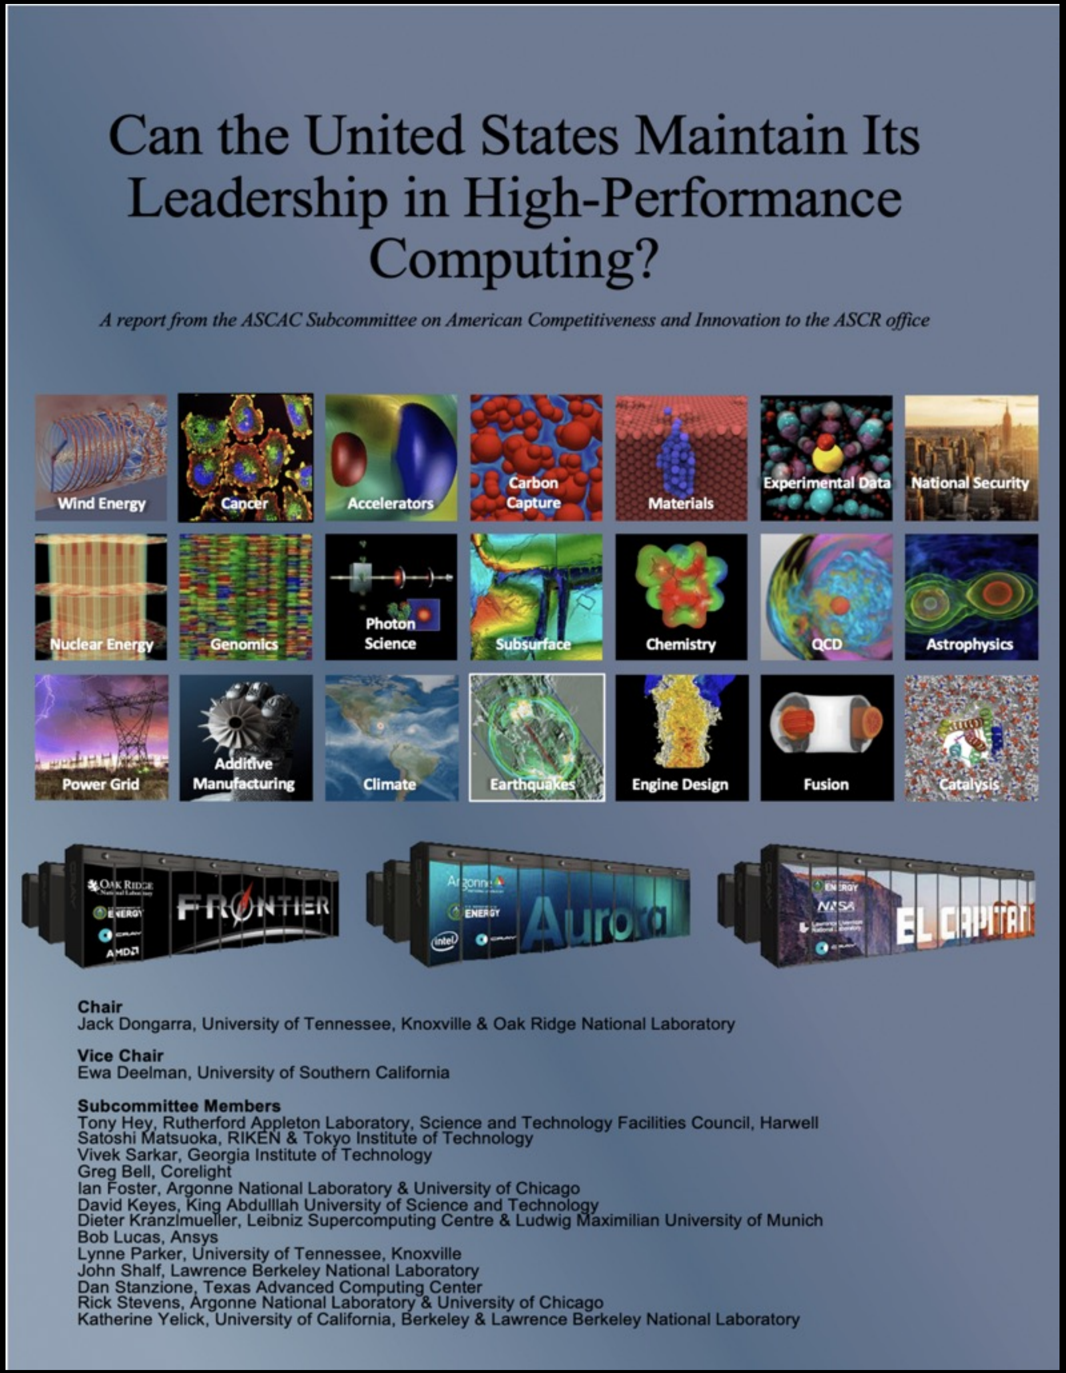

Can the United States Maintain Its Leadership in High-Performance Computing?

A report from the ASCAC Subcommittee on American Competitiveness and Innovation to the ASCR office

Chair
Jack Dongarra, University of Tennessee, Knoxville & Oak Ridge National Laboratory

Vice Chair
Ewa Deelman, University of Southern California

Subcommittee Members
Tony Hey, Rutherford Appleton Laboratory, Science and Technology Facilities Council, Harwell
Satoshi Matsuoka, RIKEN & Tokyo Institute of Technology
Vivek Sarkar, Georgia Institute of Technology
Greg Bell, Corelight
Ian Foster, Argonne National Laboratory & University of Chicago
David Keyes, King Abdulllah University of Science and Technology
Dieter Kranzlmueller, Leibniz Supercomputing Centre & Ludwig Maximilian University of Munich
Bob Lucas, Ansys
Lynne Parker, University of Tennessee, Knoxville
John Shalf, Lawrence Berkeley National Laboratory
Dan Stanzione, Texas Advanced Computing Center
Rick Stevens, Argonne National Laboratory & University of Chicago
Katherine Yelick, University of California, Berkeley & Lawrence Berkeley National Laboratory

Is this Exascale all over again?

Post-Exascale Computing

Computing demand

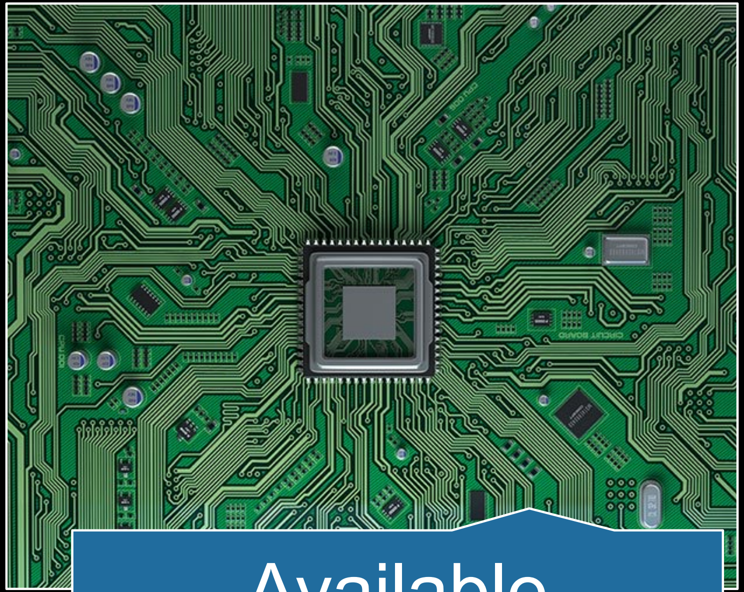

Available technology

Disruptions

Continue to Rethink Applications

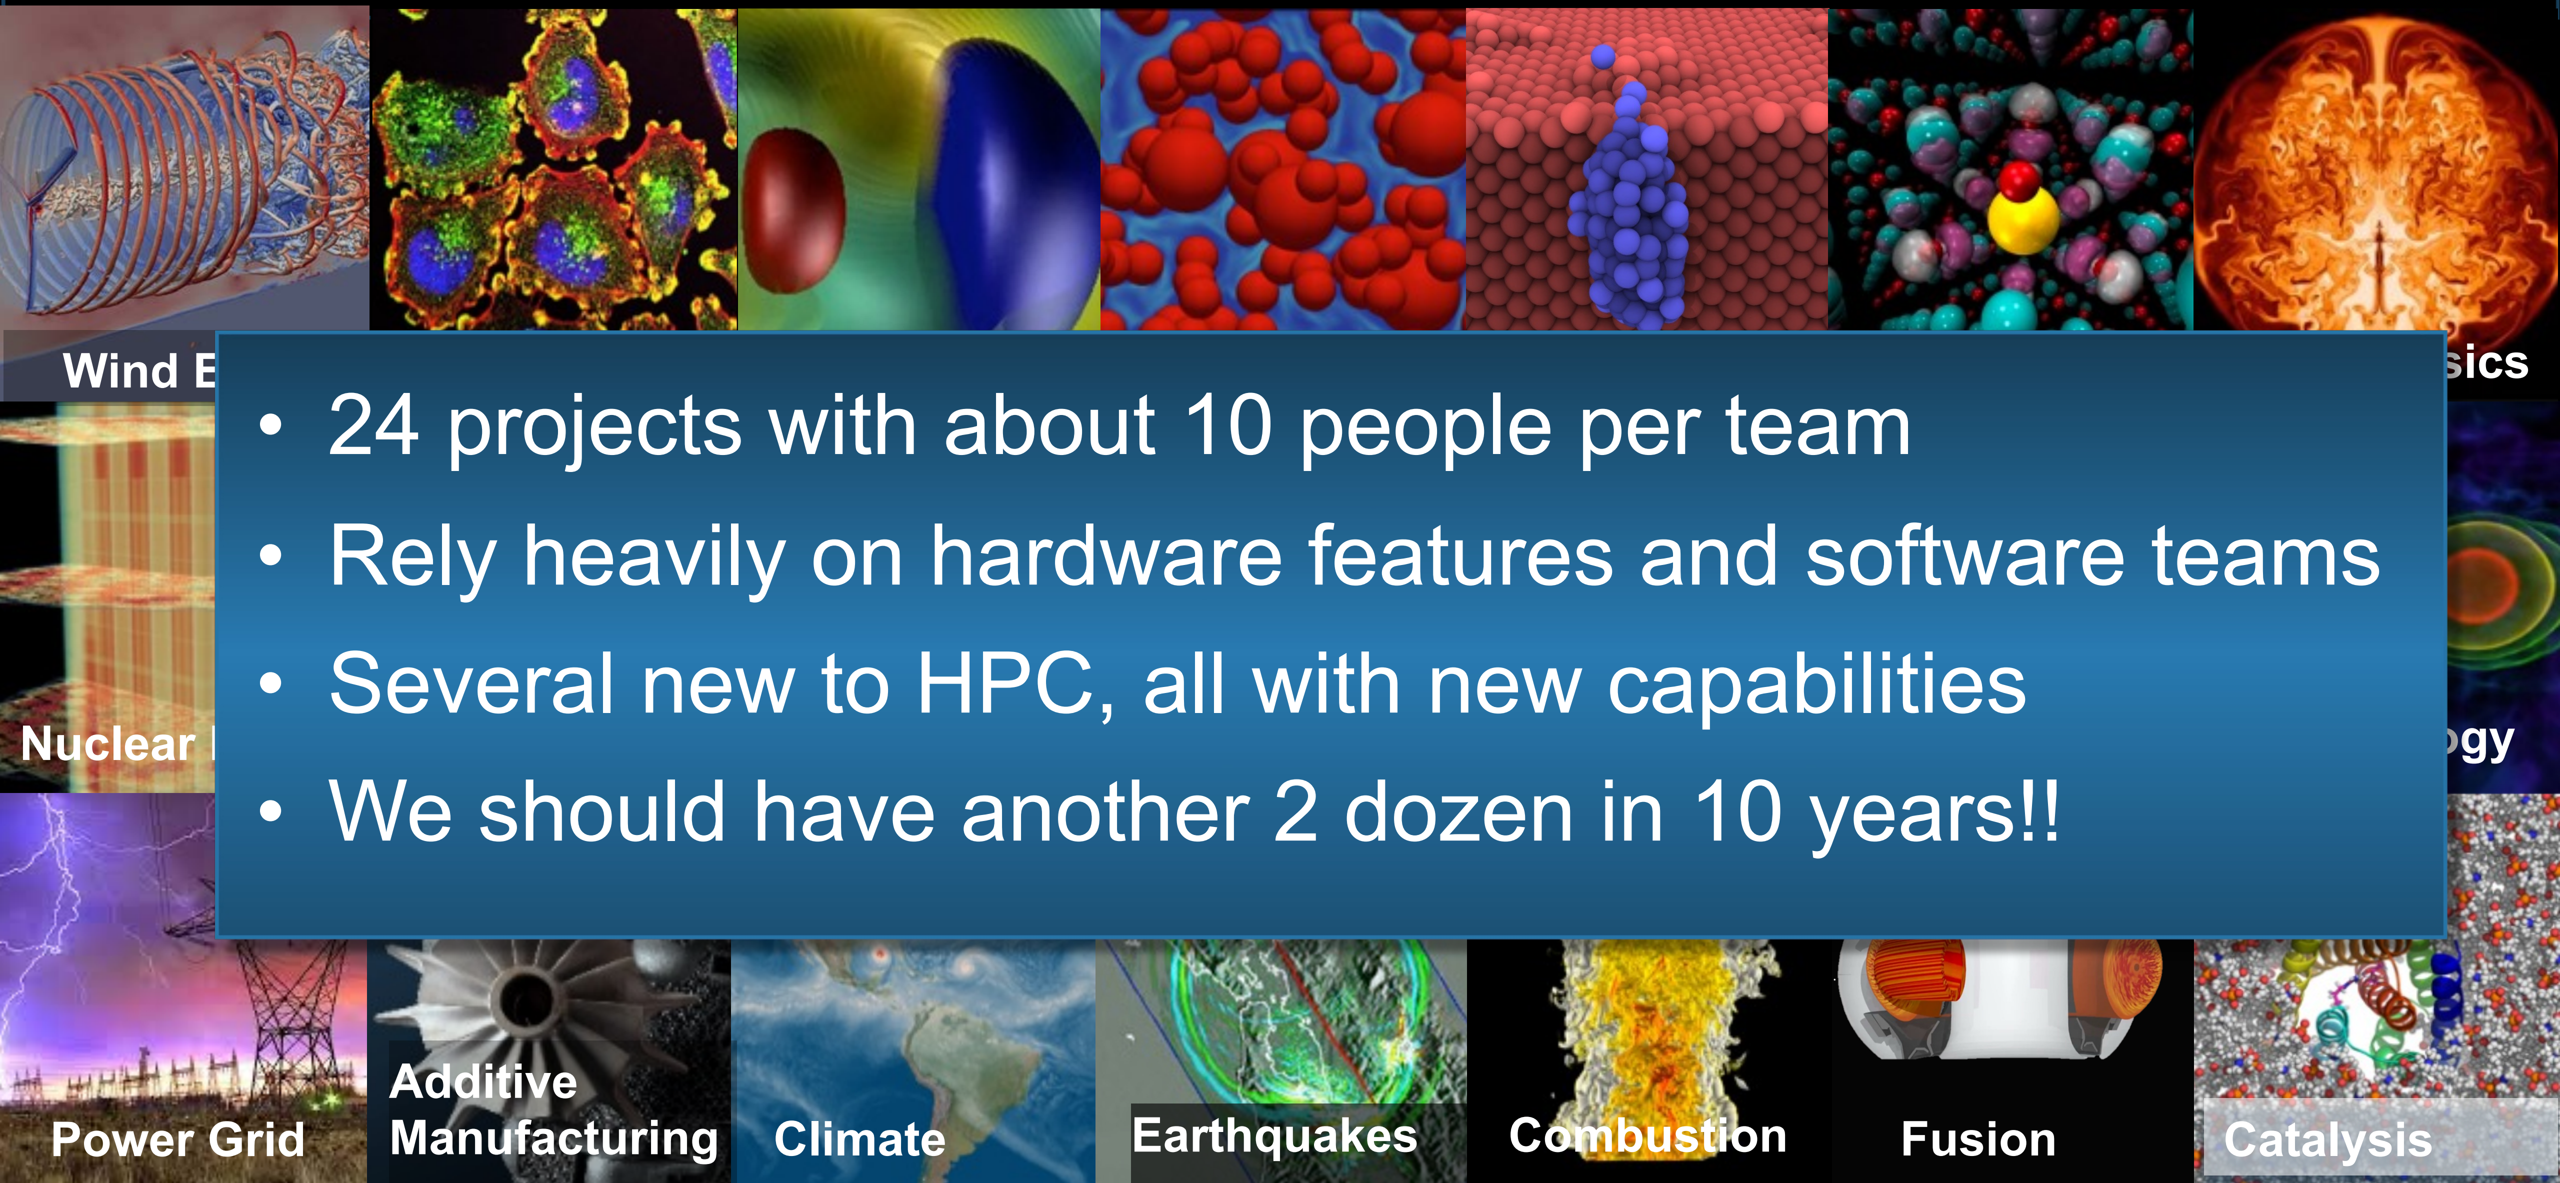

- 24 projects with about 10 people per team
- Rely heavily on hardware features and software teams
- Several new to HPC, all with new capabilities
- We should have another 2 dozen in 10 years!!

Scientific Computing Circa 2007

Exascale report from 2007 Town Halls
Entirely focused on modeling and simulation

~~Scientific Computing is often used synonymously with Simulation and HPC~~

Simulation ⊂ Scientific Computing ⊂ HPC

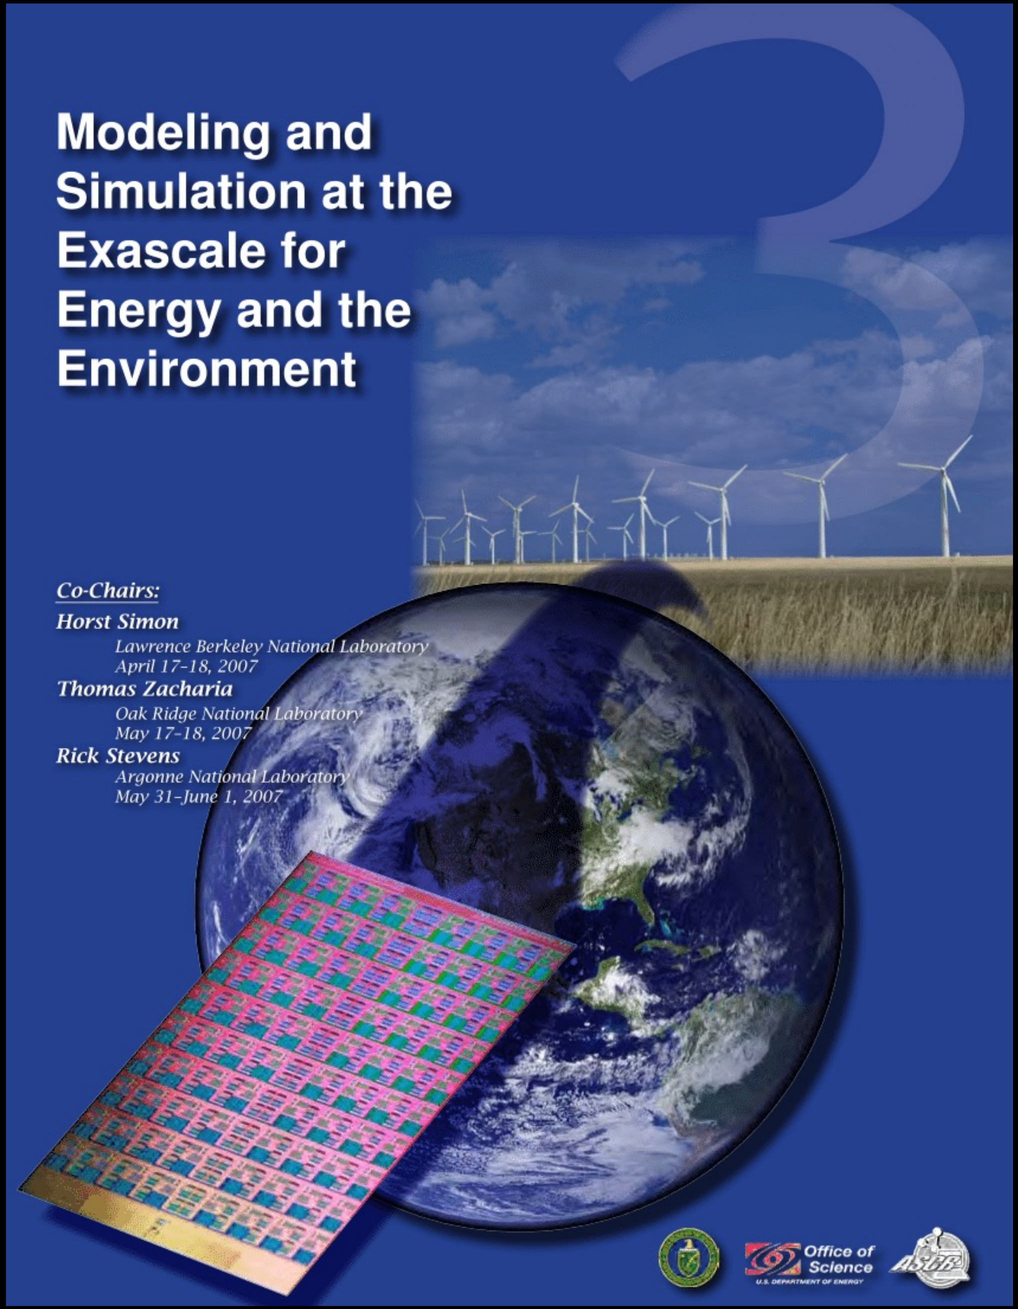

Runtime of “hero” calculations are too long

Number of Nodes	Memory Footprint	Wall-Clock Time
2400	~300–400 TB	6 months
4990	~600 TB	3–4 months
288	~20 TB	1 month
3250	104 TB	5.8 days
512	32.8 TB	2 months

Iterative design does not happen on 6 month cycles

Weak Scaling has Diminishing Returns

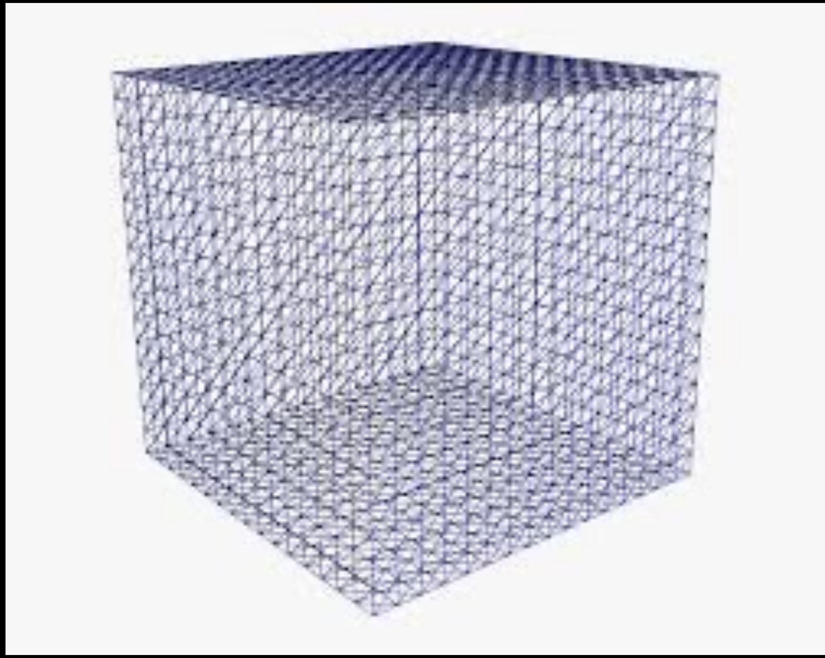

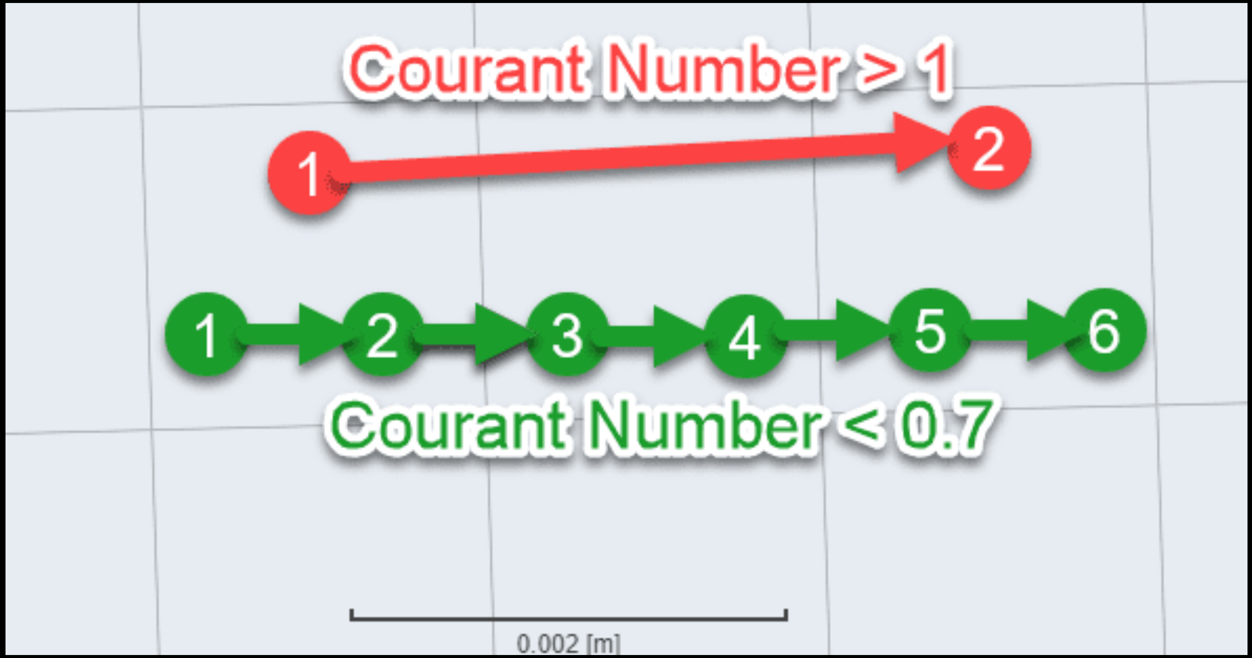

Increase resolution by 10x in each dimension

Increase cores by 1000x

Runtime increases ☹

New demands for HPC in Science

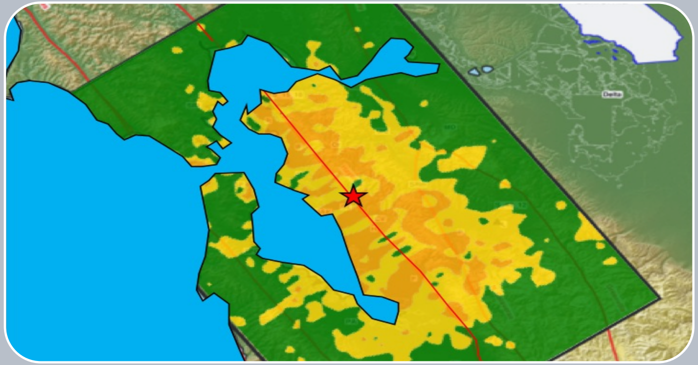

Simulation

From atoms to the universe

Data

Images, text, to genomes

Learning

Interpret, infer and automate

JGI User Science and ExaBiome's MetaHipMer Assembler

Peatland bog (0.6 TB)

Mycorrhizal fungi (1.5 TB)

Prescribed fire (1.5 TB)

Subtropical soils (1.5 TB)

Dark matter fungi (1.7 TB)

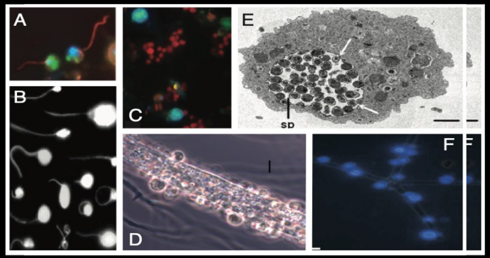

Coastal mangroves (2.1 TB)

Mountainous watershed (2.7 TB)

Soil carbon cycle (3.3 TB)

Great Redox Experiment (8TB)

Lake Mendota Time Series (25 TB)

Analyzing huge data sets with exascale tools lead to discovery of new species and protein families … new science!

Machine Learning Drives Computational Demand

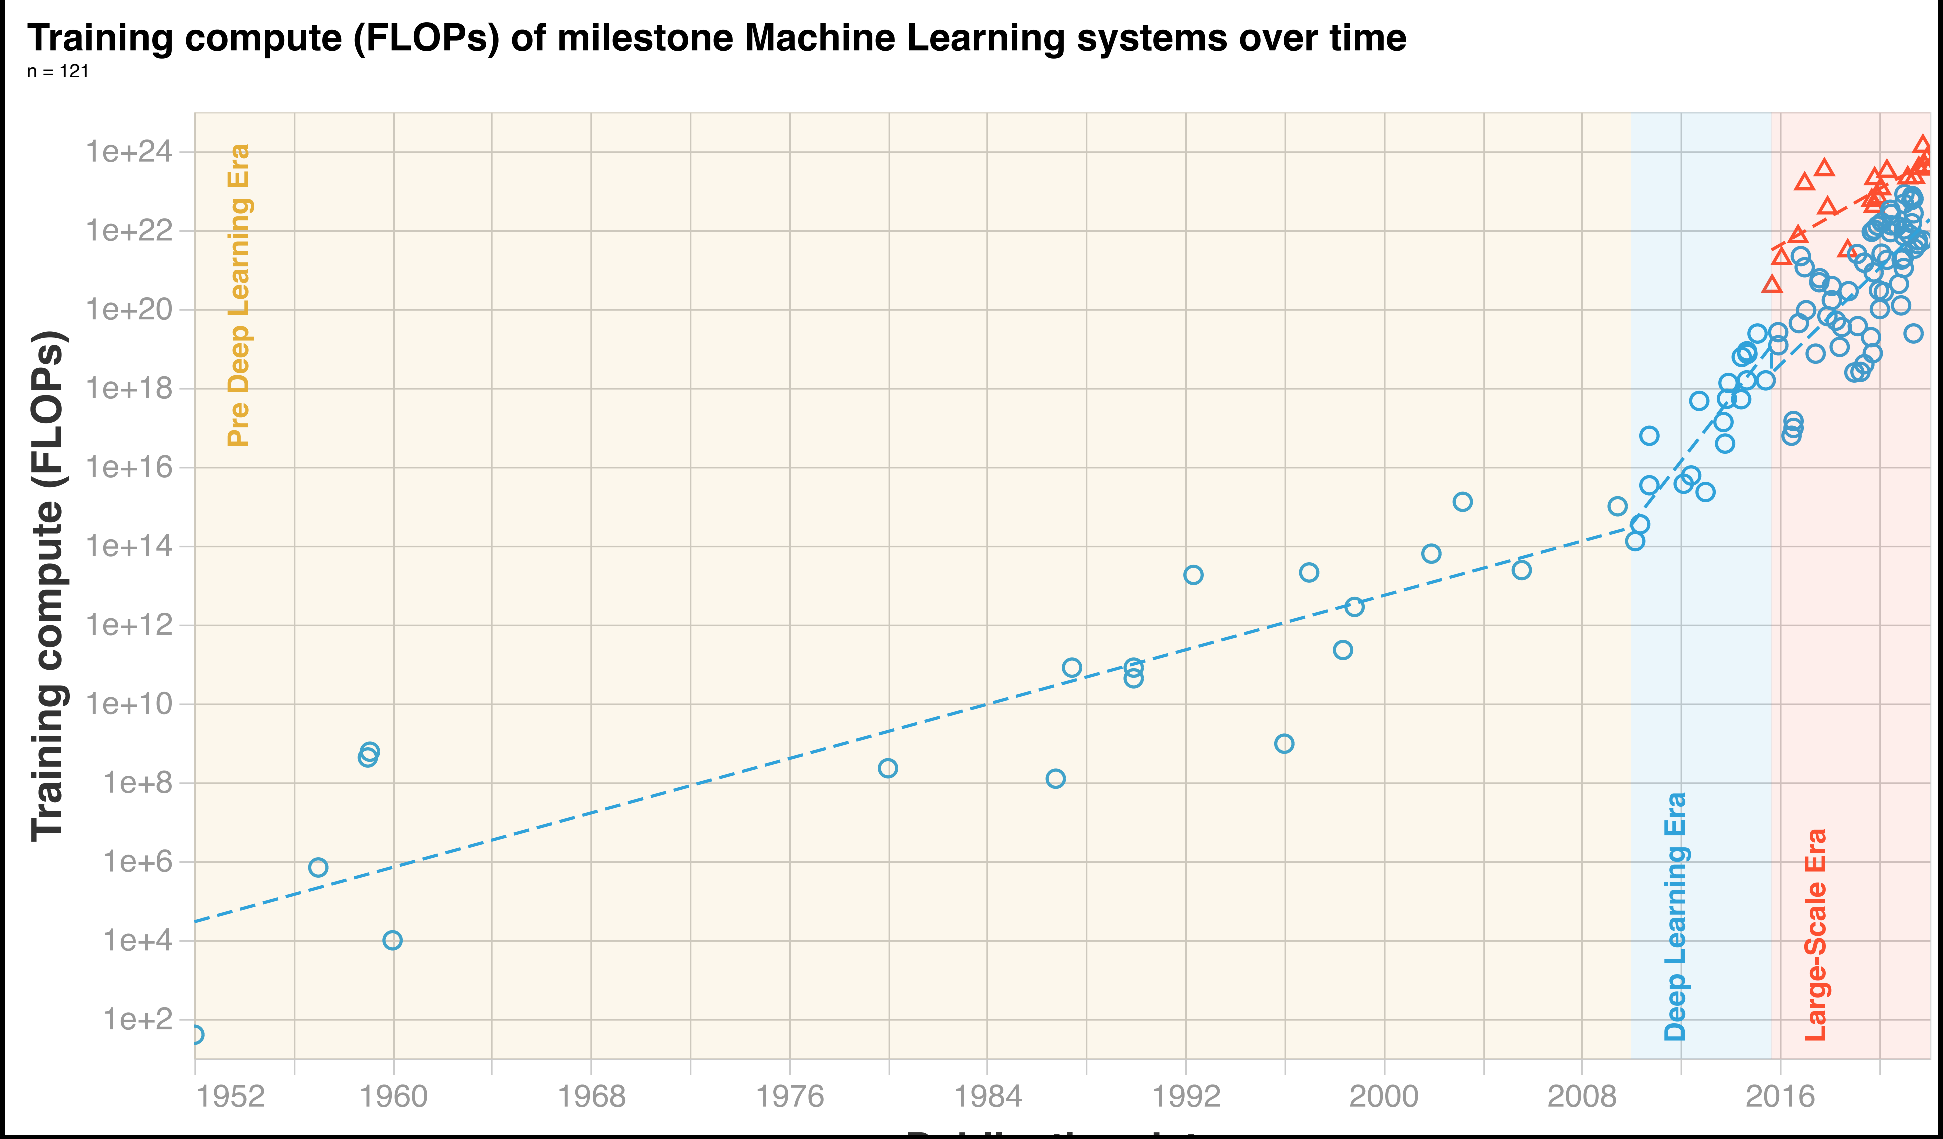

Computing Requirements in Machine Learning

300,000x increase from 2011 (AlexNet) to 2018 (AlphaGoZero)

A petaflop/s-day = 10^{15} neural net operations per second for one day, ~= 10^{20} operations

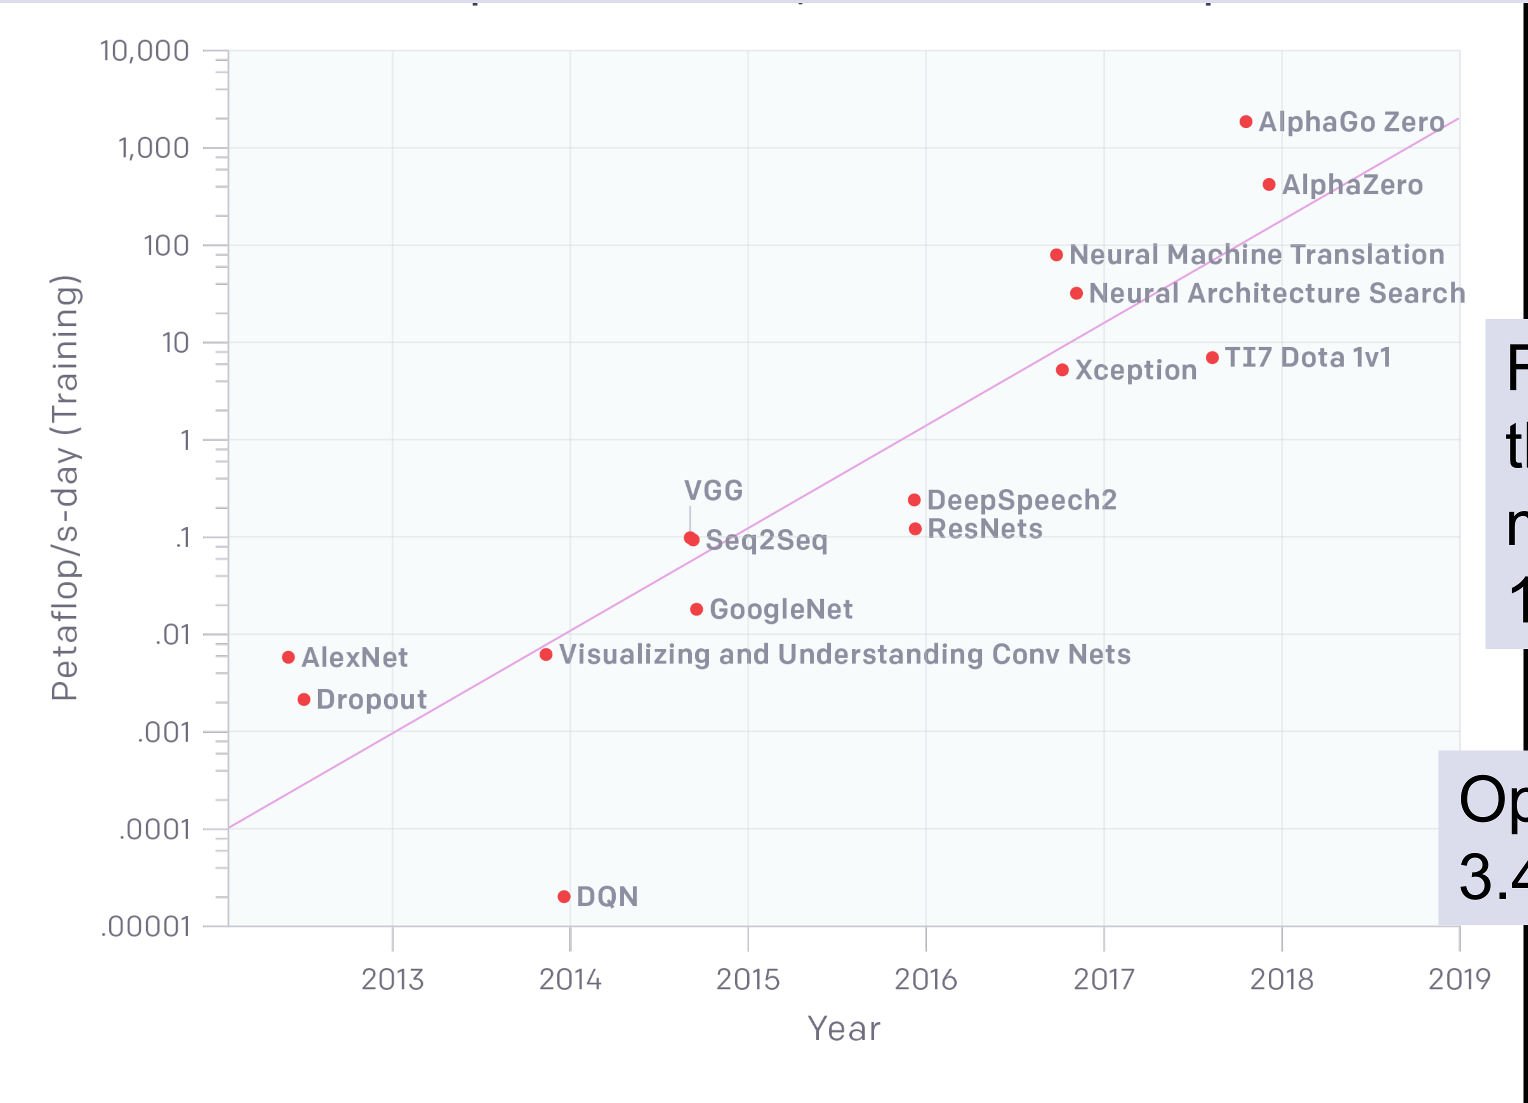

From 2011-2018 the fastest Top500 machine grew < 15x

OpenAI estimates 3.4-month doubling!

https://blog.openai.com/ai-and-compute/

Reproducibility and Efficiency in Science through Cloud Labs

"Plugging an experiment into a browser forces researchers to translate the exact details of every step into unambiguous code"

https://www.theguardian.com/

Automation in Self-Driving Laboratories

E.g., Strateos Cloud Lab
14K square feed
200+ instruments

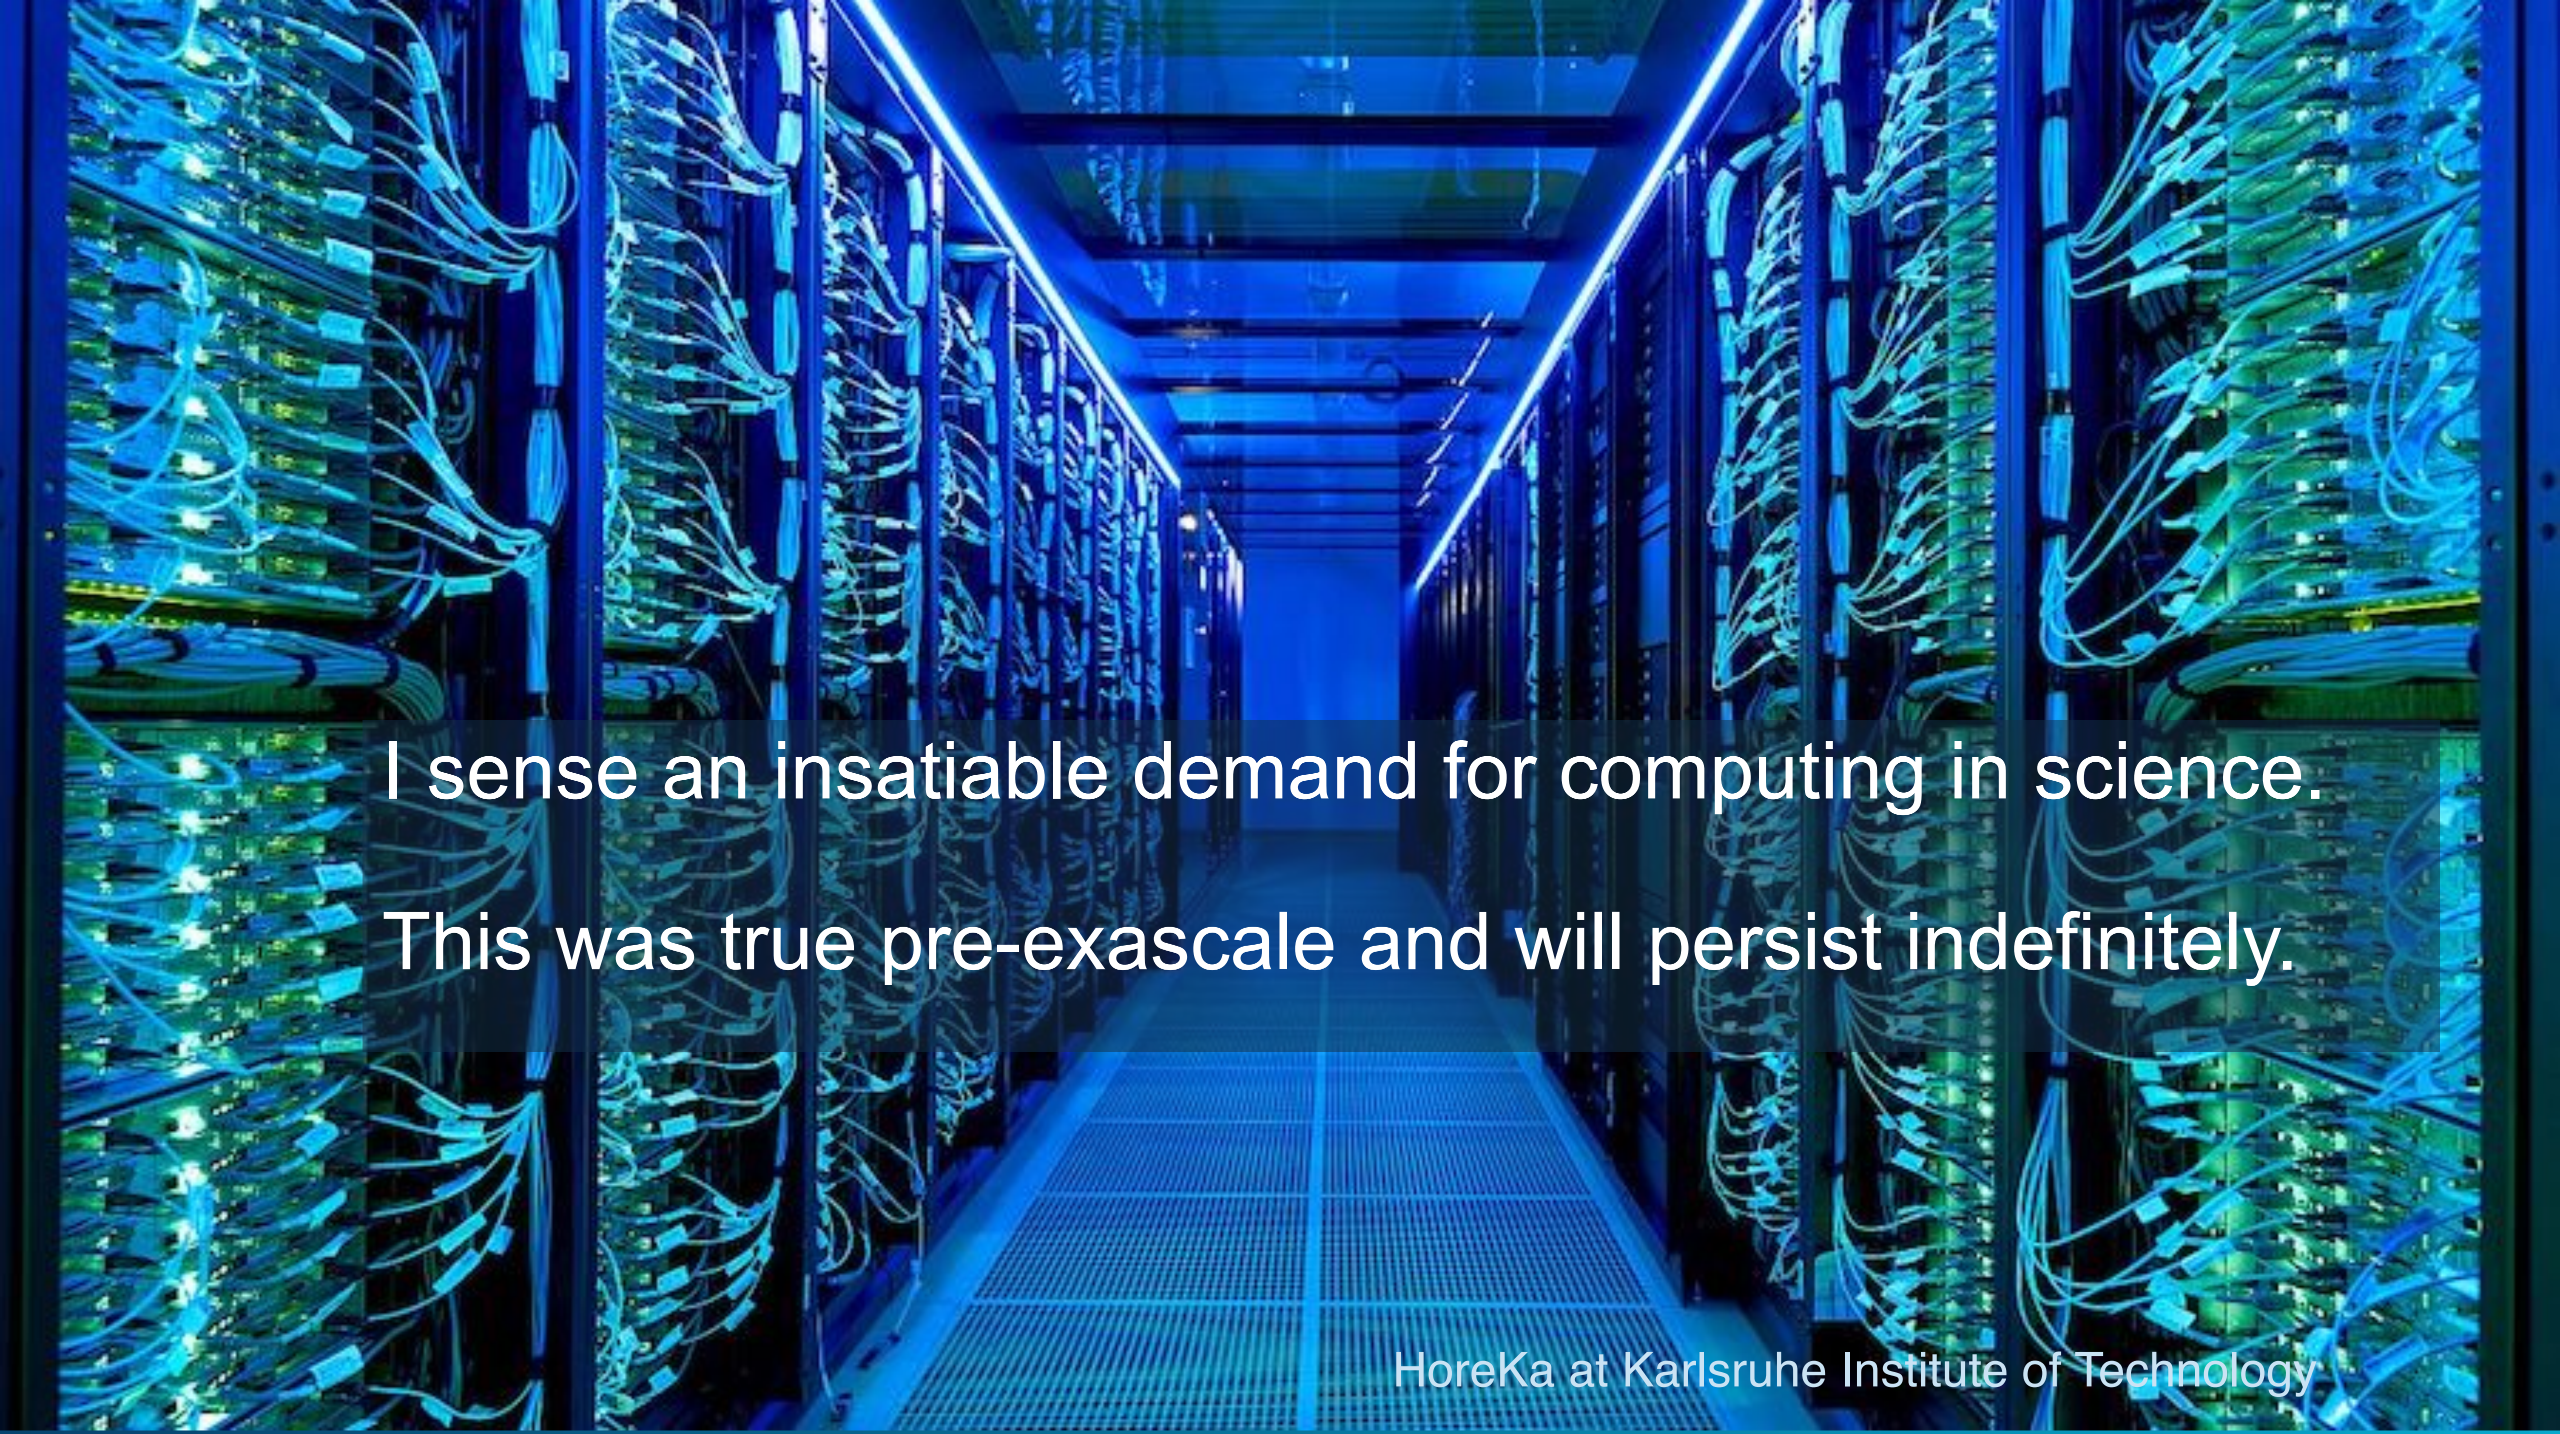

HoreKa at Karlsruhe Institute of Technology

Post-Exascale Computing

Computing demand

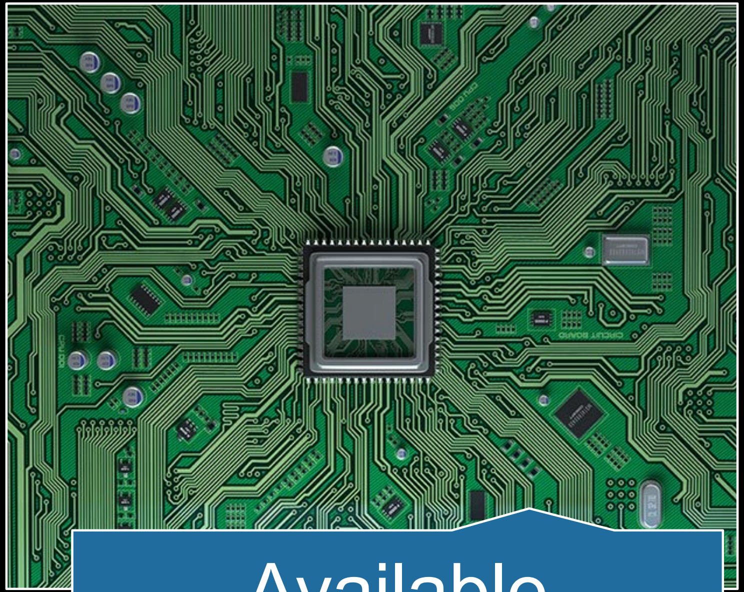

Available technology

Disruptions

122 YEARS OF MOORE'S LAW

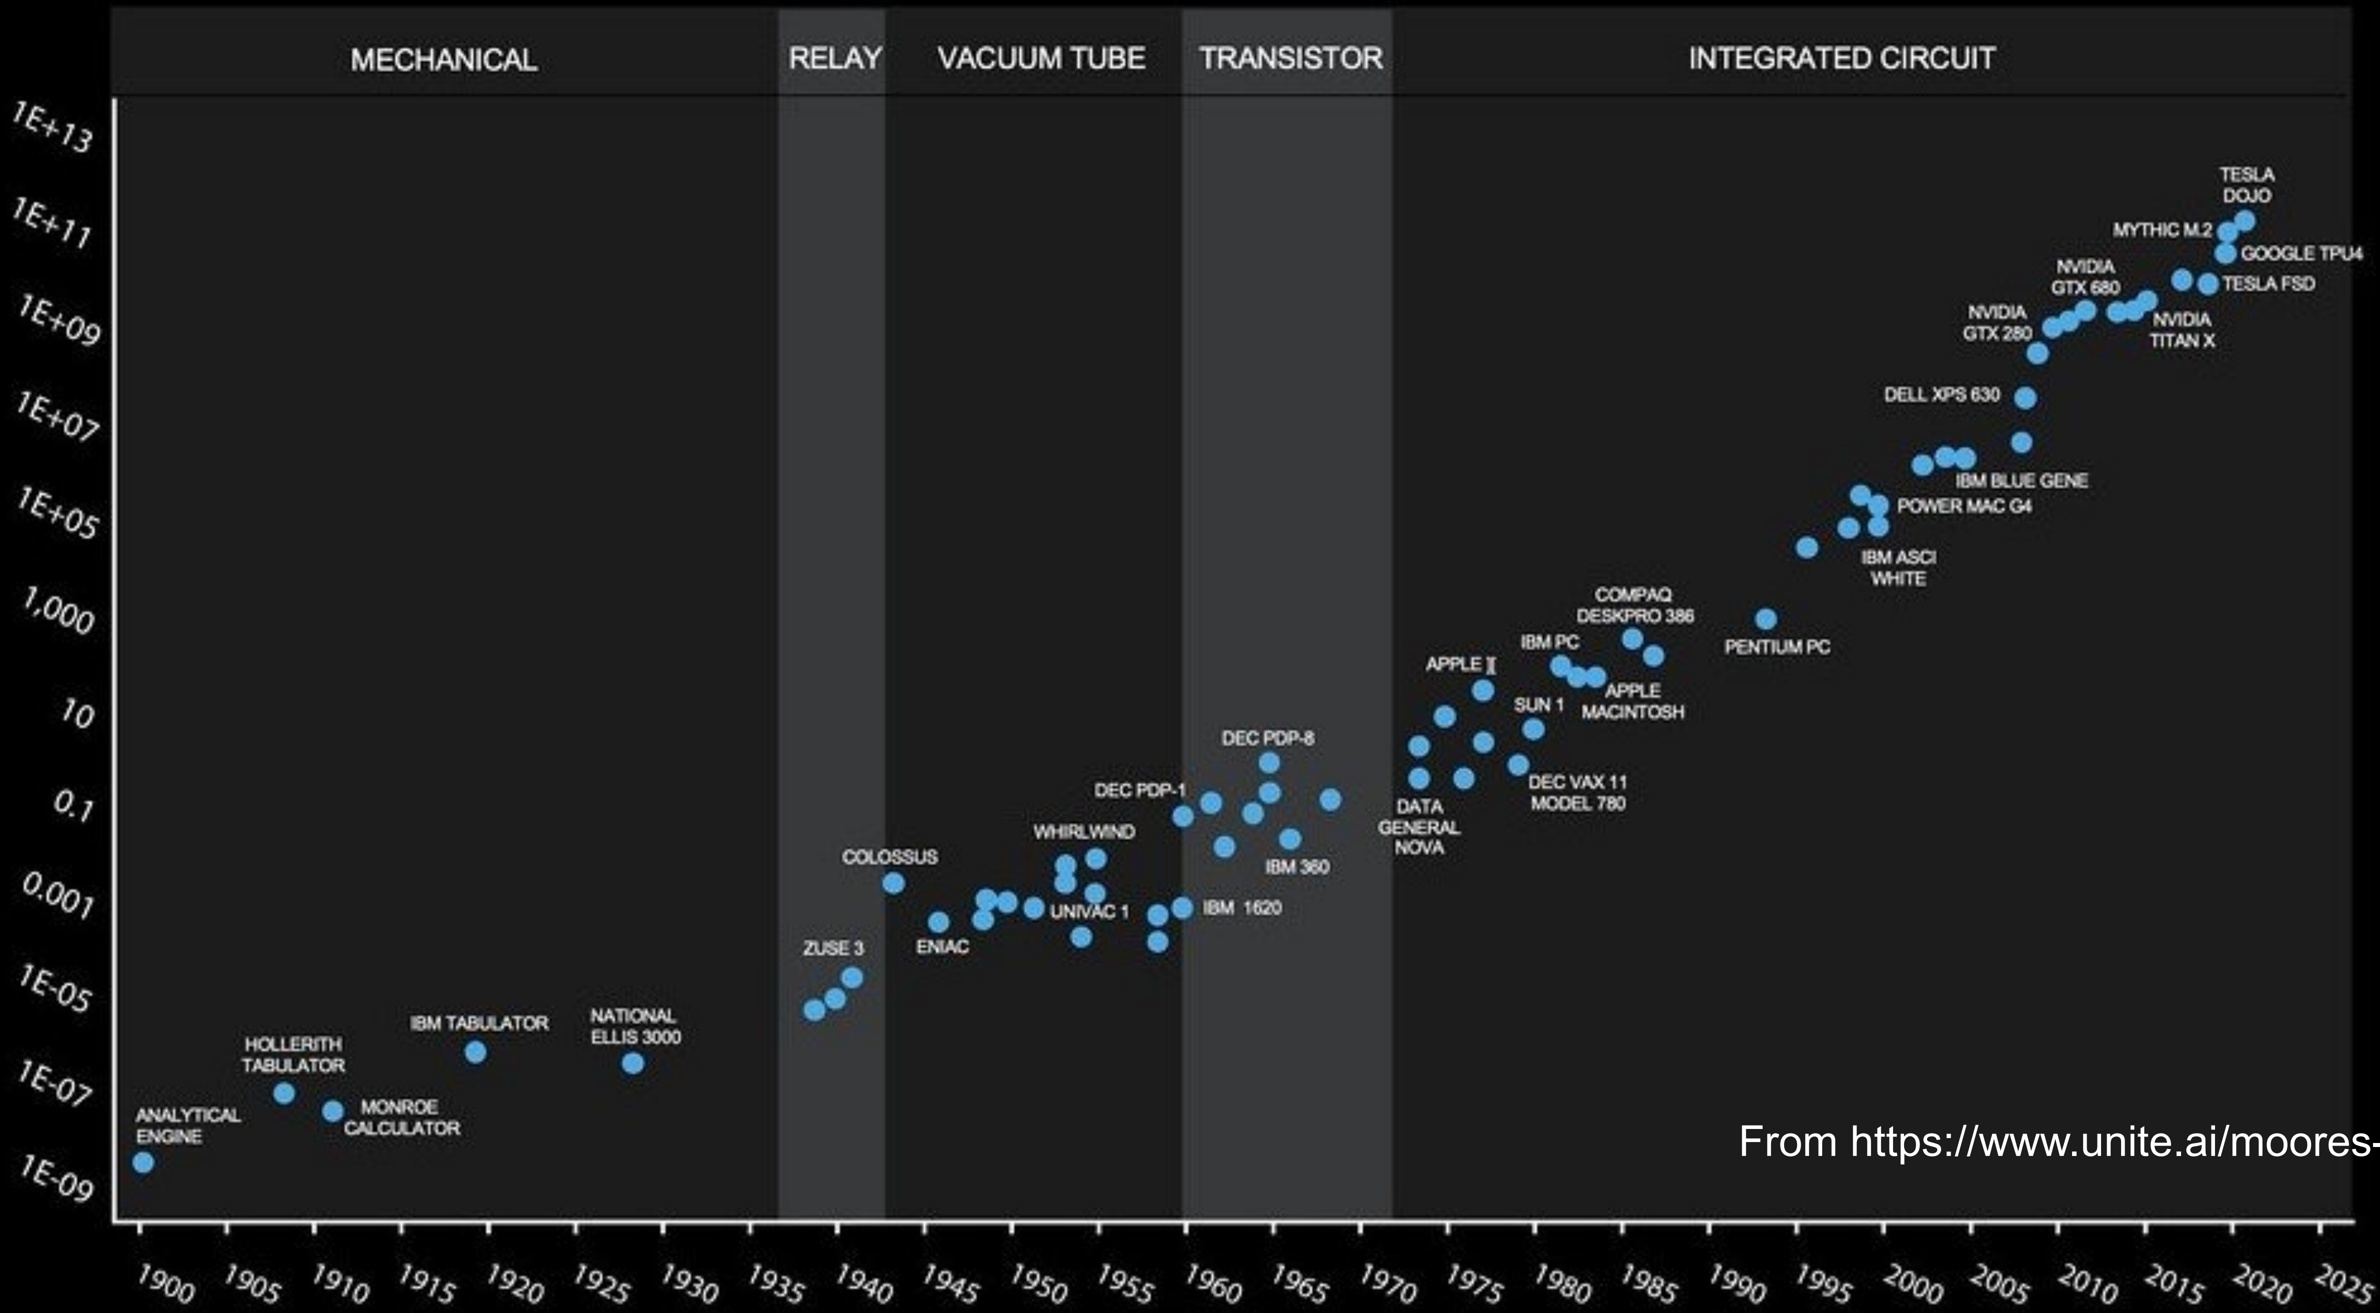

From https://www.unite.ai/moores-law/

10 µm	6 µm	3 µm	1.5 µm	1 µm	800 nm	600 nm	350 nm	250 nm	180 nm	130 nm	90 nm	65 nm	45 nm	32 nm	22 nm	14 nm	10 nm	7 nm	5 nm	3 nm

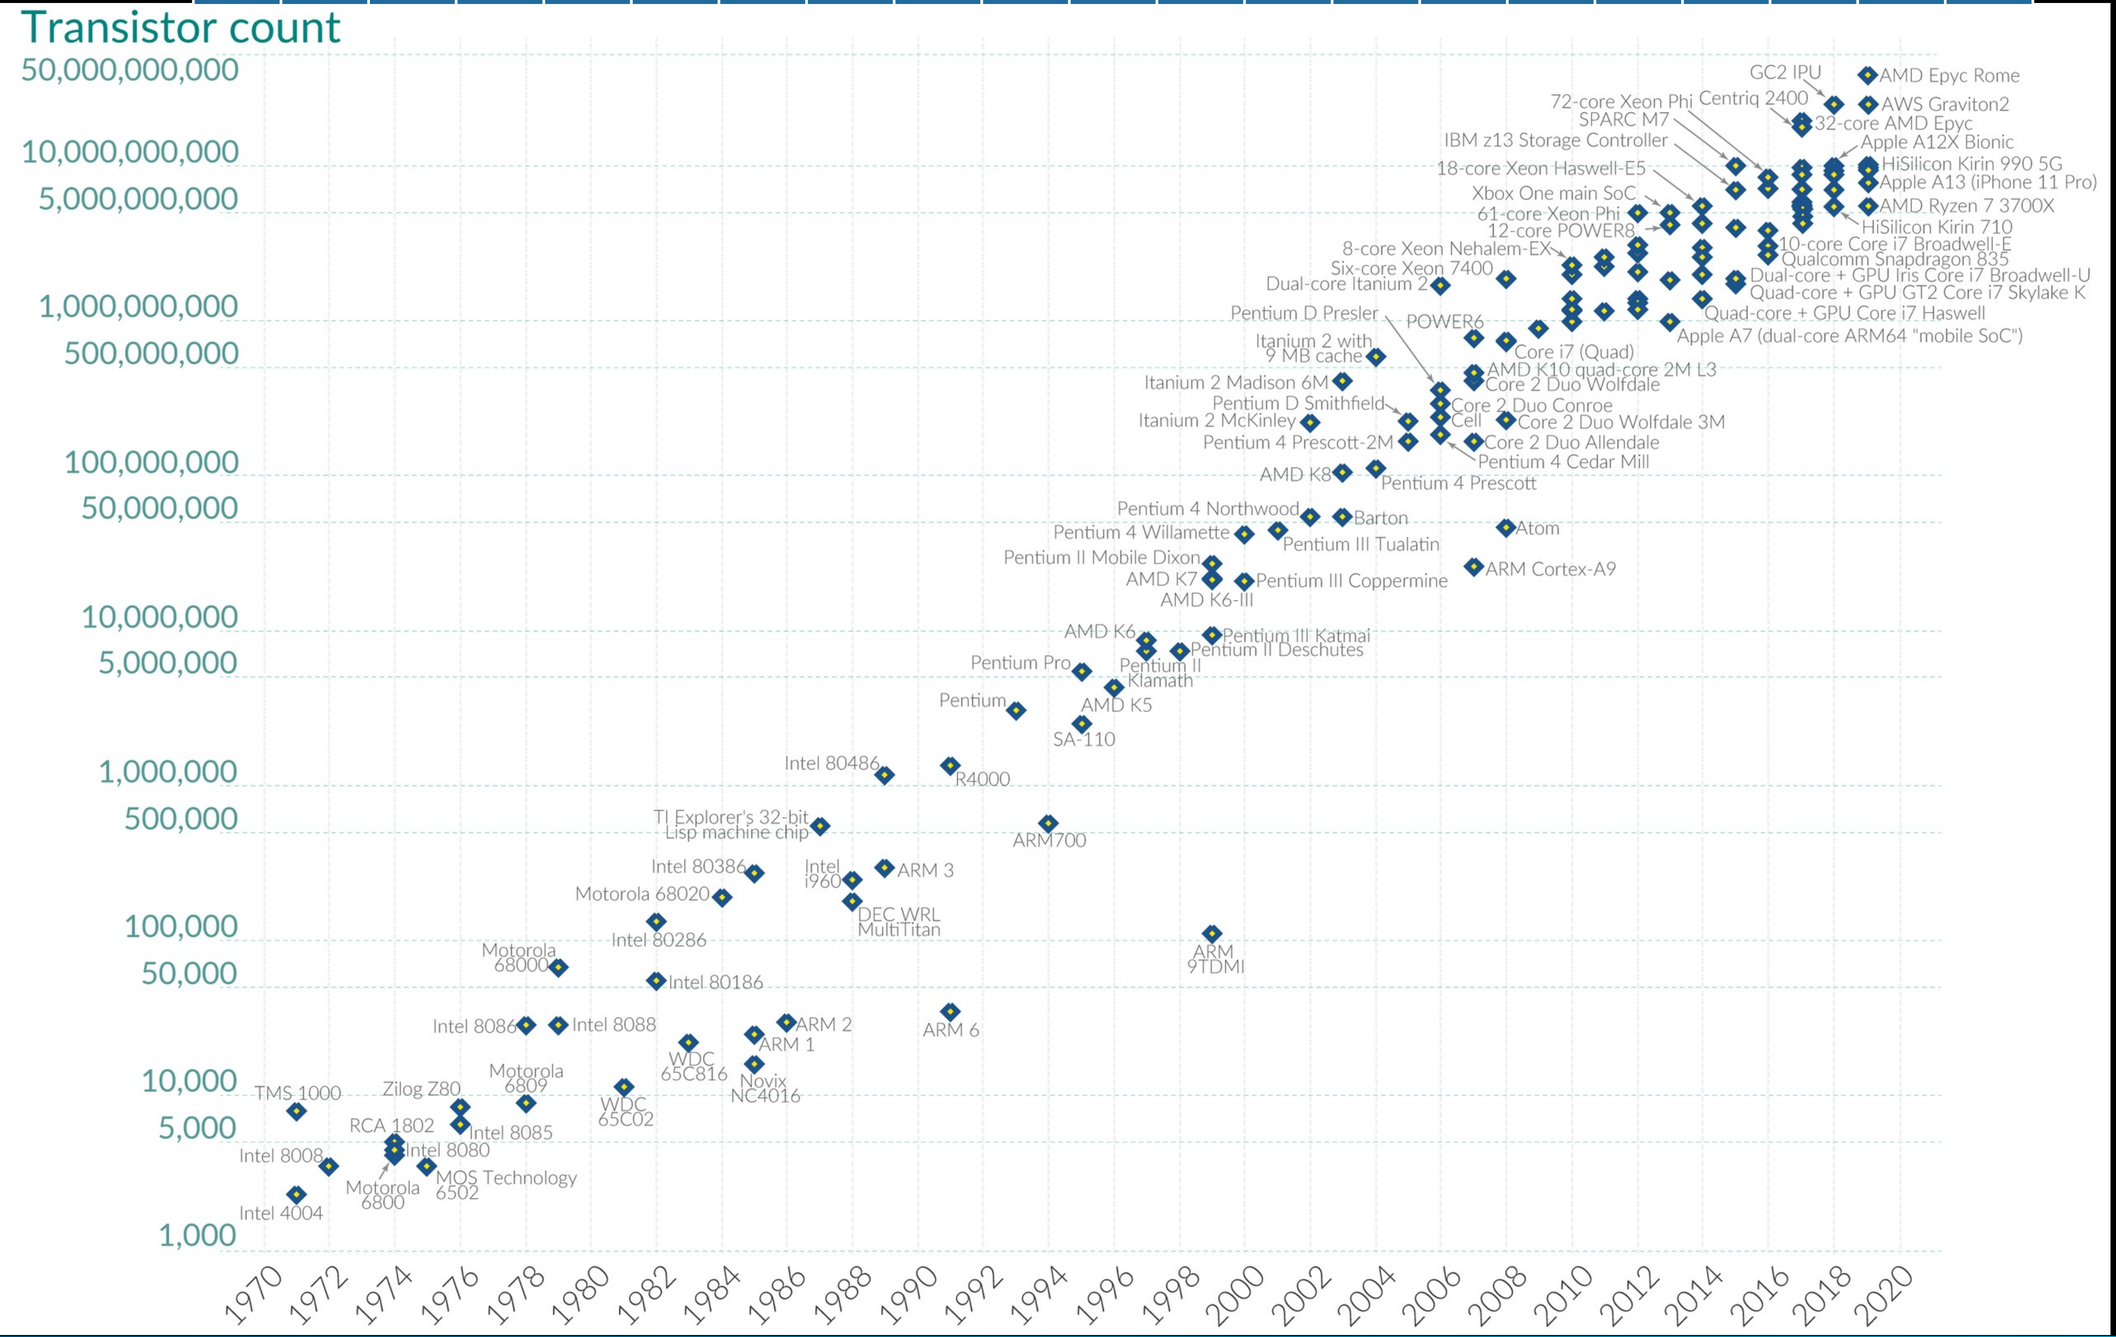

It's hard to think exponentially

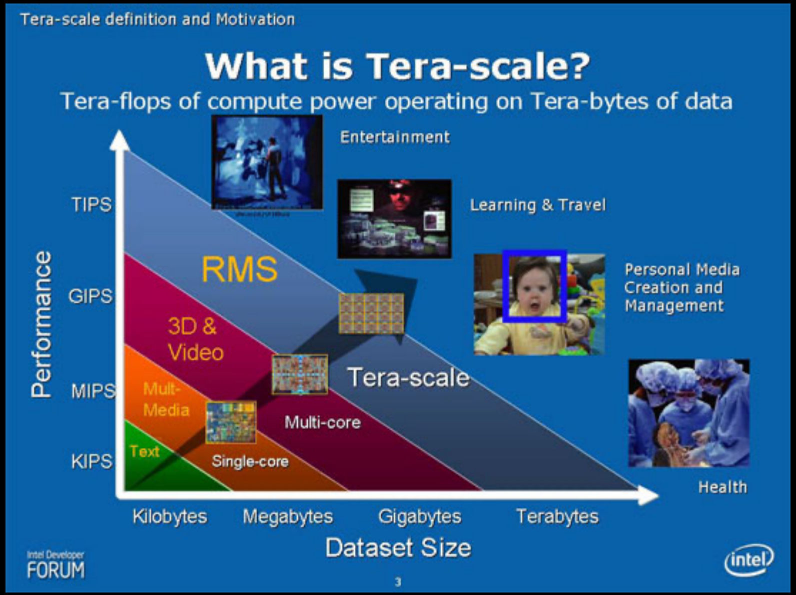

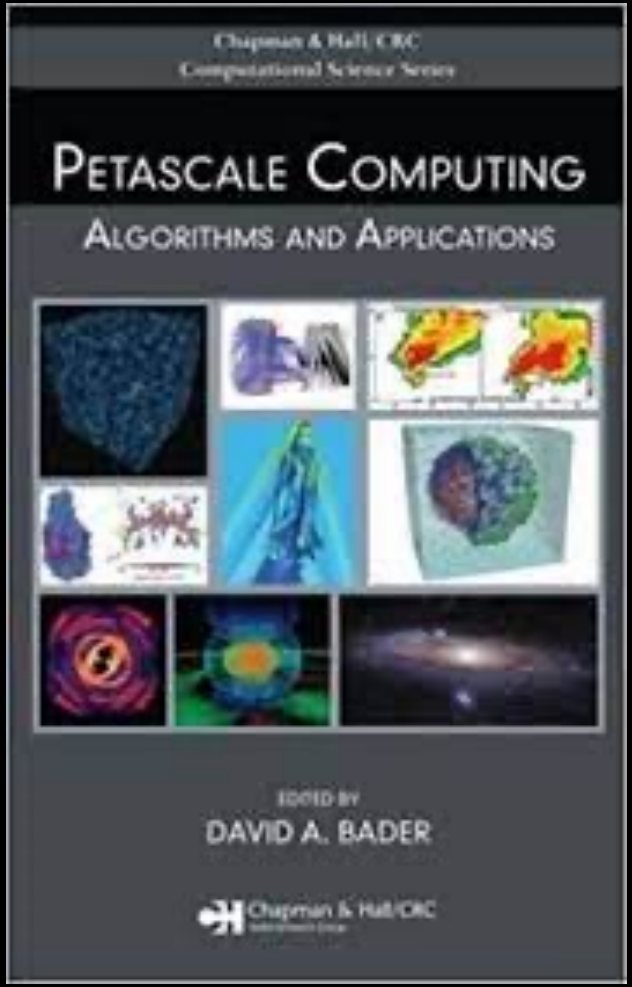

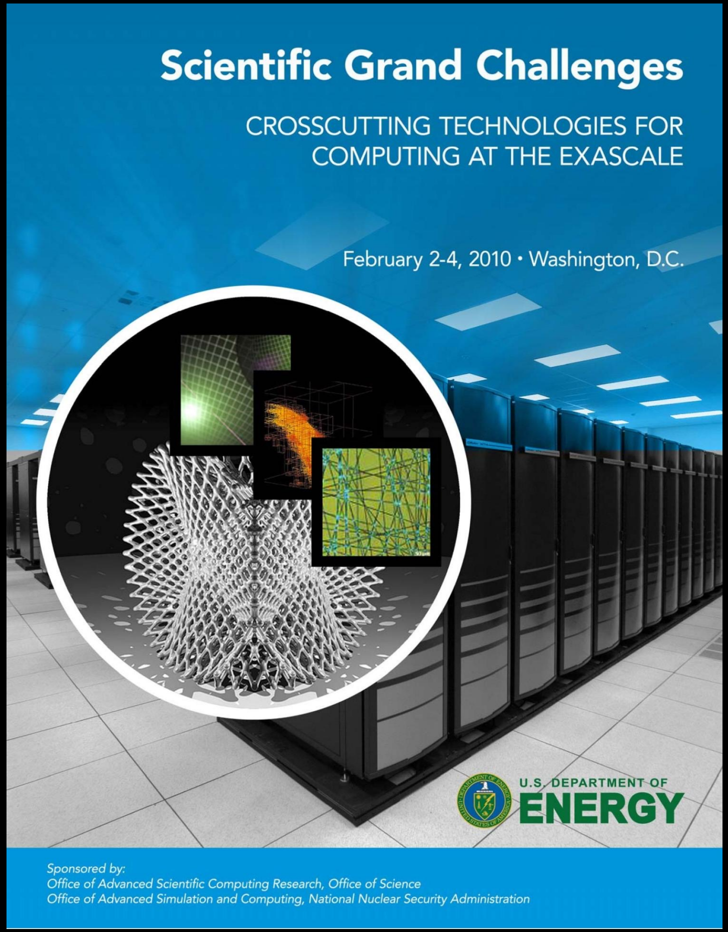

But it’s also hard to stop

Dennard Scaling is Dead; Moore's Law Will Follow

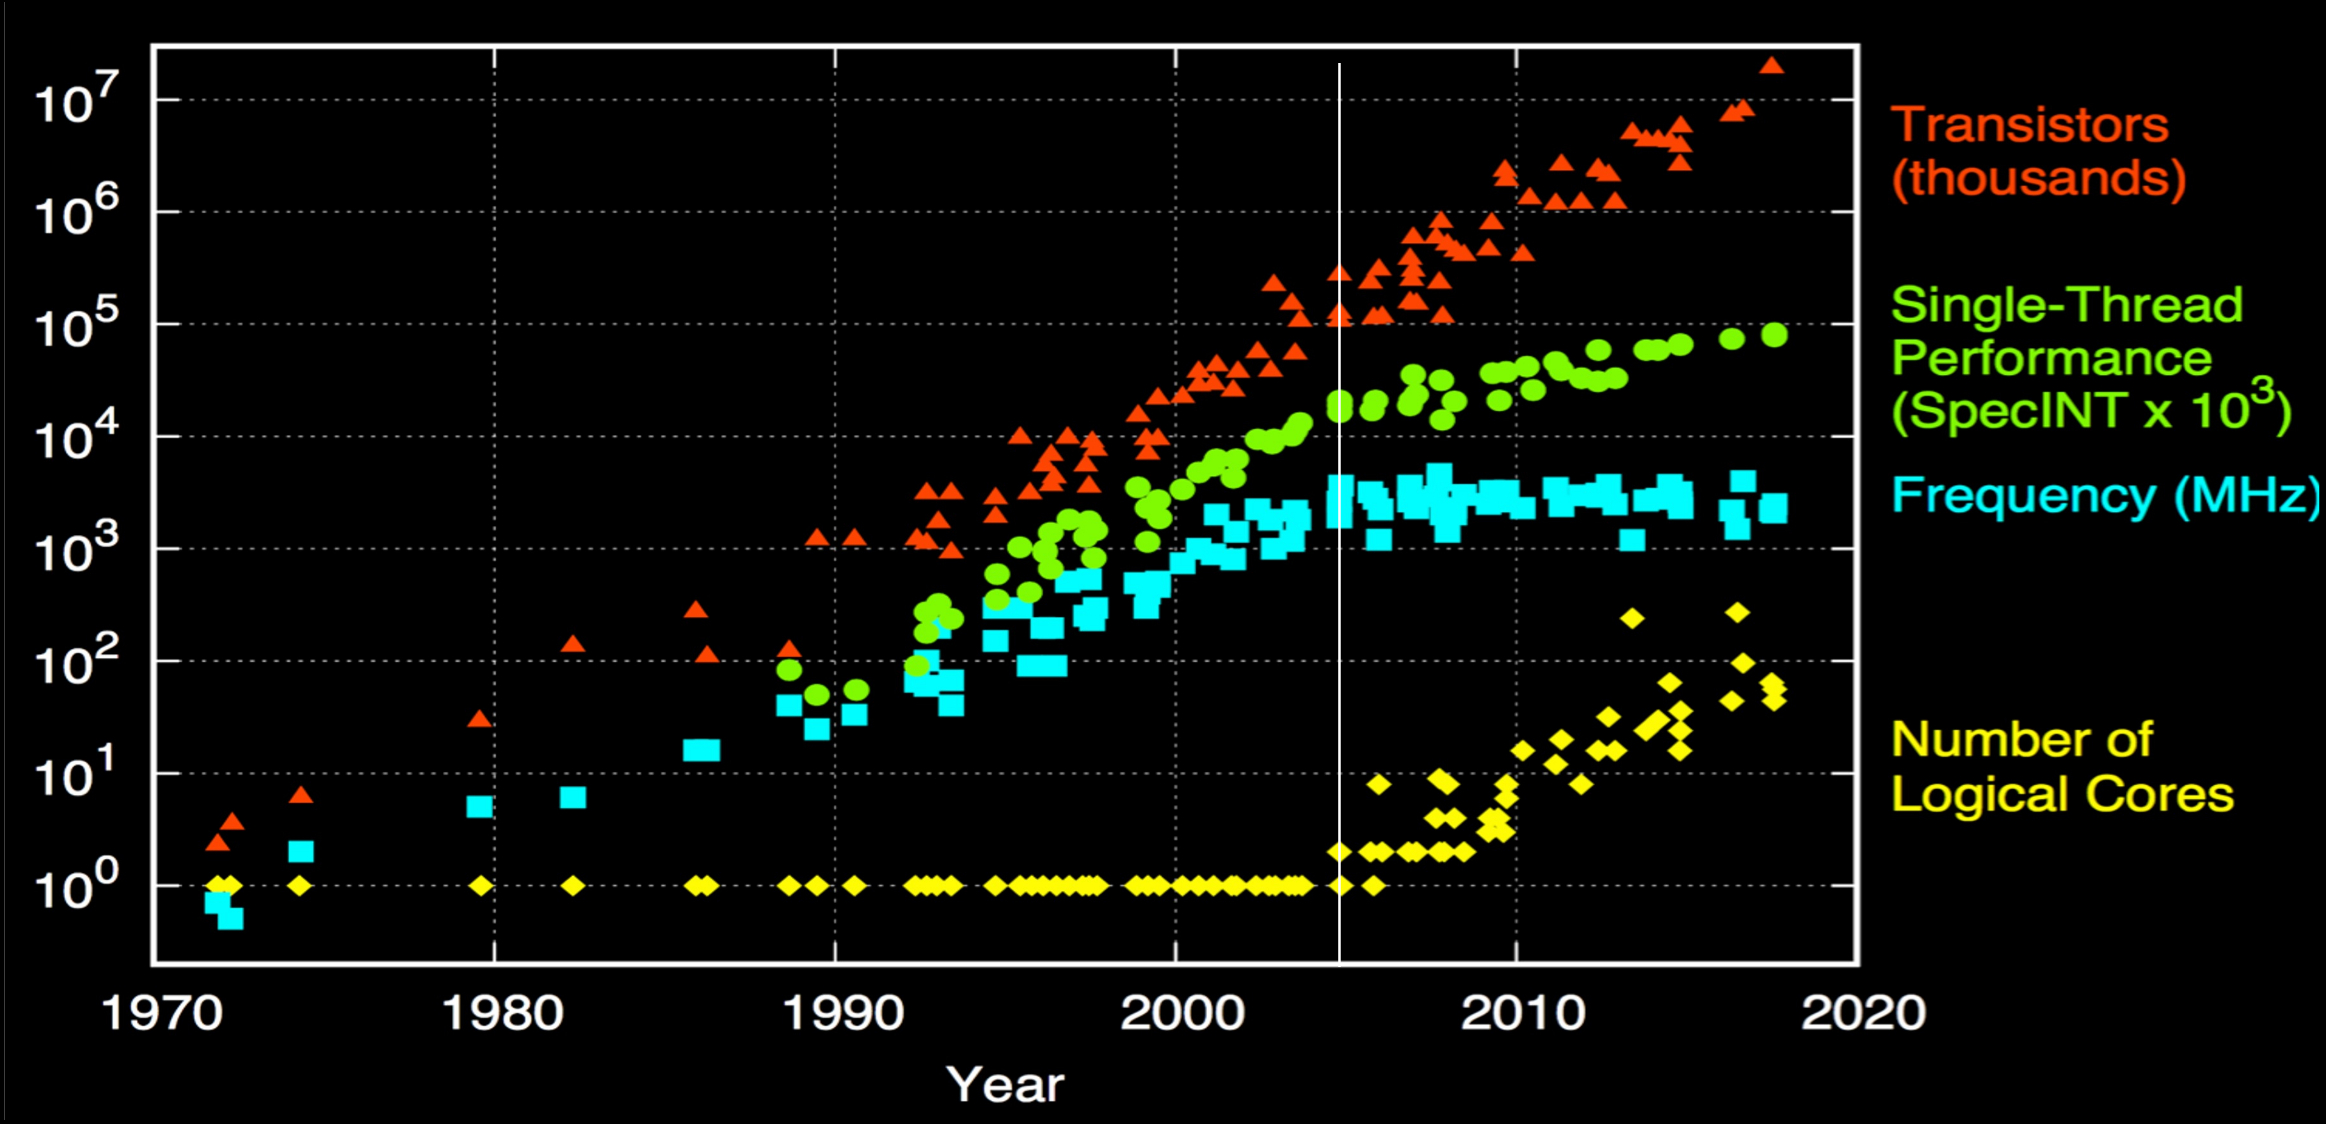

M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, C. Batten, and K. Rupp

Prediction of Atlas computing +$1B

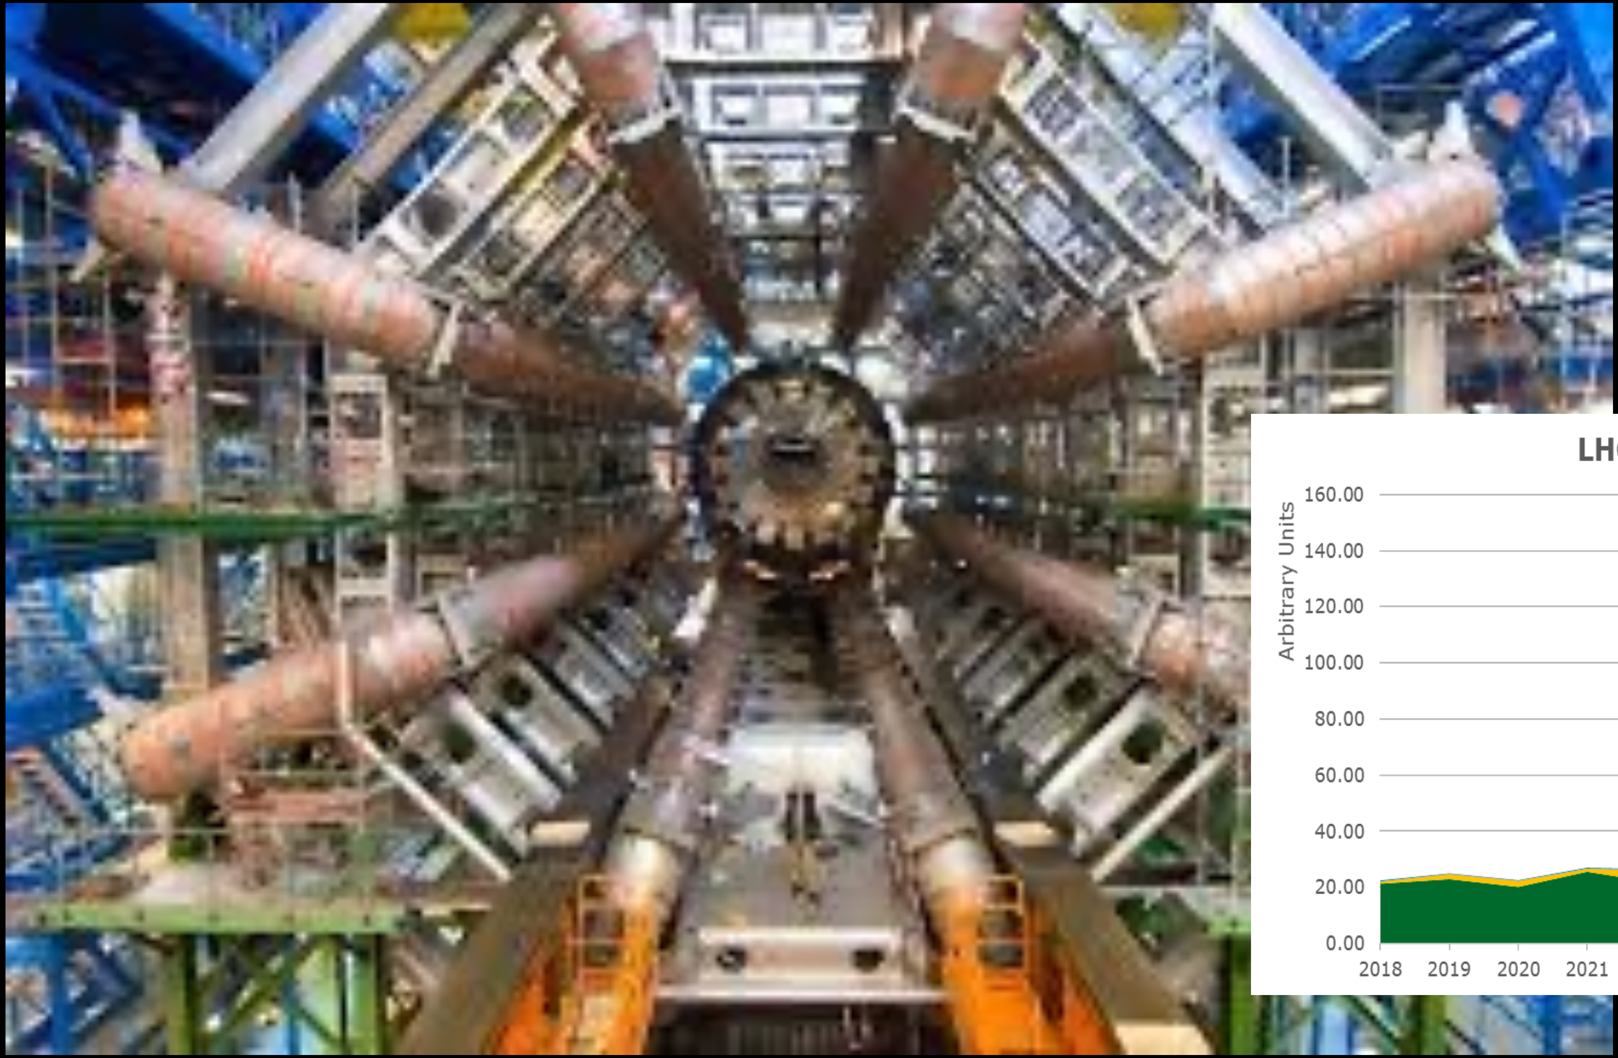

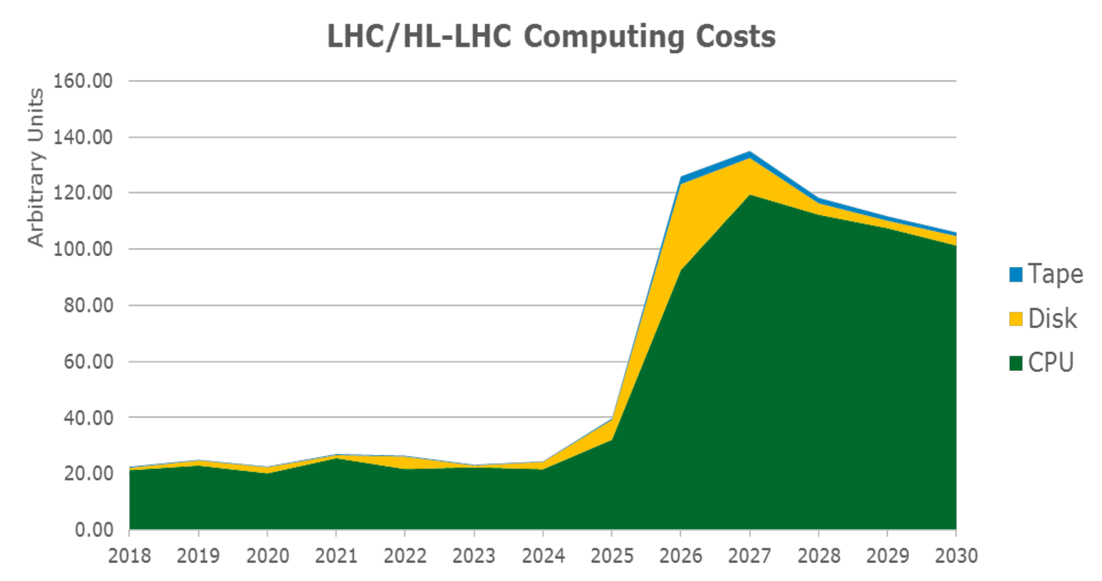

End Game for Moore's Law: Parallelization and Specialization

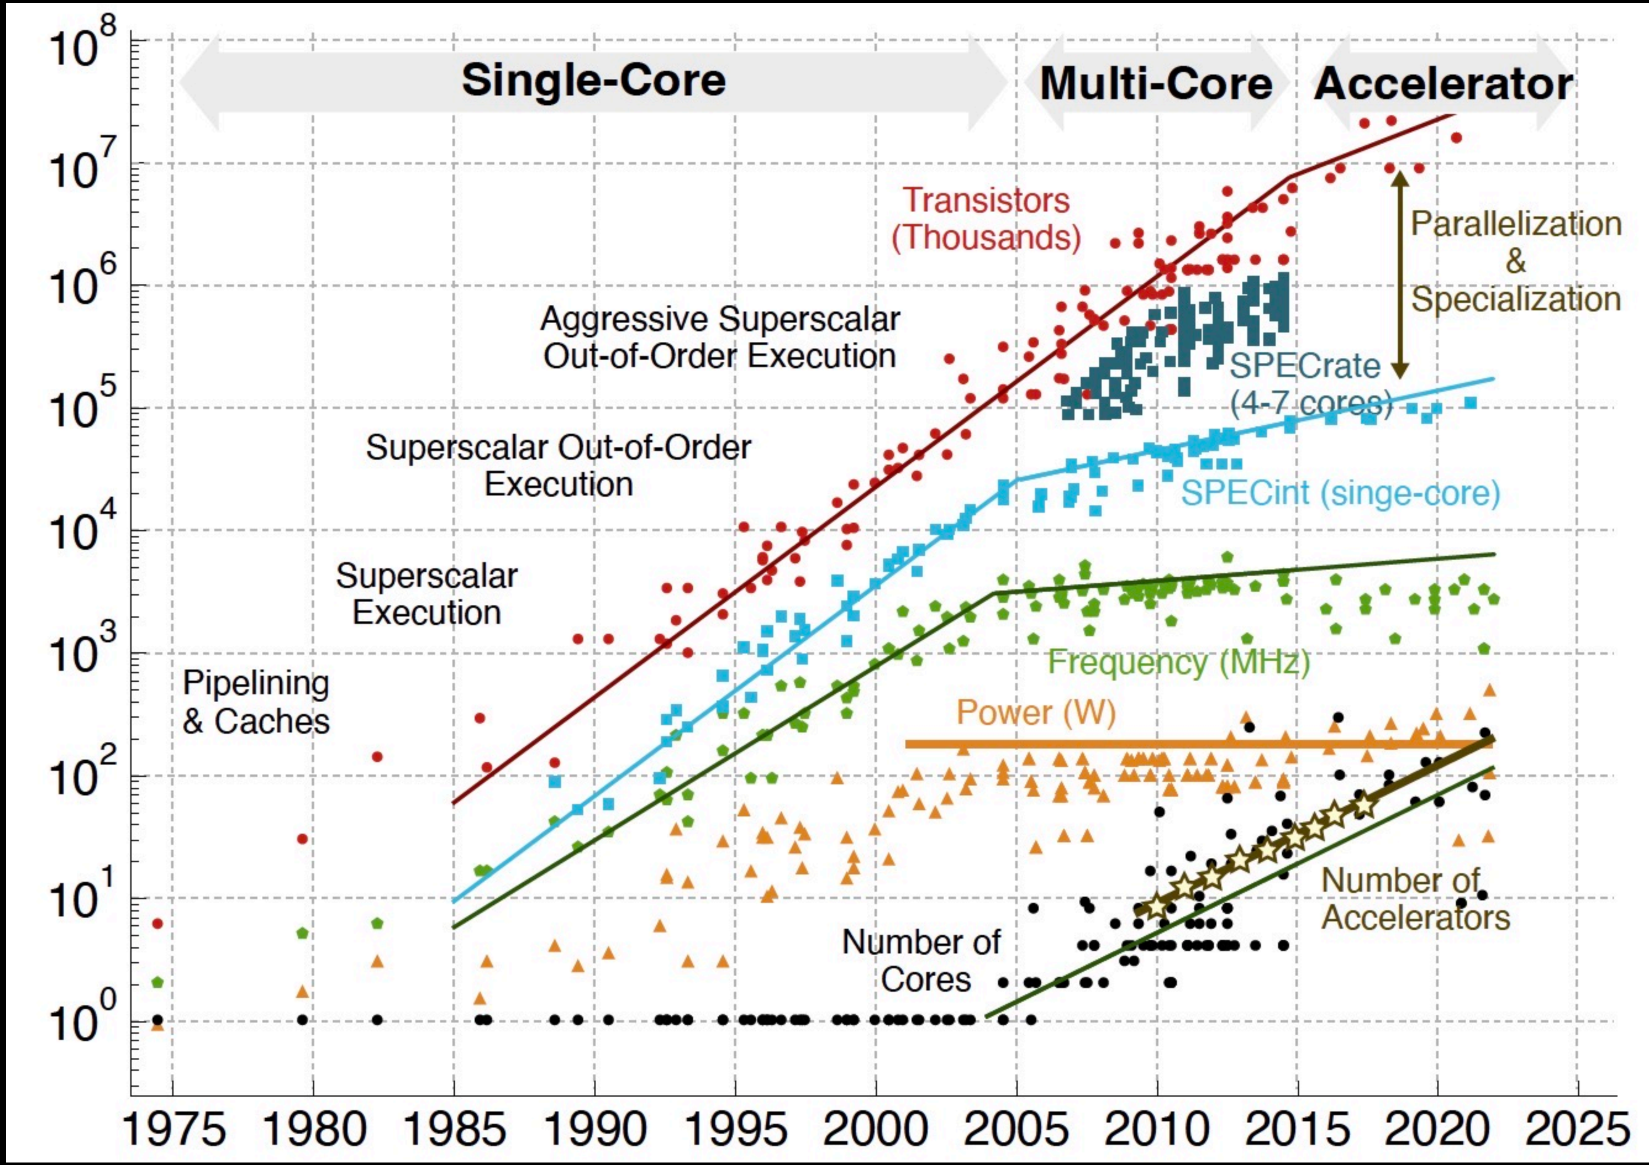

From Chris Batten, Cornell ENGRI 1210 citing

C. Batten, M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, K. Rupp & [Y. Shao, IEEE Micro'15] & [C. Leiserson, Science'20]

Exascale Architecture Plans (2008)

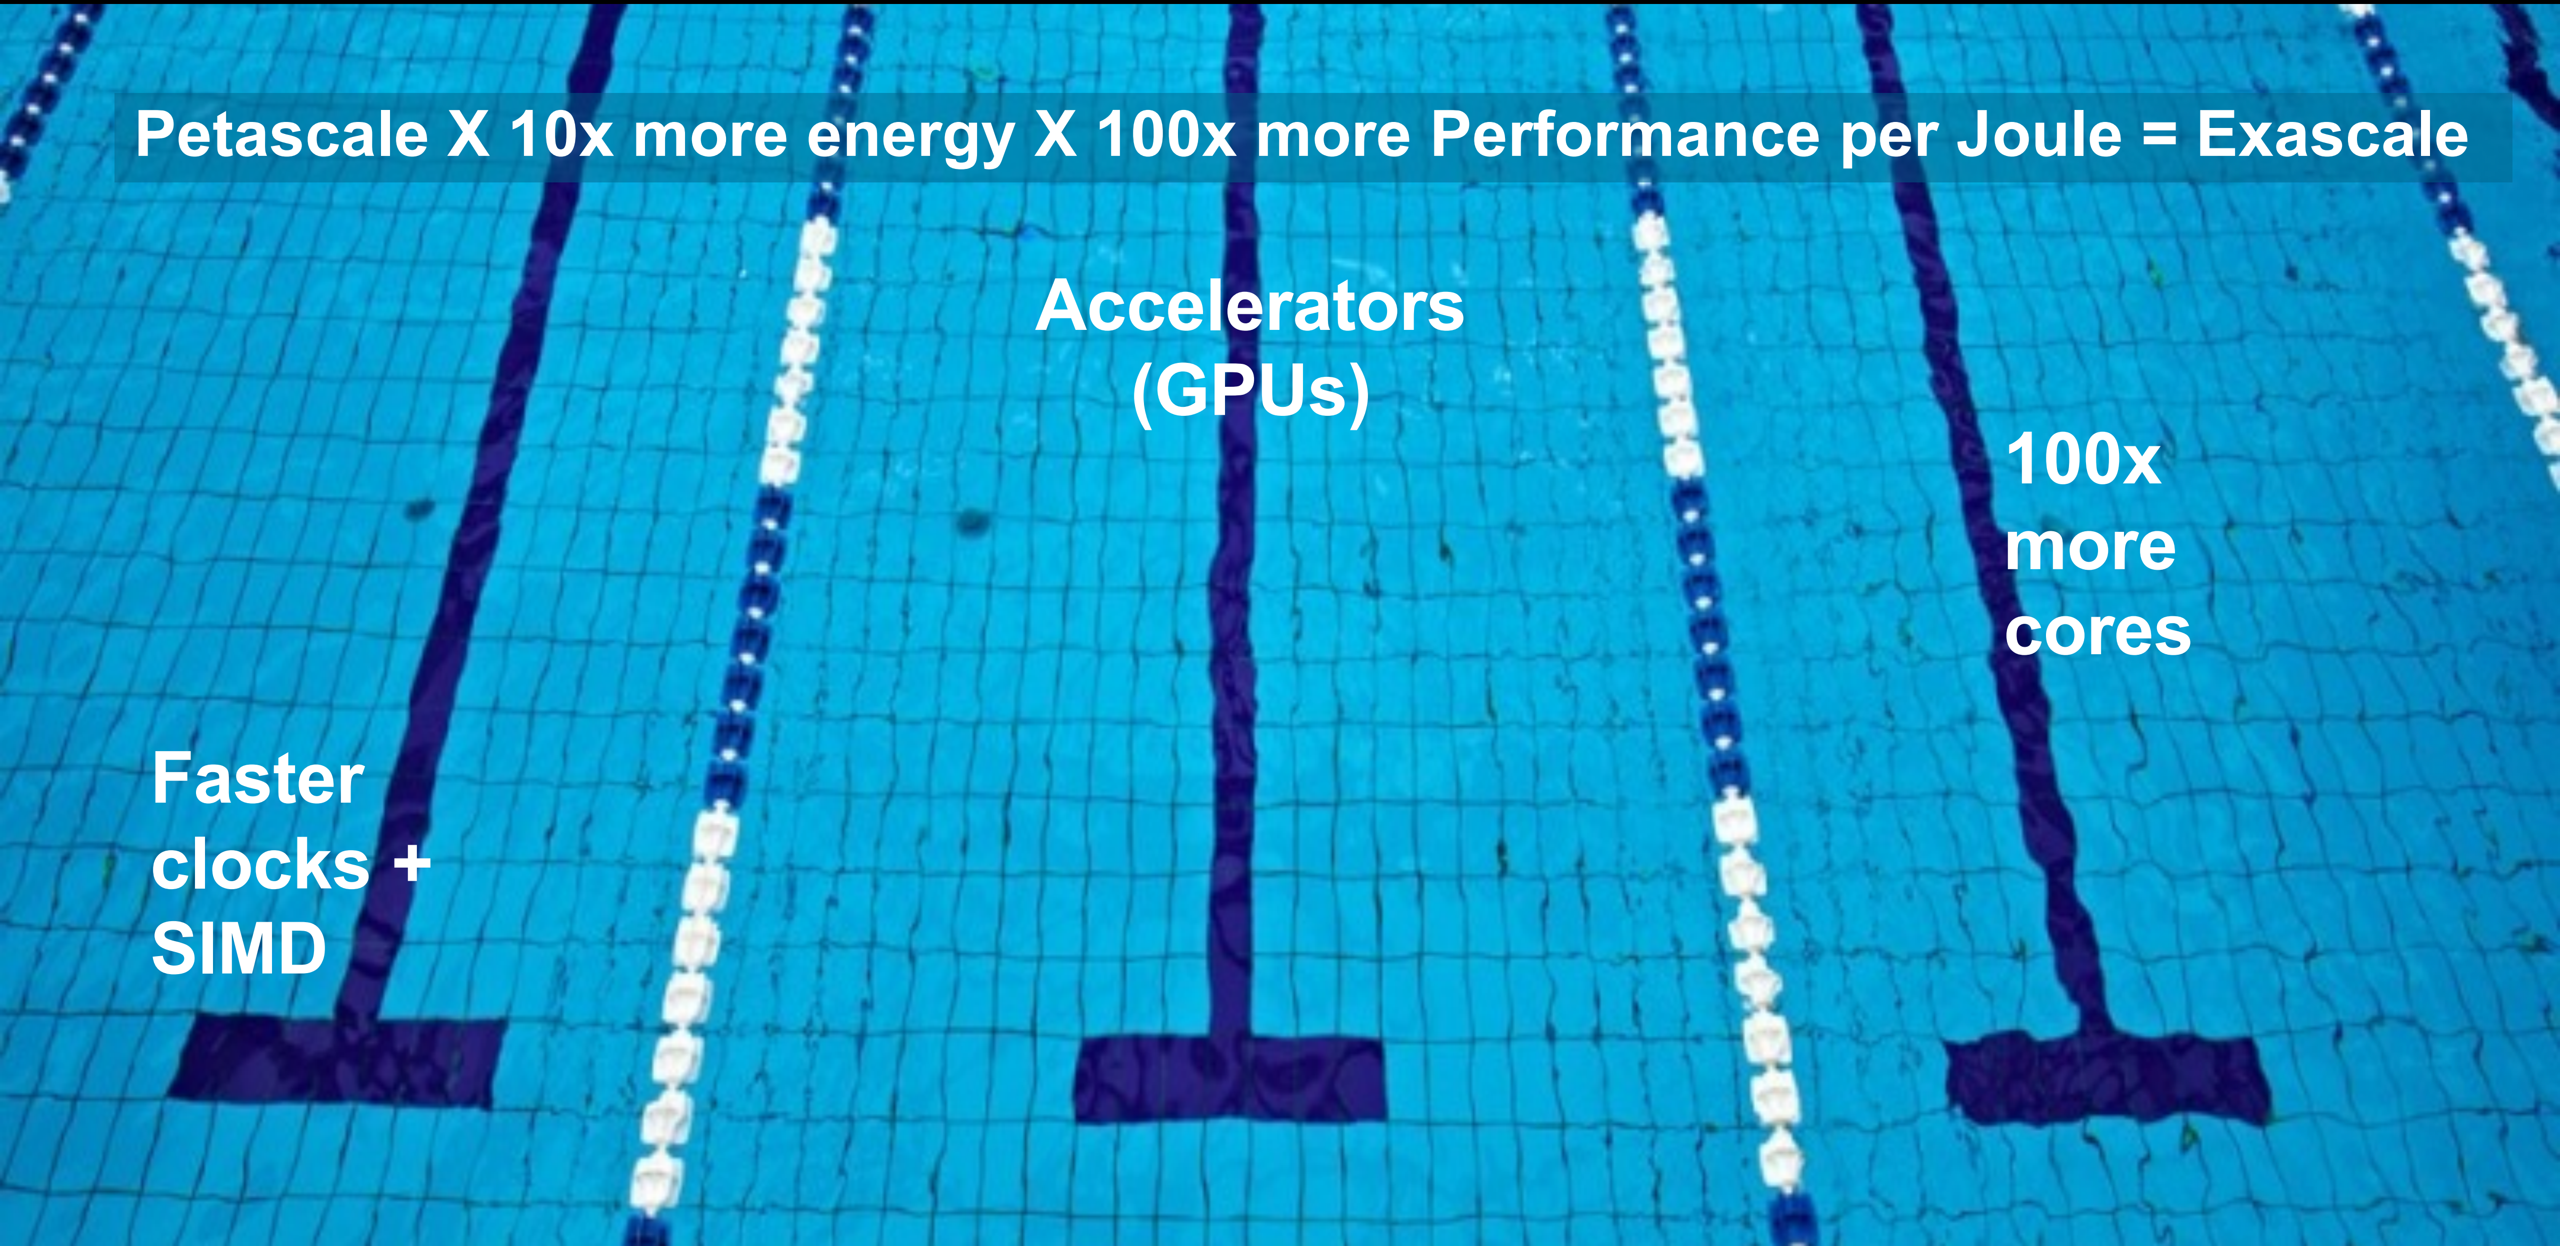

Exascale Era Architectures (US DOE Office of Science)

US DOE Office of Science Systems

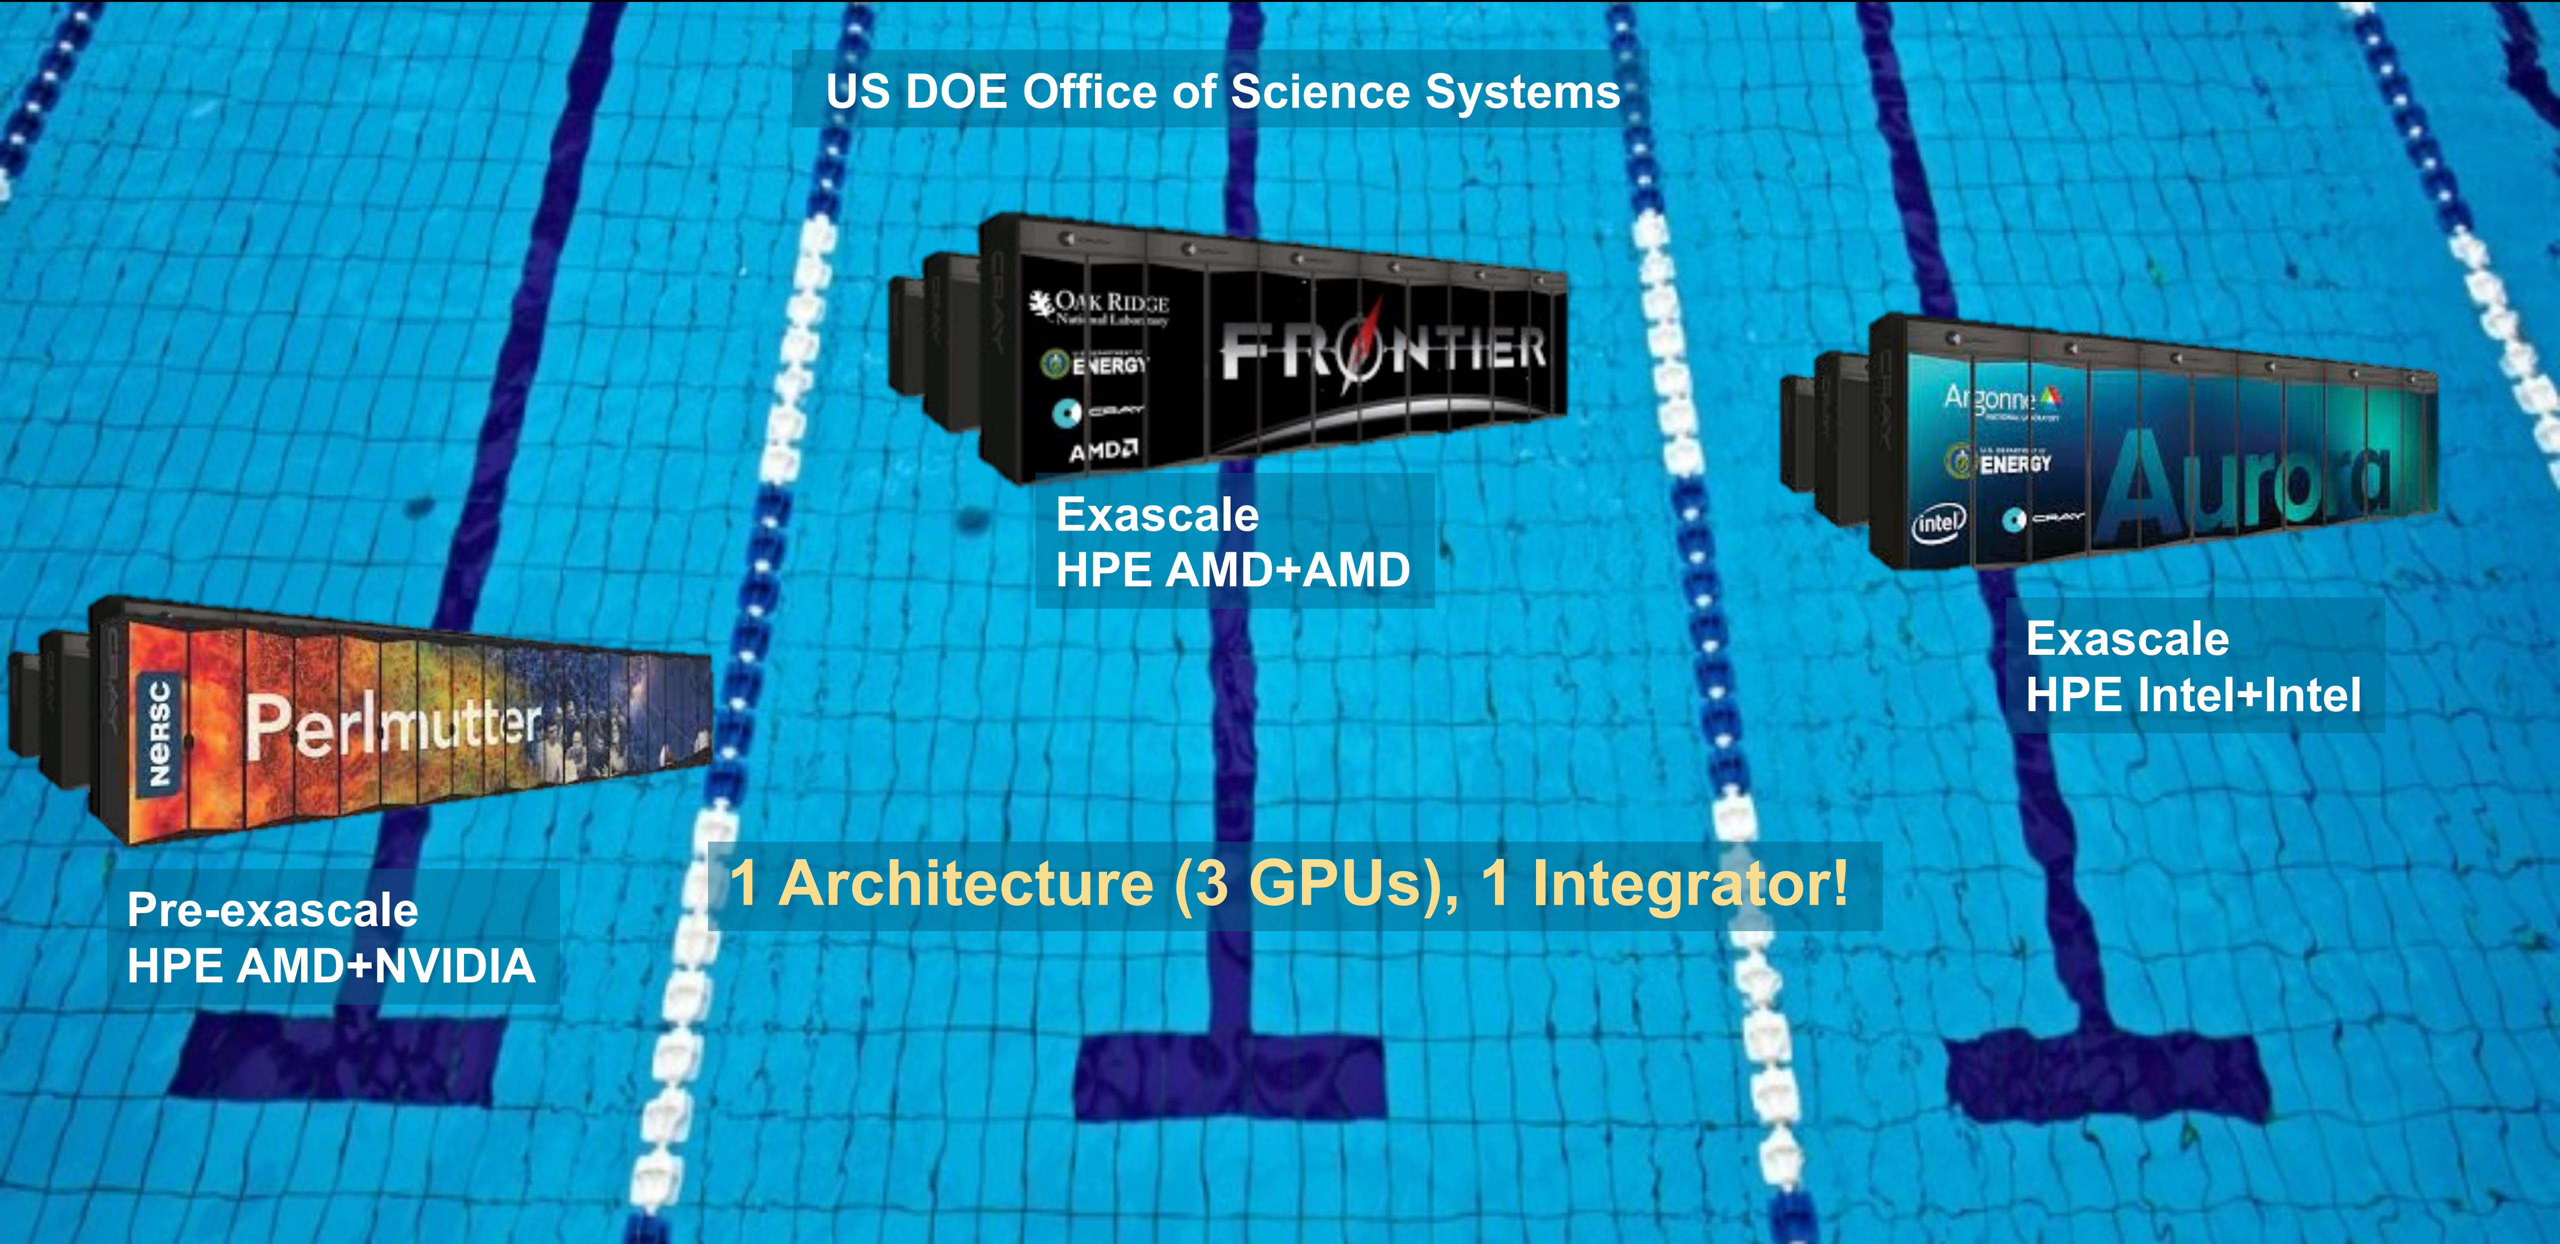

Pre-exascale
HPE AMD+NVIDIA

Exascale
HPE AMD+AMD

Exascale
HPE Intel+Intel

1 Architecture (3 GPUs), 1 Integrator!

Growth of Accelerators in HPC

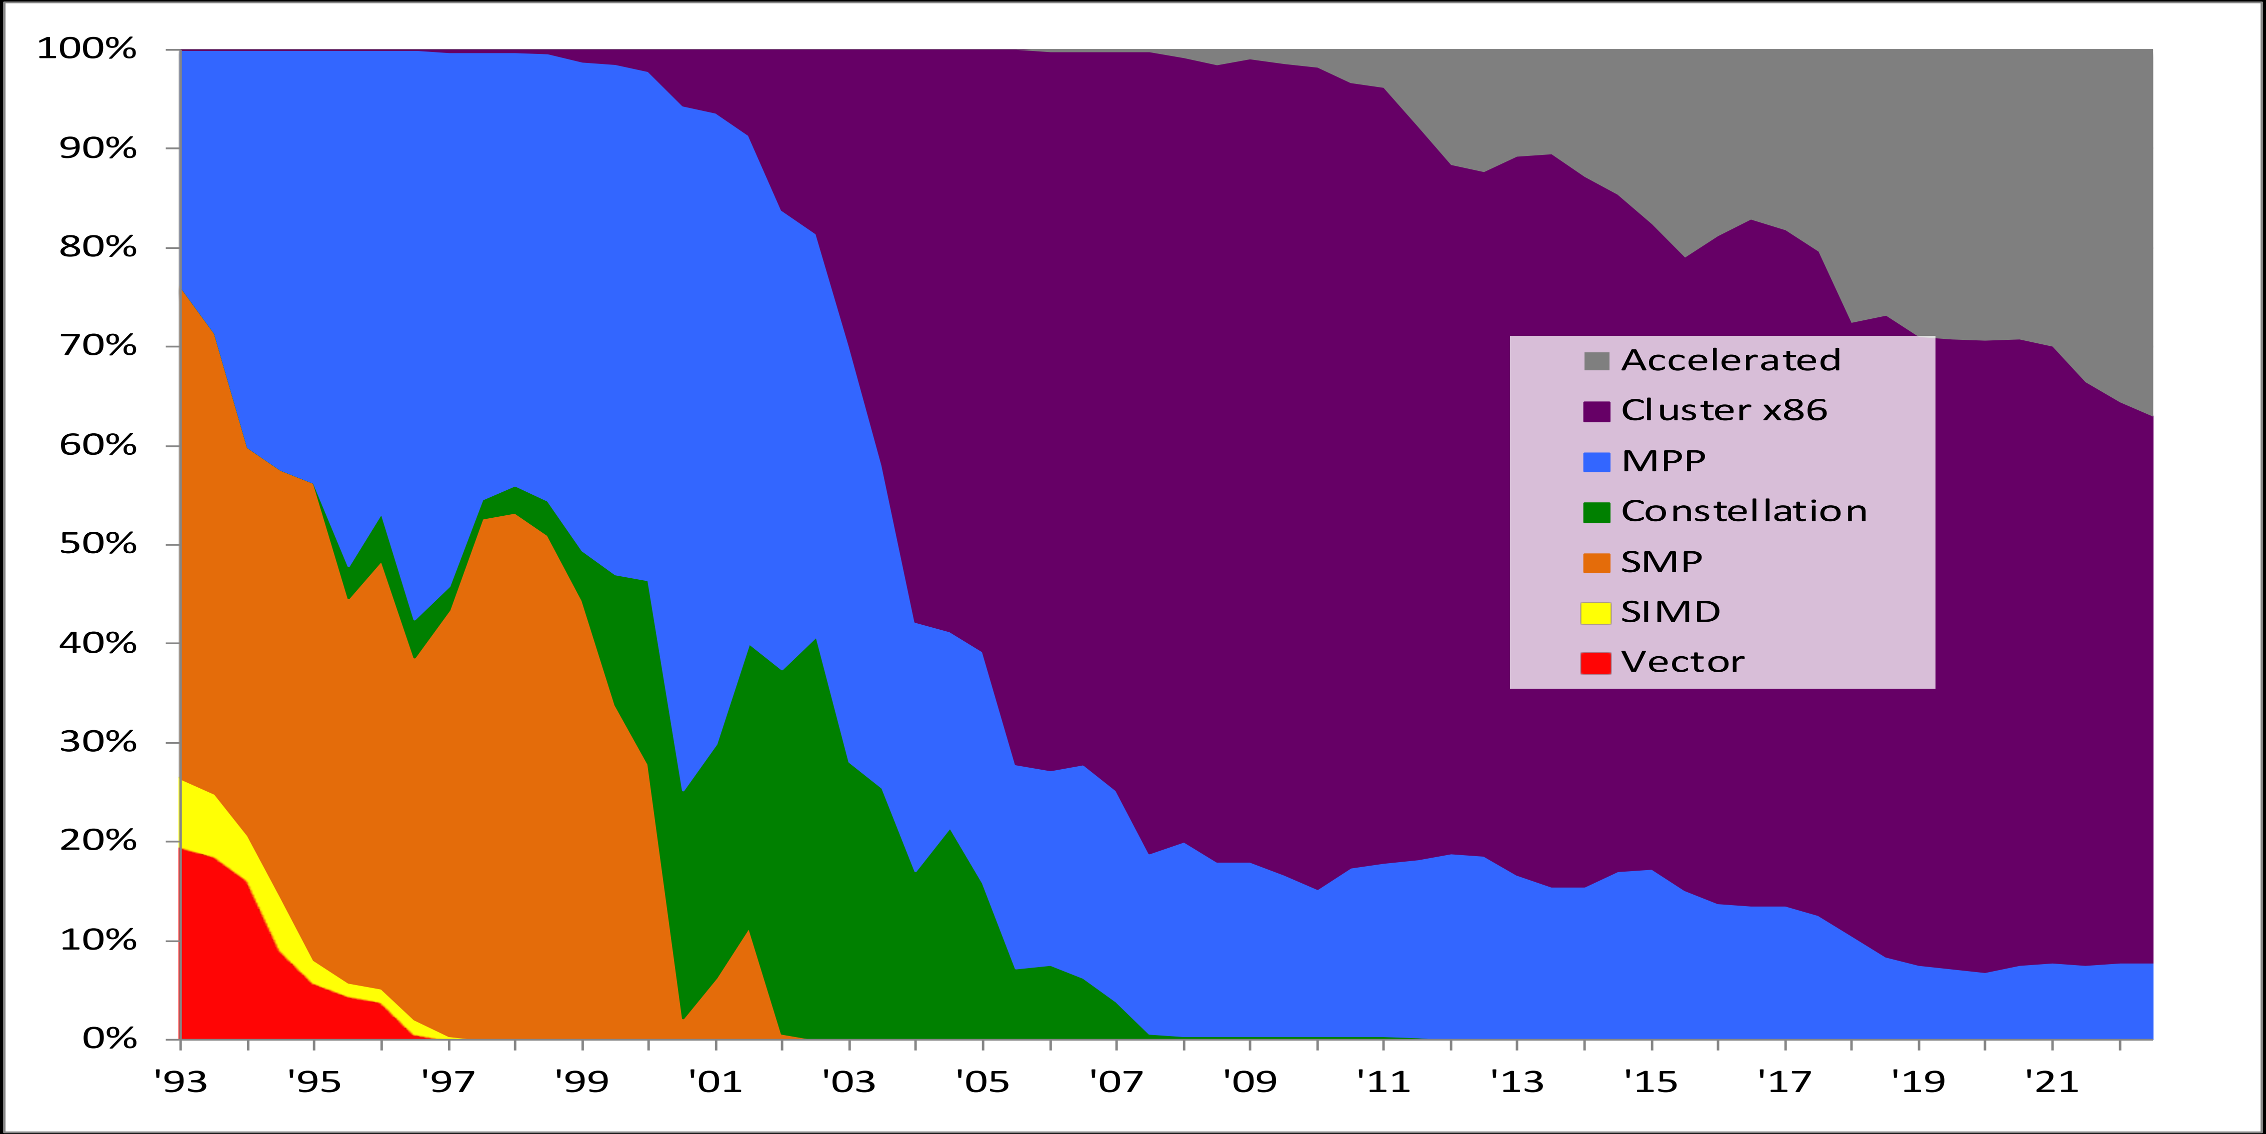

AI Chip Landscape

More on https://basicmi.github.io/AI-Chip/

Tech Giants/Systems

Google
Microsoft
aws
IBM
facebook
HUAWEI
Baidu 百度
Alibaba Group 阿里巴巴集团
FUJITSU
NOKIA
TOSHIBA
Hewlett Packard Enterprise
DELL

IC Vender/Fabless

intel
SAMSUNG
NVIDIA
QUALCOMM
AMD
NXP
ST
XILINX
MEDIATEK
BROADCOM
MARVELL
Rockchip 瑞芯微电子

IP/Design Service

arm
SYNOPSYS
Imagination
cadence
CEVA
VeriSilicon
SiFive
ARTERIS IP
alchip
GUC
FARADAY
eSilicon

Startup in China

Cambricon 寒武纪科技
地平线 Horizon Robotics
BITMAIN
intellifusion 云天励飞
ChipIntelli
Think Force
Canaan
云知声 Unisound
Rokid
AISPEECH 思必驰 专注人性化的智能语音
肇观电子 NextVPU
Enflame
亿智科技
清微智能 TSING MICRO

Startup Worldwide

cerebras
WAVE COMPUTING
Graphcore
habana
thinci
SambaNova SYSTEMS
KALRAY
LIGHTELLIGENCE
HAILO Empowering Intelligence
Esperanto TECHNOLOGIES
Tenstorrent
MYTHIC
Preferred Networks
brainchip
PEZY Computing
GREENWAVES TECHNOLOGIES
AIMOTIVE
KONIKU
Tachyum
flexlogix
SYNTIANT
gyrfalcon technology
NOVUMIND

扫码访问AI芯片文章

Compiler

GLOW
tvm
NVIDIA TensorRT
ONNC
nGraph Compiler stack (Beta)
plaidML

Benchmarks

MLPerf
AI - Benchmark
AI Matrix.

Everyone is Making AI Chips

NVIDIA
AMD
Intel
IBM

Traditional chip makers

"Software" companies

Facebook + Intel
Amazon (Echo, Oculus)
Google (TPU, Pixel)
Apple (SoCs)
Microsoft ("AI chip")

Not everyone is selling those chips!

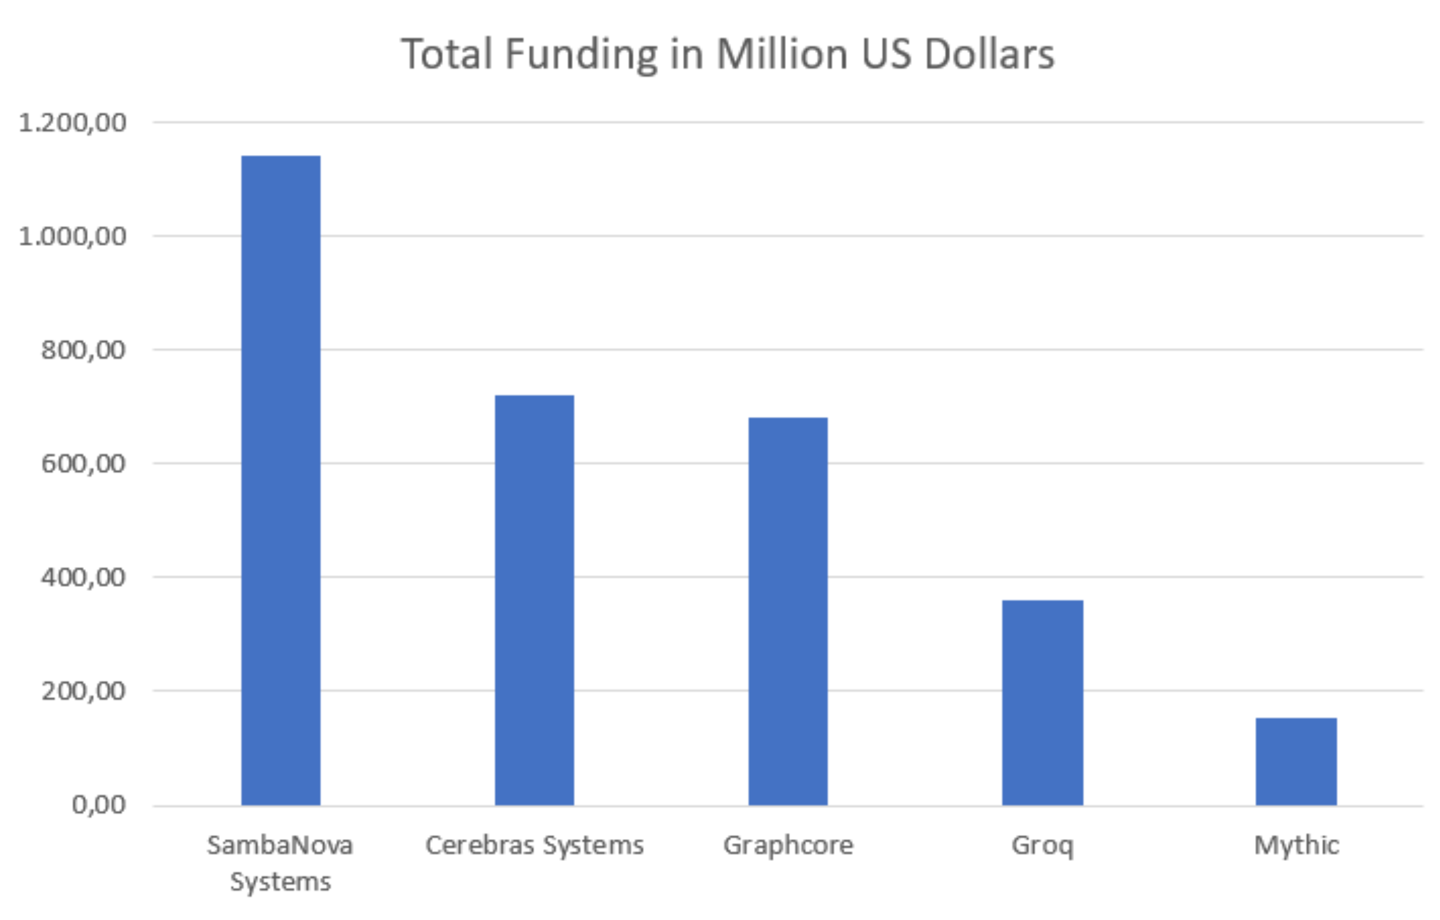

Graphcore, Nervana Cerebras, Wave Computing, Horizon Robotics, Cambricon, DeePhi, Esperanto, SambaNova, Eyeriss, Tenstorrent, Mythic, ThinkForce, Groq, Lightmatter

Specialization for the masses?

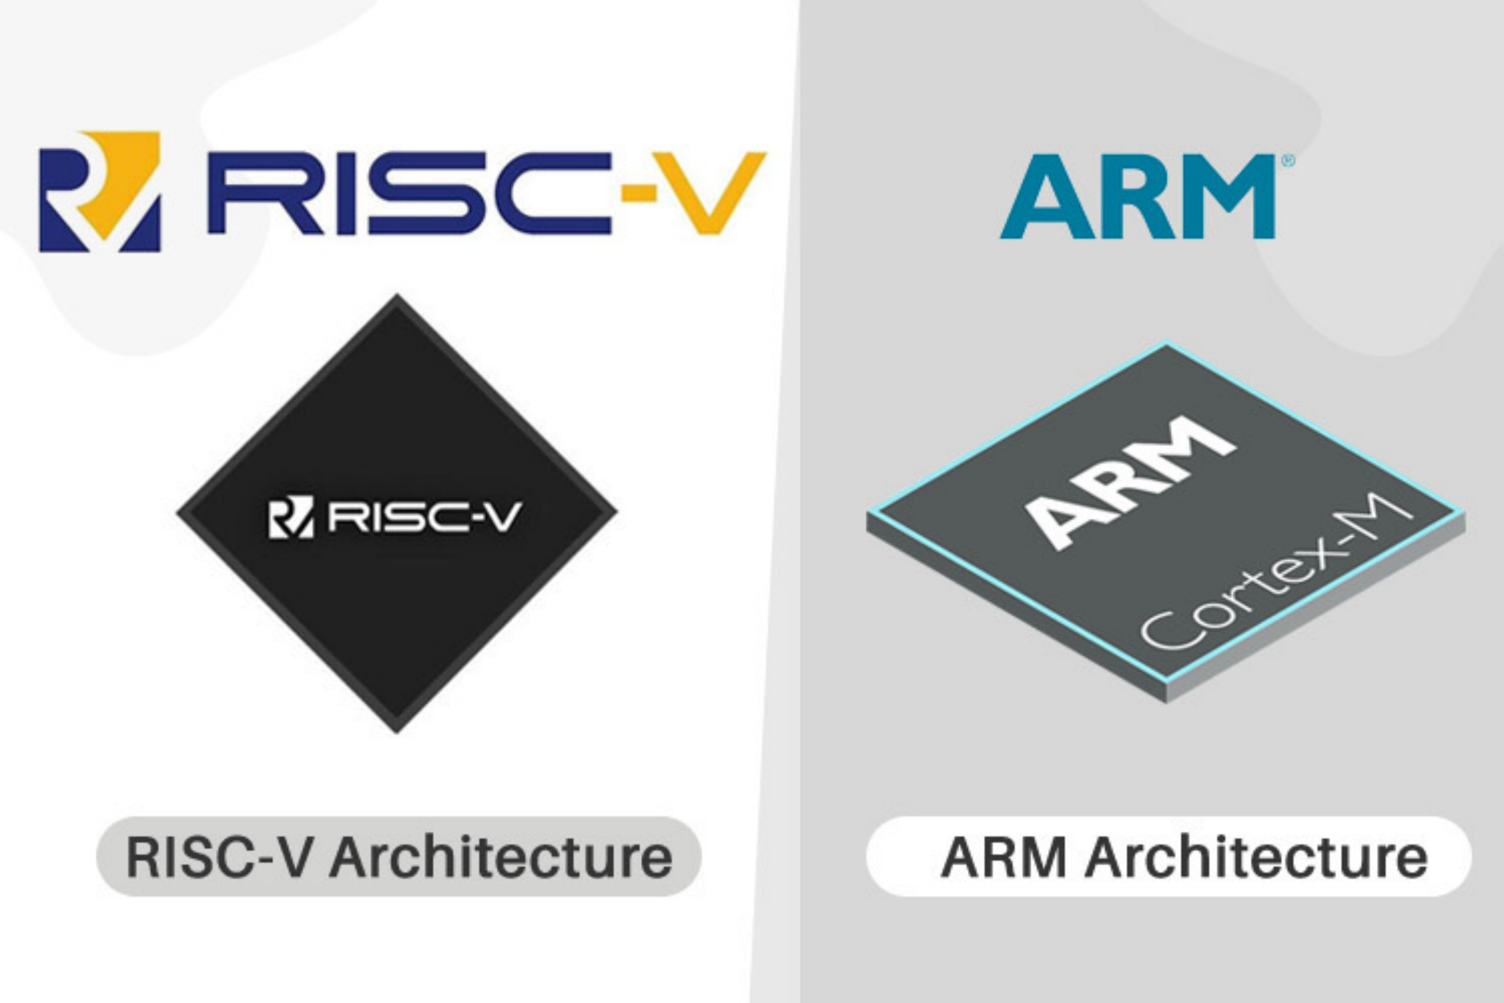

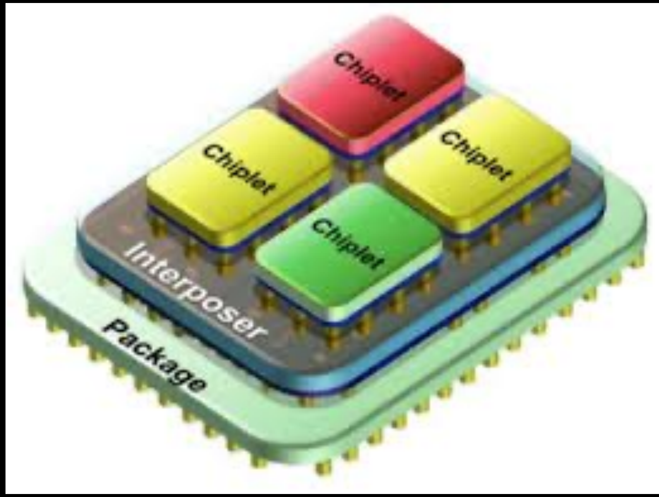

Specialization: Is deep learning the only application?

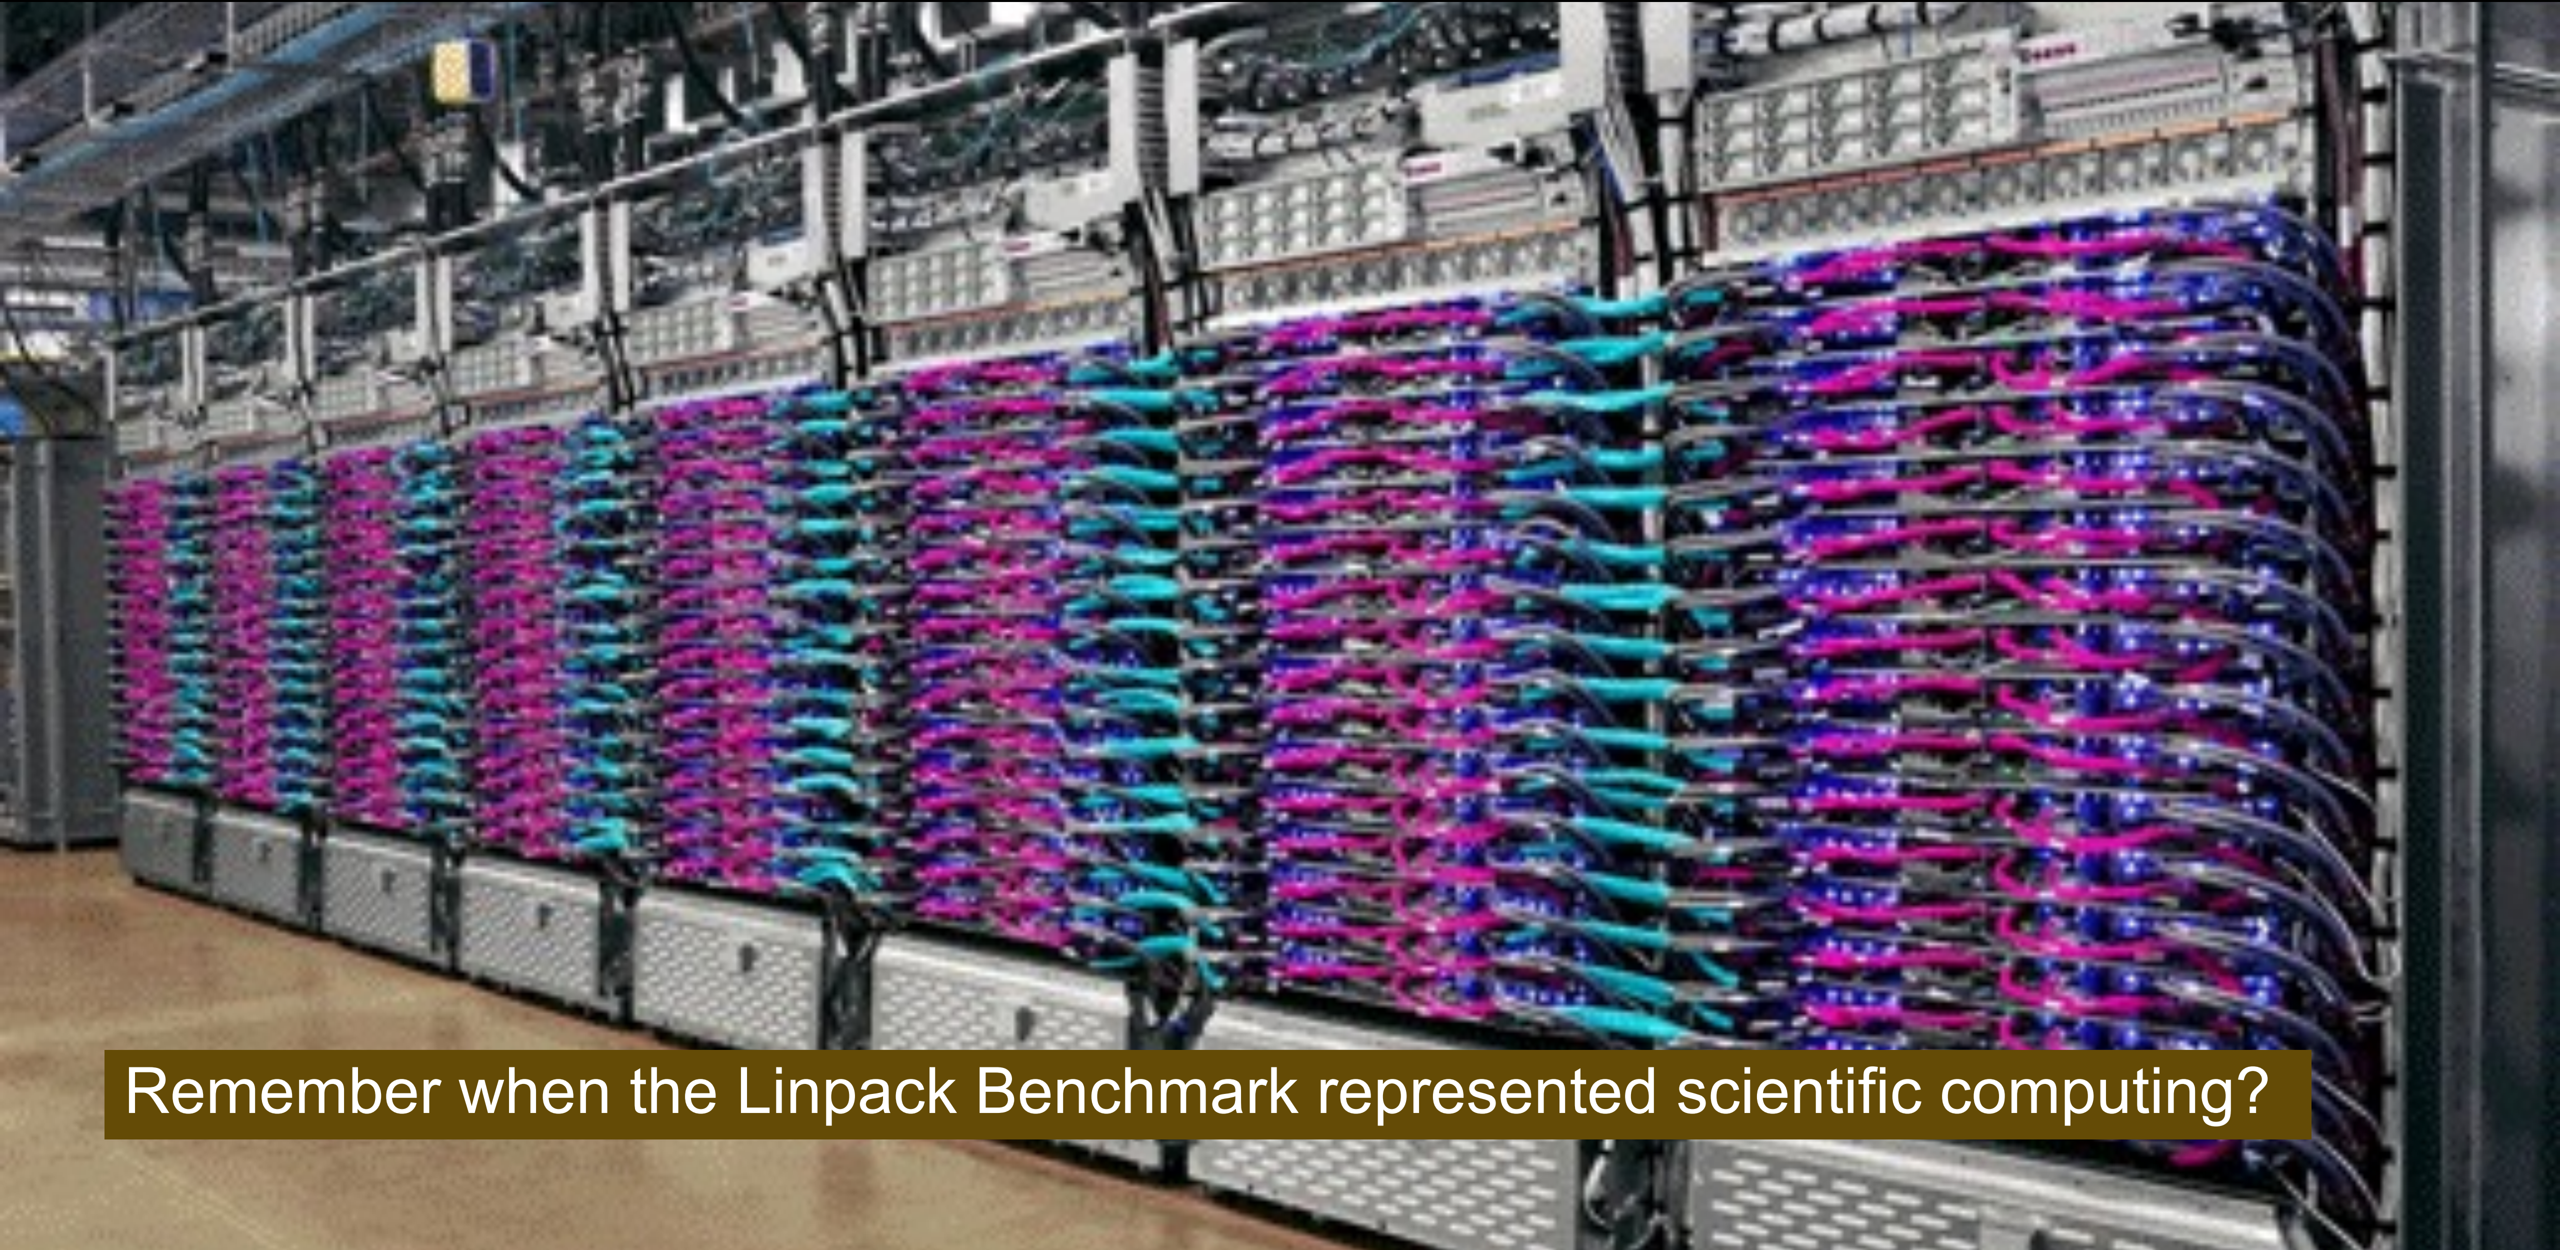

Remember when the Linpack Benchmark represented scientific computing?

Cloud "Hyperscalars" dominate computing

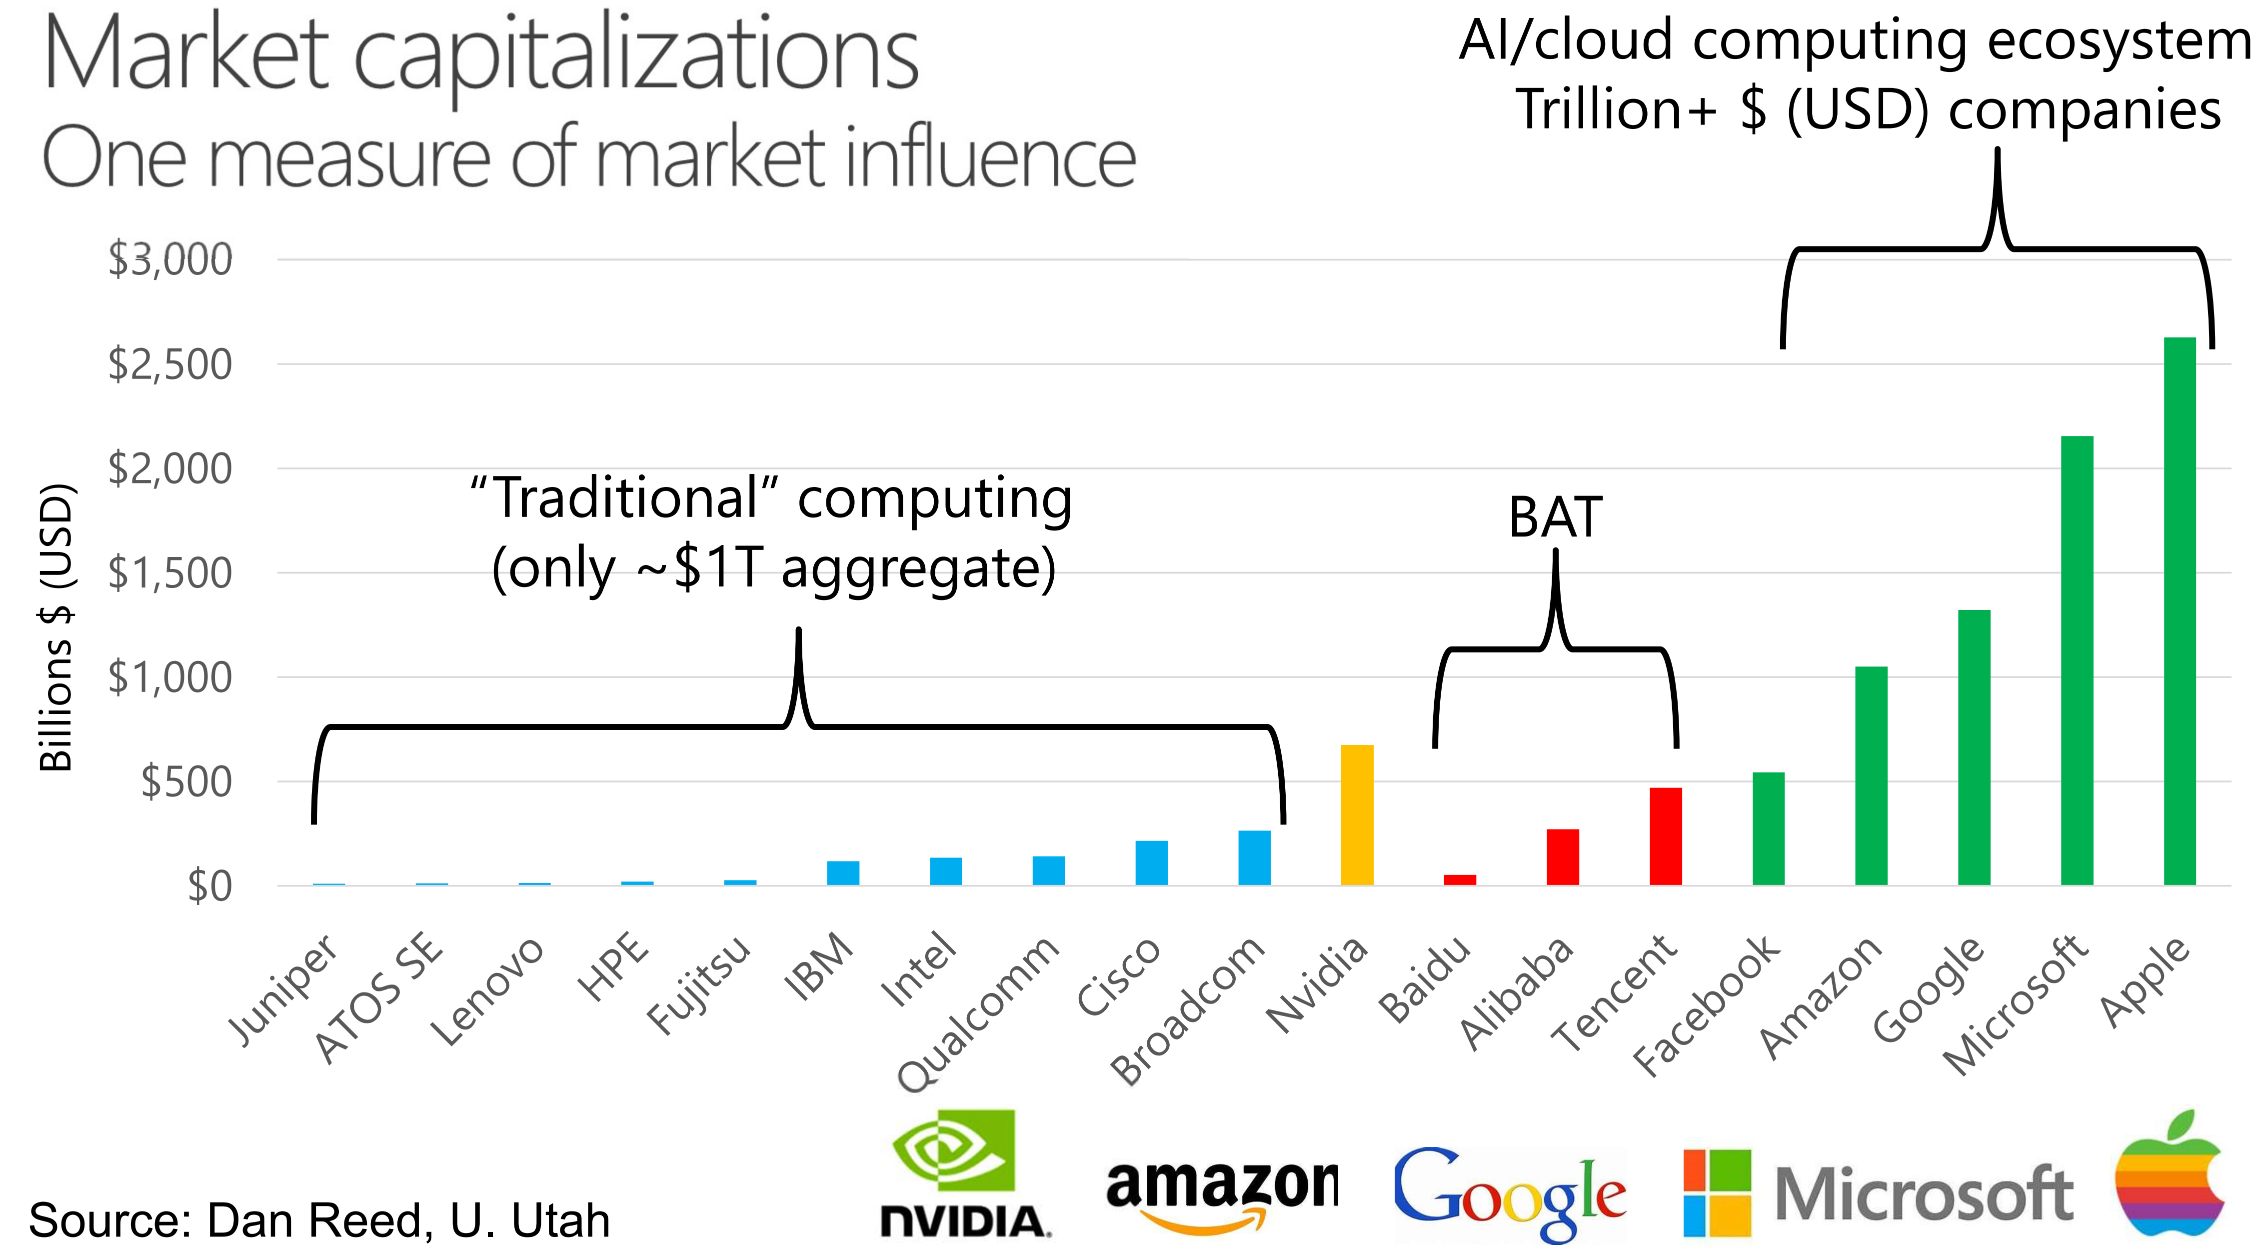

Source: Dan Reed, U. Utah

Technology and Marketplace: Radically Different!

What's a post-Exascale strategy for the science community?

Beat them

- Design processors for science

  More Co-Design and

  don't forget the math and software

Join them

- Leverage AI Hardware

  for AI in Science

  andSimulation ?

Post-Exascale Computing

Computing demand

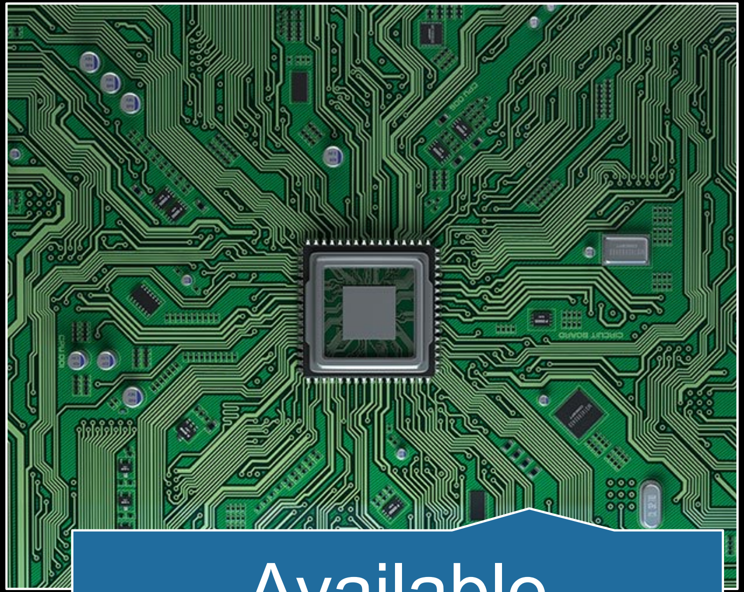

Available technology

Disruptions

AI for Science

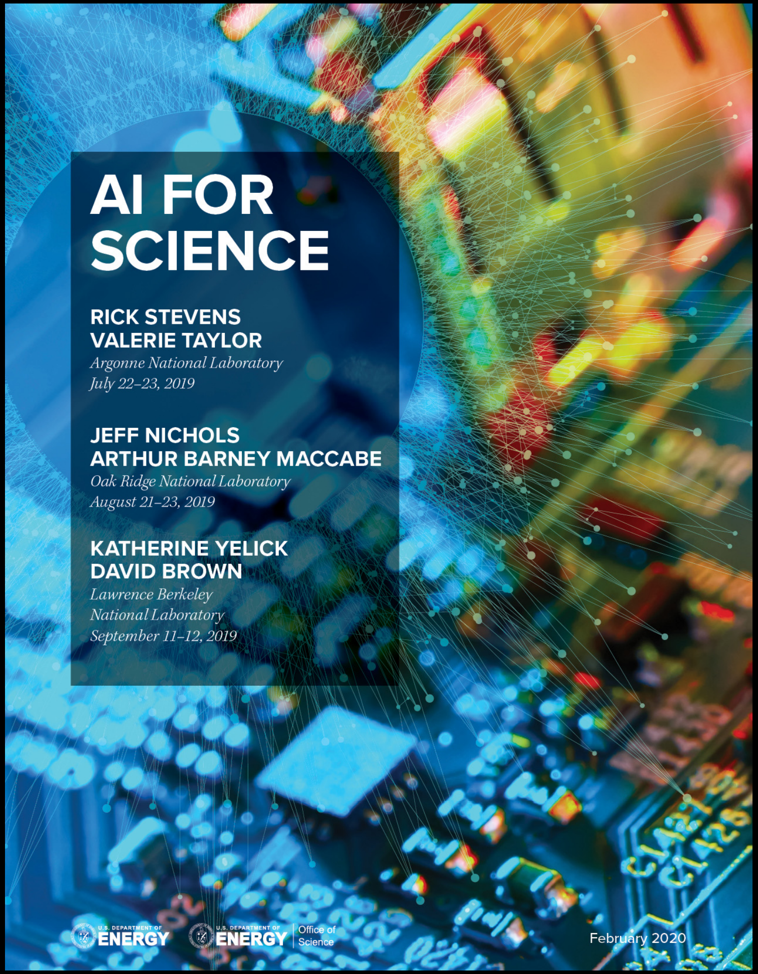

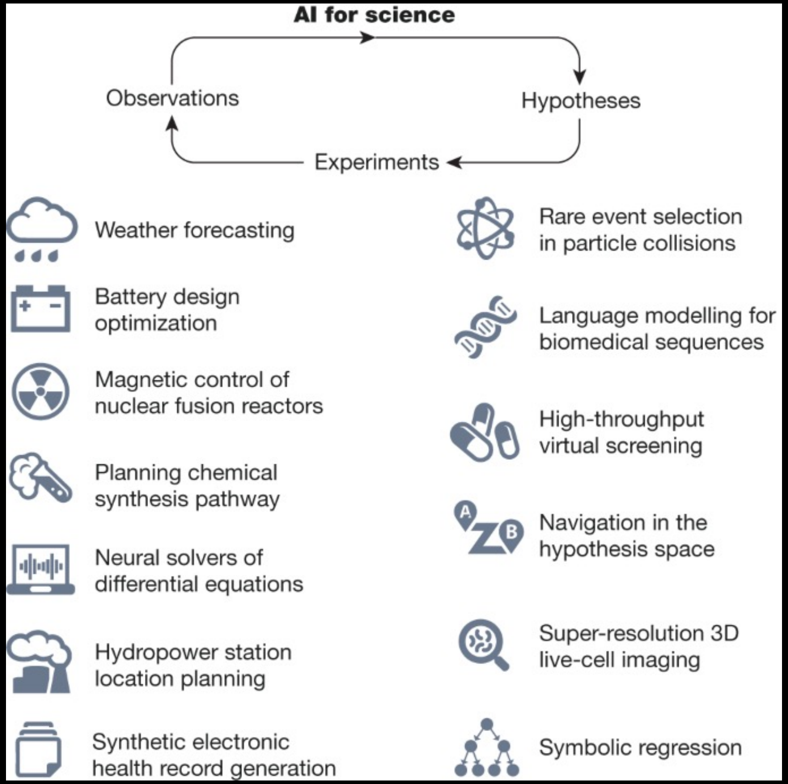

Scientific discovery in the age of artificial intelligence, 2023

Analyze Images to Find Cats

Classification

Localization

Detection

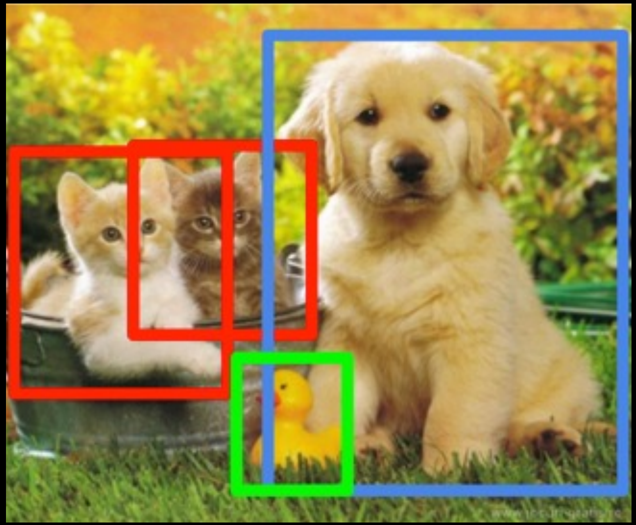

Segmentation

Source: Prabhat

Analyze Simulations to Find Hurricanes

Classification

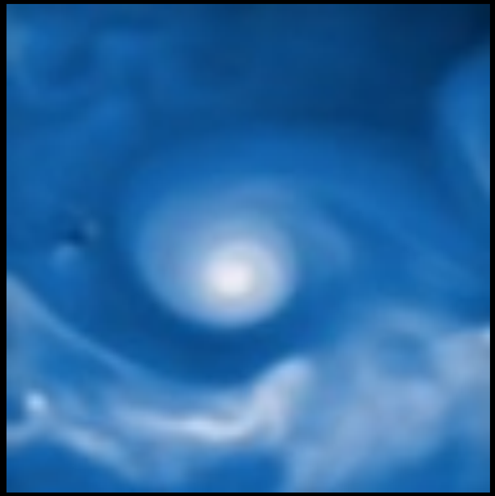

Localization

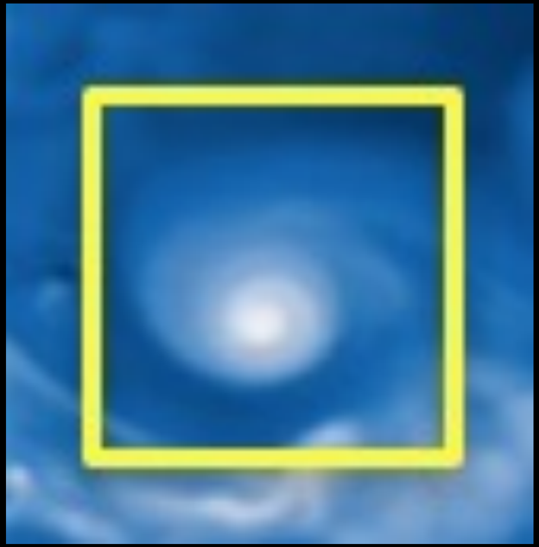

Detection

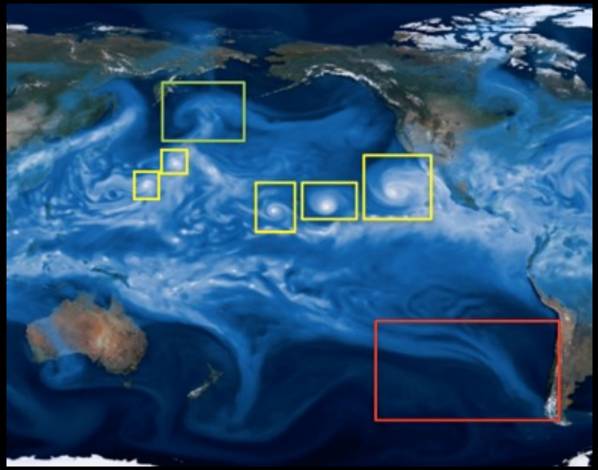

Segmentation

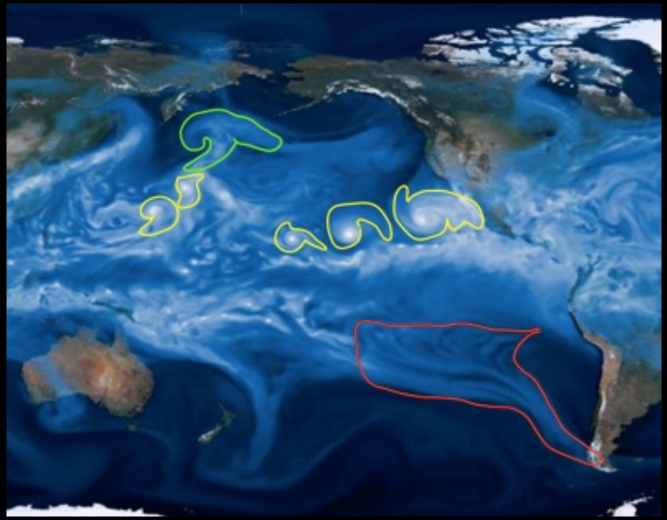

Extending image-based methods to complex, 3D, scientific data sets is non-trivial!

Source: Prabhat

Deep Learning: like adding 4,000 extra tons of detectors!

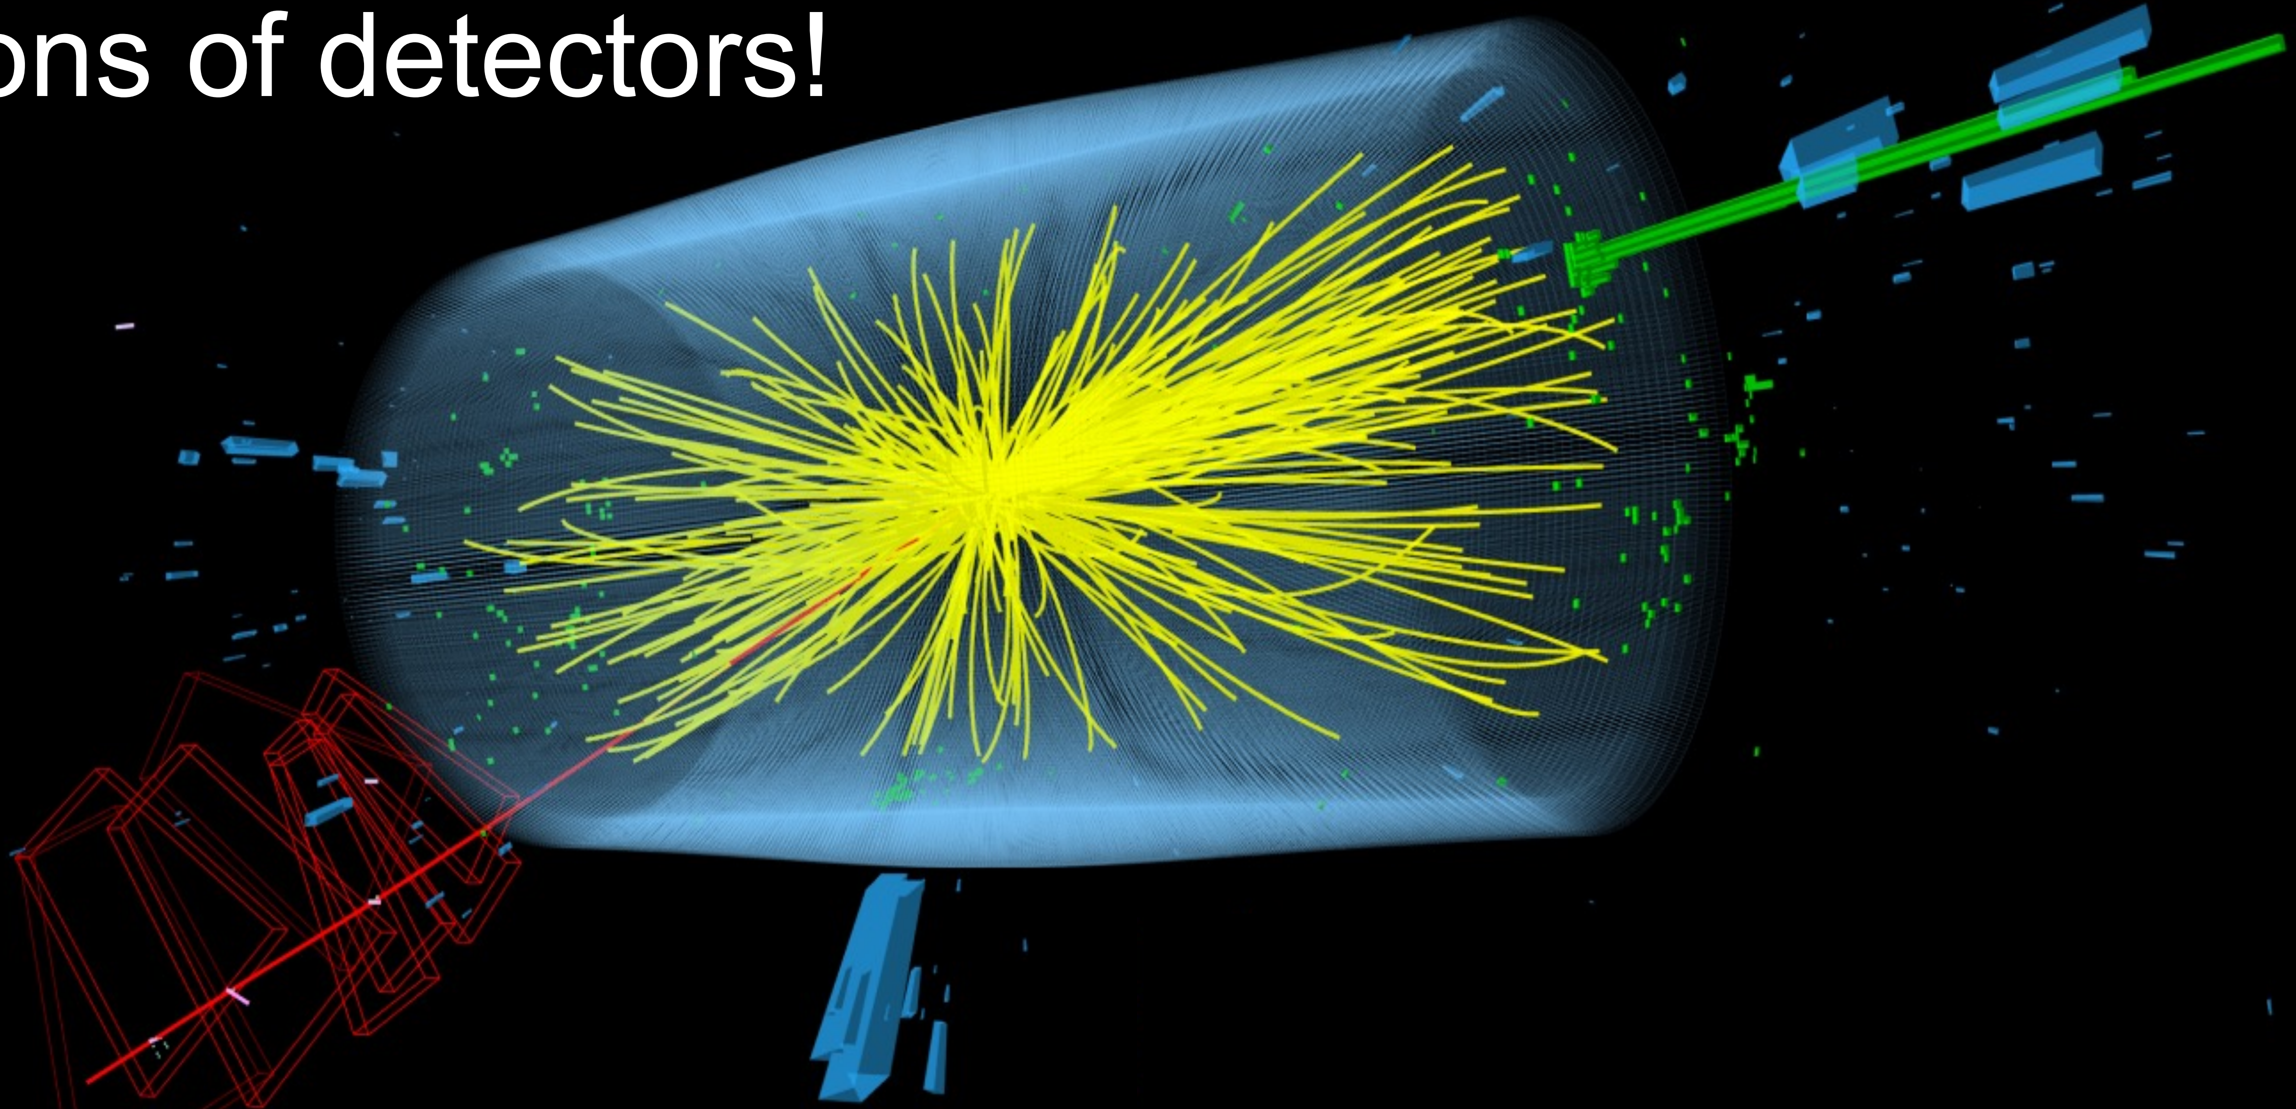

Based on 8/12/2016 slide by Joe Lykken at Fermilab

CNNs for Materials with Physical Laws

Physics-aware learning

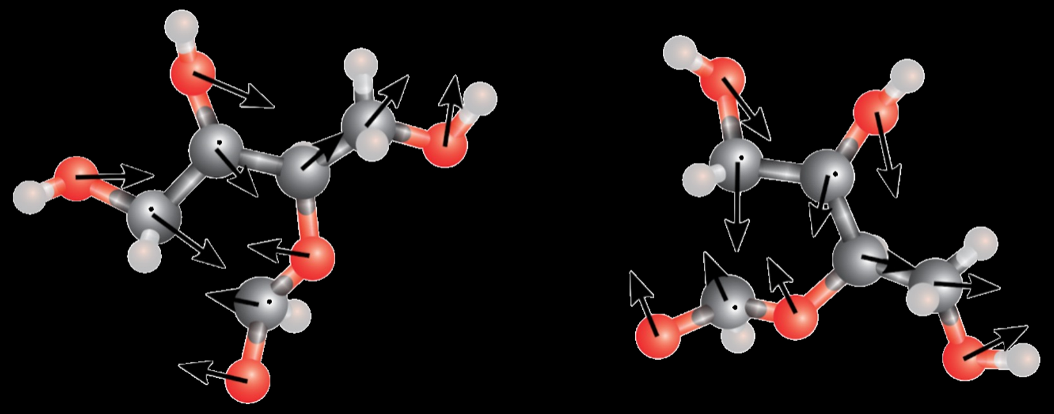

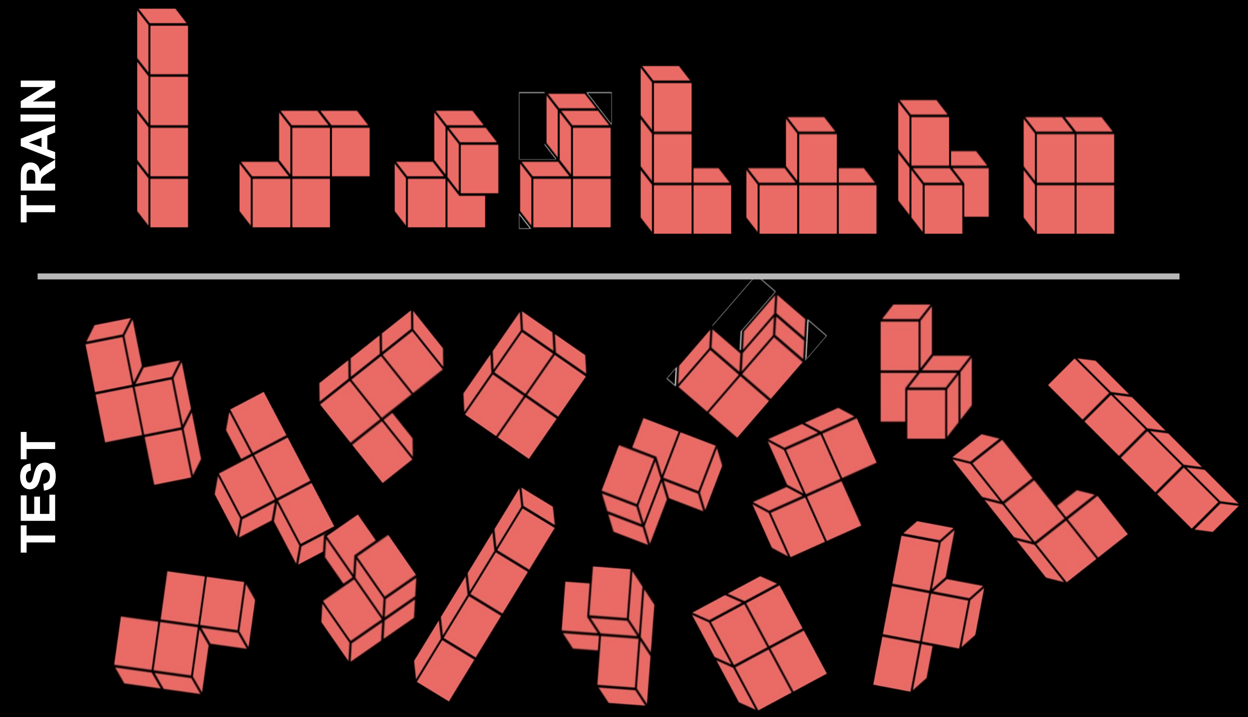

A network with 3D translation- and 3D rotation-equivariance

Slides from Tess Smidt and Risi Condor; E.g., 2018 paper by Thomas, Smidt, Kearnes, Yang, Li, Kohlhoff, Riley

Using NLP on scientific publications

Analyze 3.3 million abstracts from materials science papers

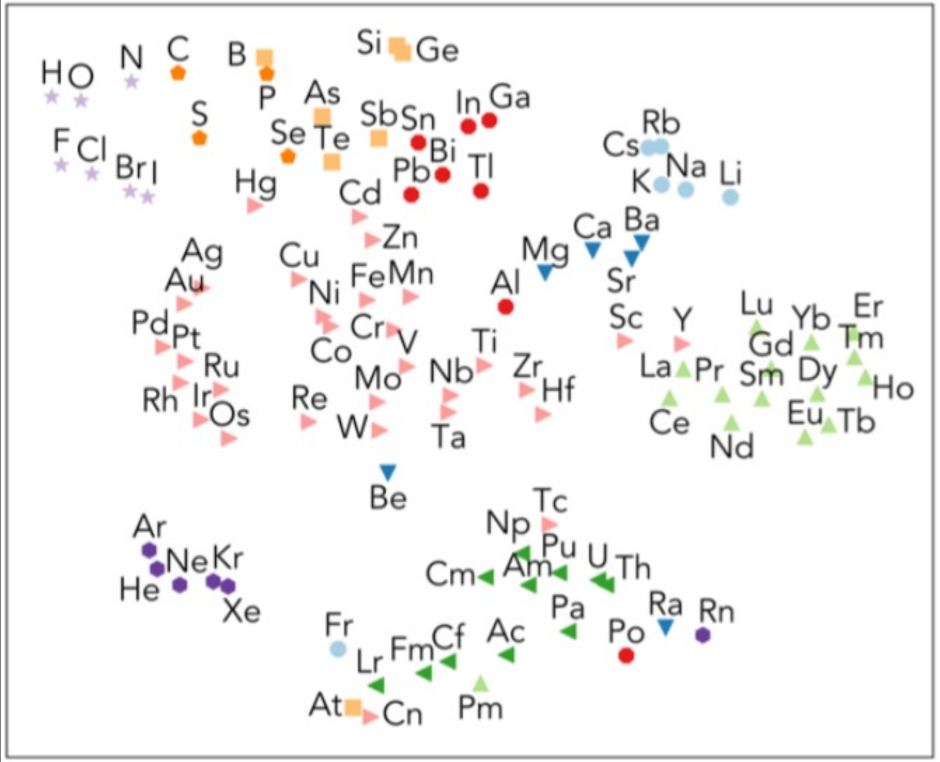

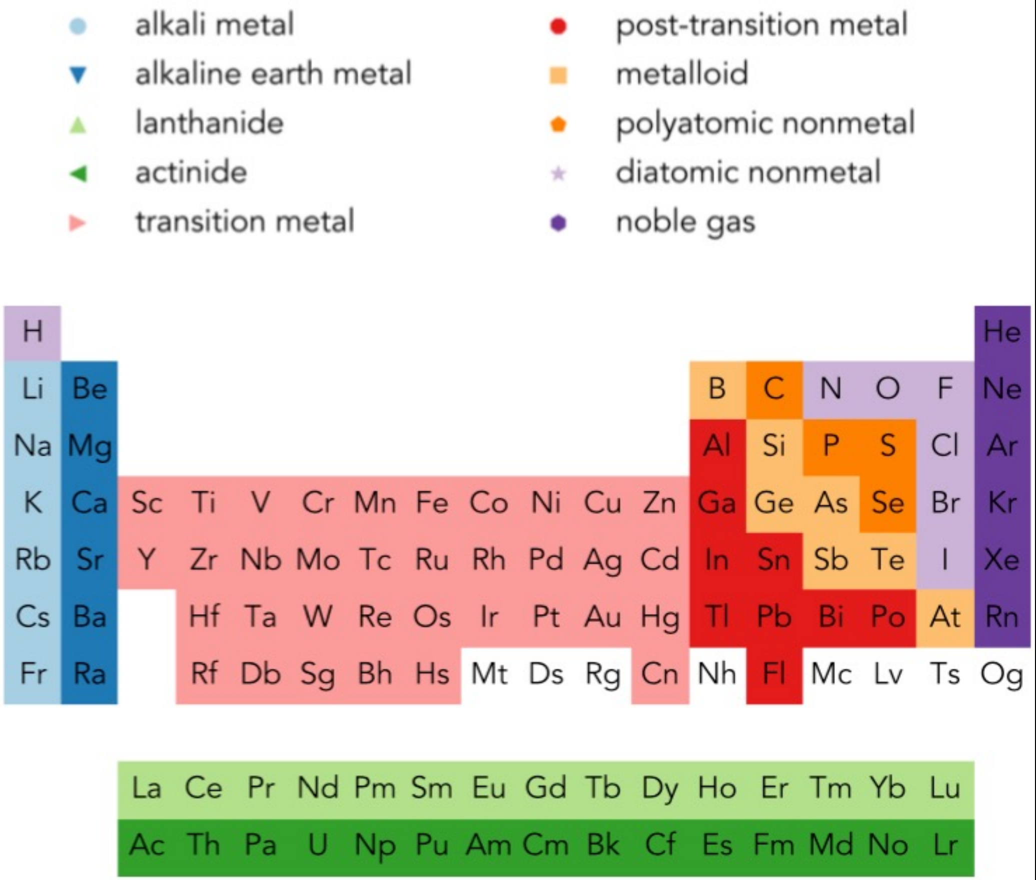

Word2vec's representation of the elements, projected onto two dimensions

Vahe Tshitoyan, Leigh Weston, John Dagdelen, Anubhav Jain

Filtering, De-Noise and Curating Data

AmeriFlux & FLUXNET: 750 users access carbon sensor data from 960 carbon flux data years; Developing ML to denoise data.

Arno Penzias and Robert Wilson discover Cosmic Microwave Background in 1965

Generate Videos

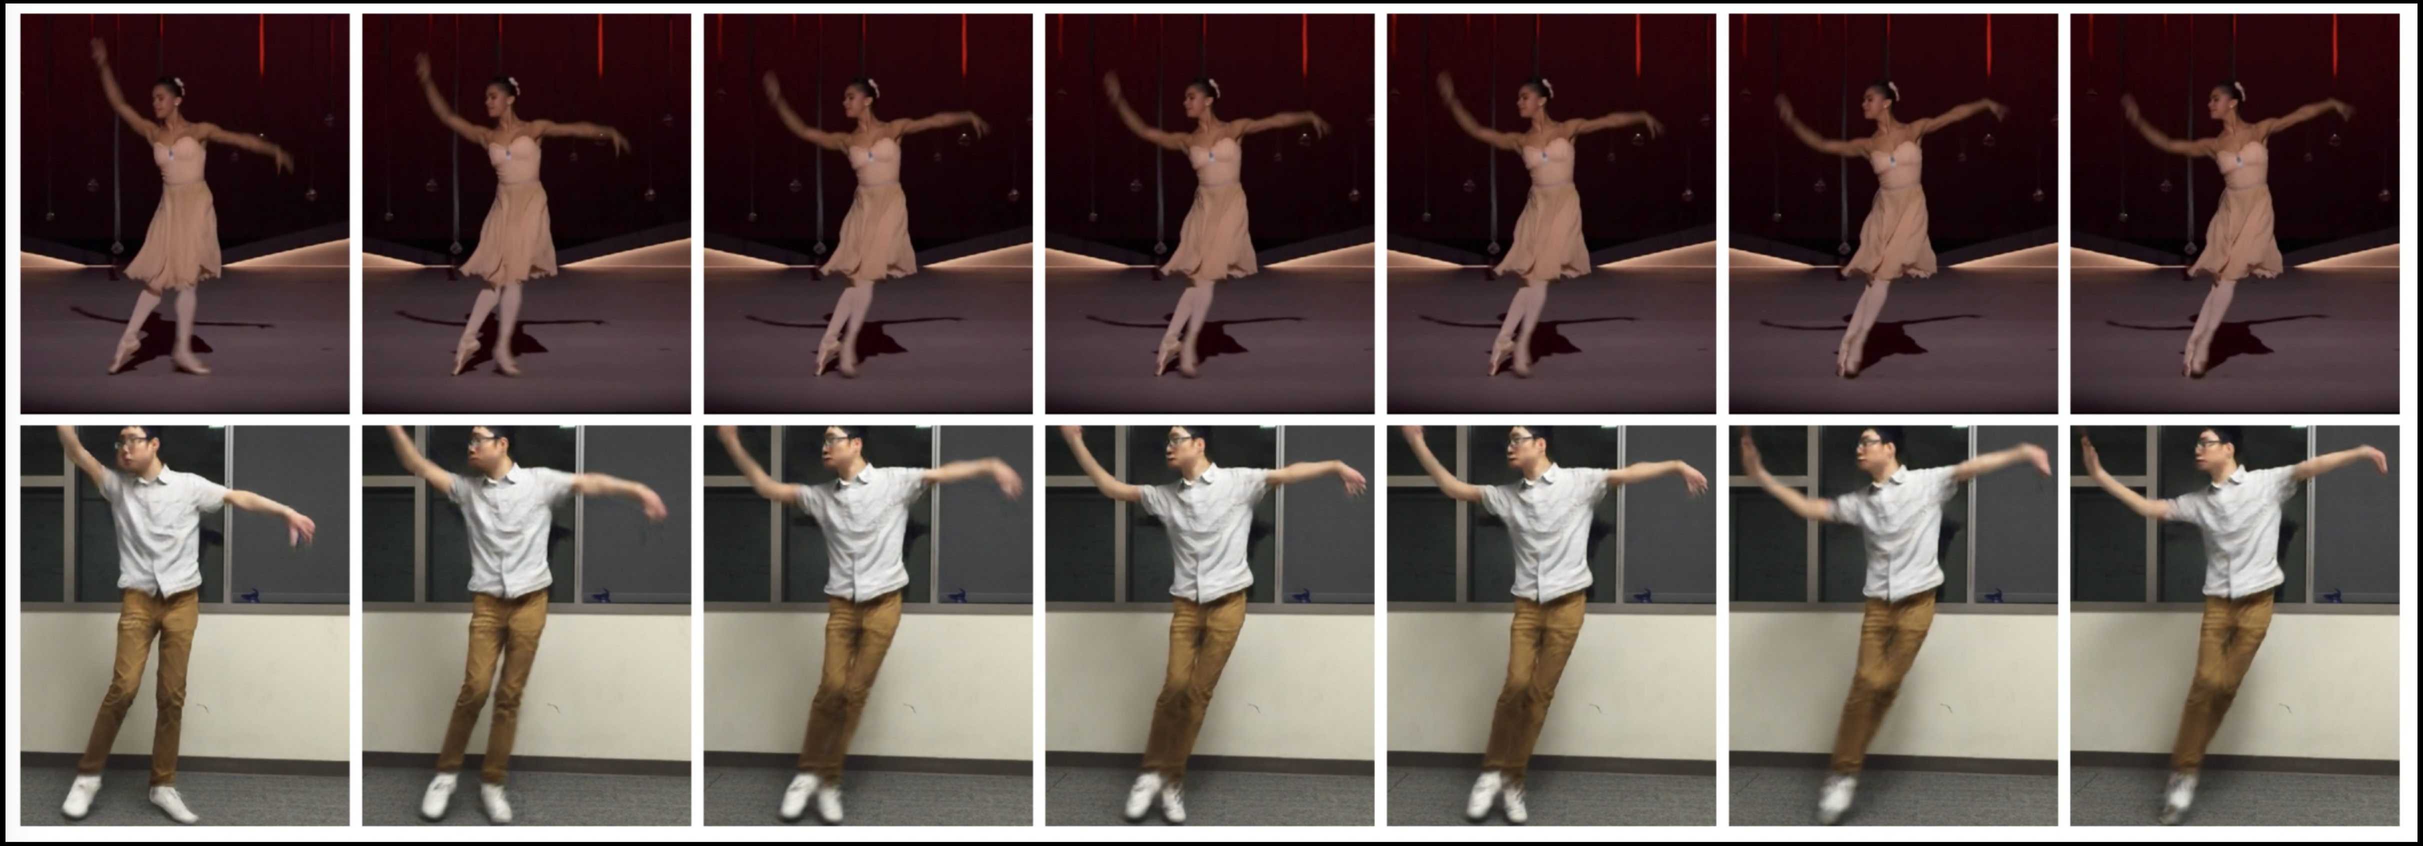

Caroline Chan, Shiry Ginosar, Tinghui Zhou, Alexei A. Efros, UC Berkeley

Generate Data from Expensive Experiments

Generate convergence maps of weak gravitational lensing, to help in understanding the physical laws governing the universe.

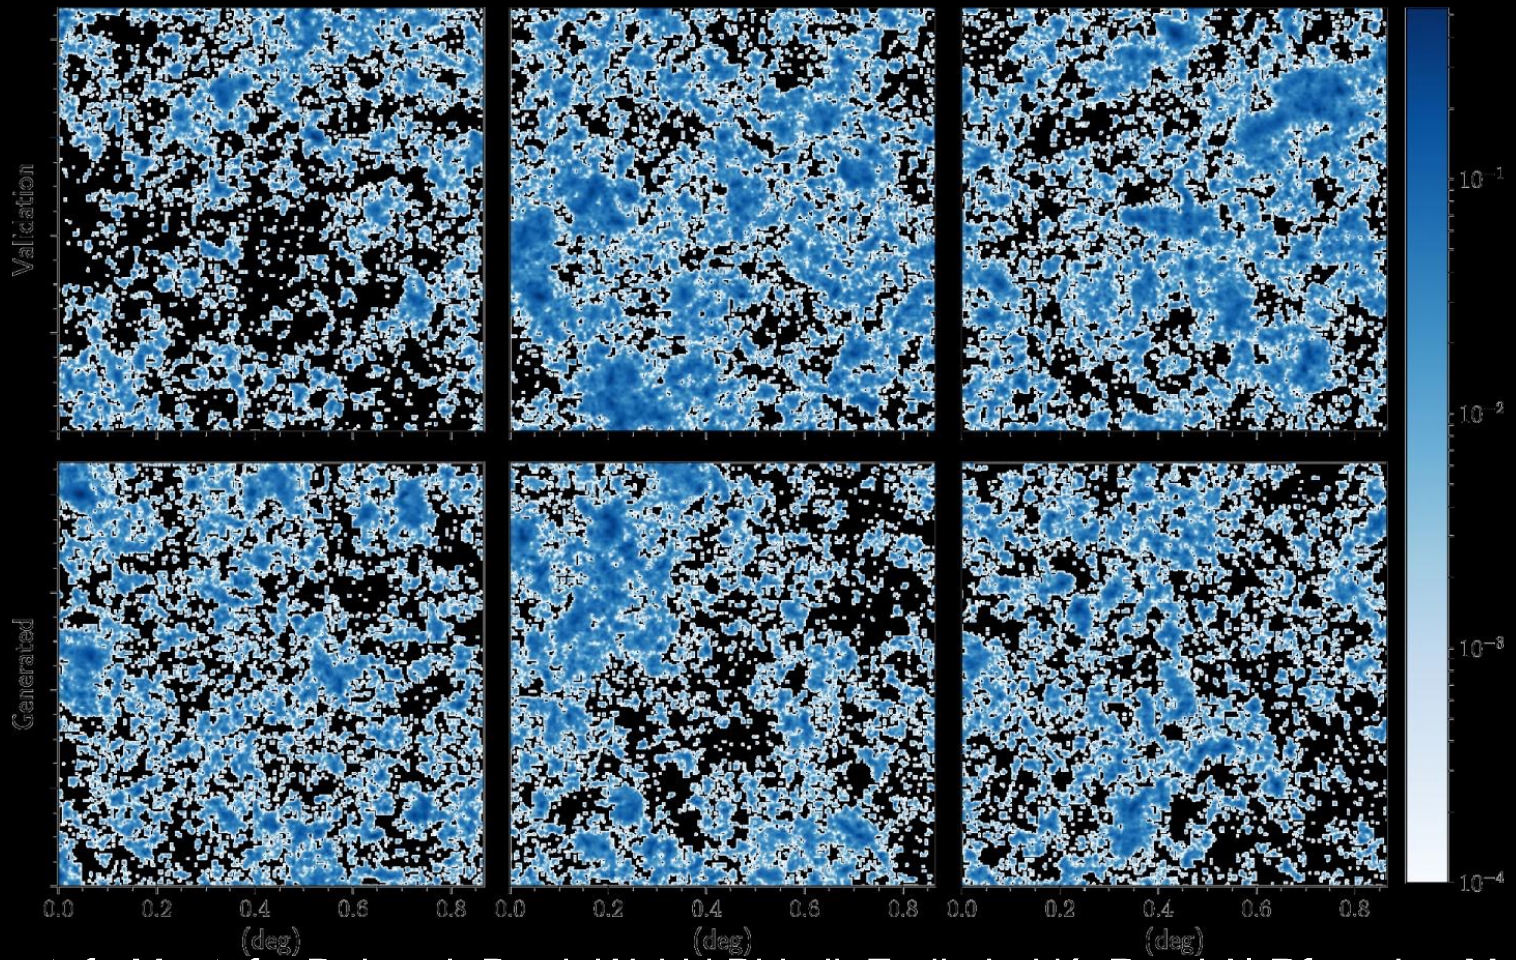

CosmoGAN: Mustafa Mustafa, Deborah Bard, Wahid Bhimji, Zarija Lukić, Rami Al-Rfou, Jan M. Kratochvil

Automated experiments

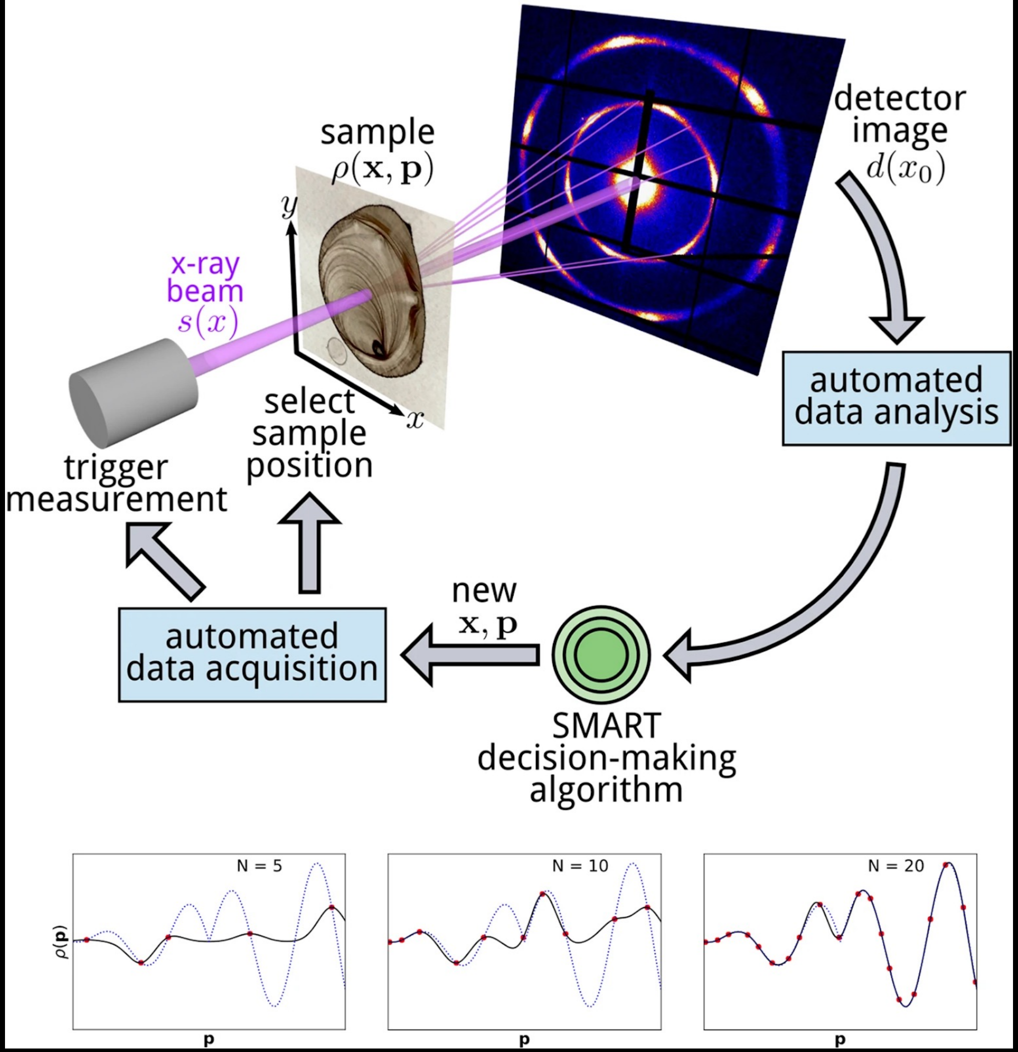

Utilization and robustness

- AI-based autonomous discovery
- Decisions based on small datasets
- Uncertainty estimates

Source: CAMERA Project, PI James Sethian
Slide input: Lavanya Ramakrishna

Digital Twins

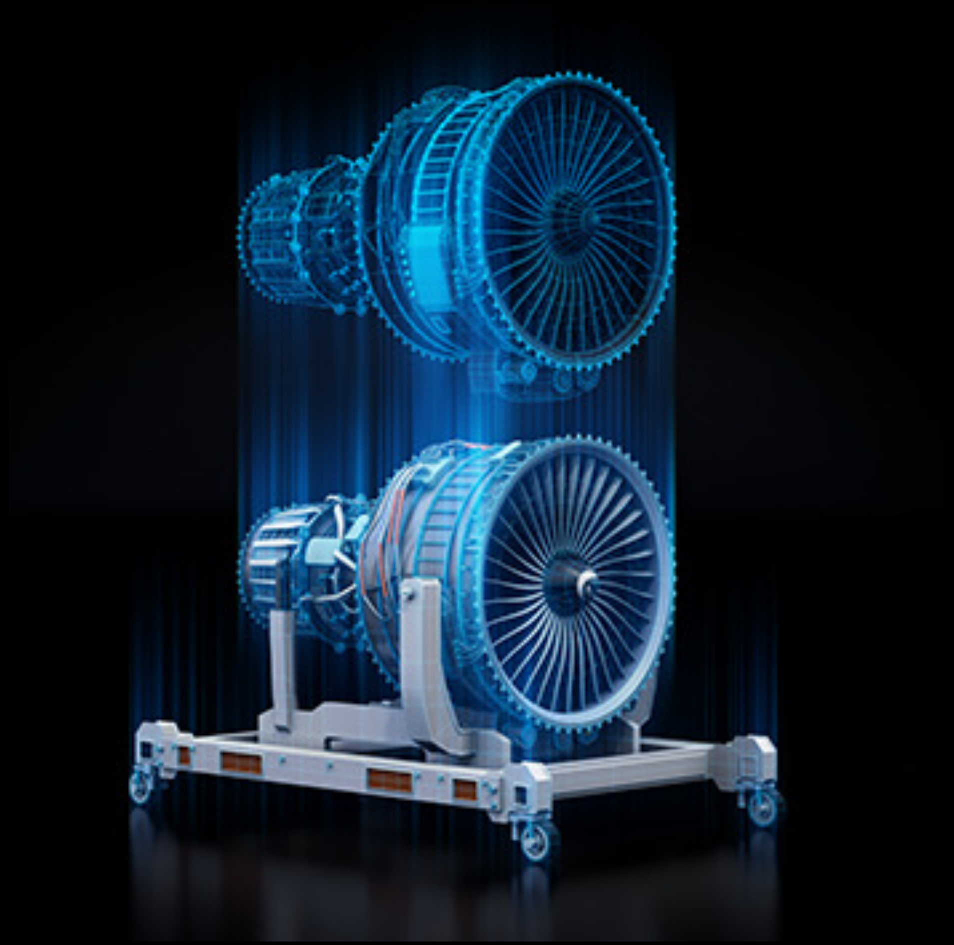

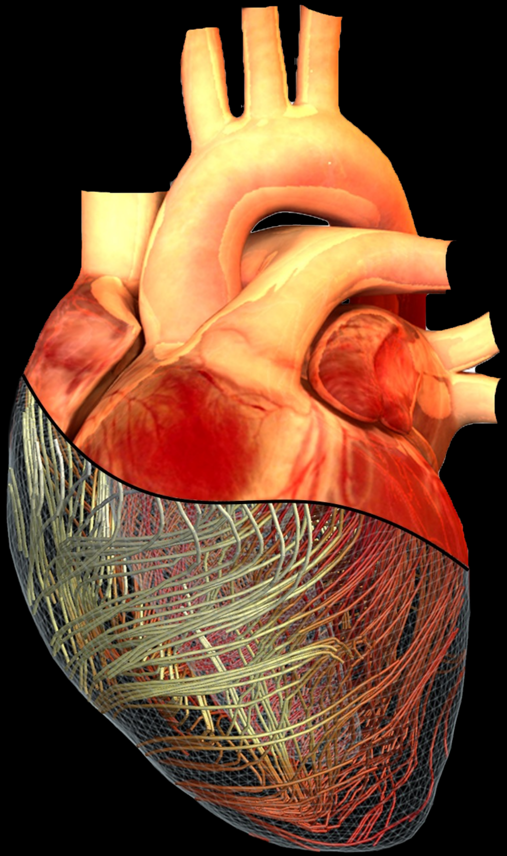

- Simulations
- Sensors / data
- Multi-level
- Real-time

AI in Science

Draft NAIRR report is out

The economic model is key

Cloud vs HPC : **It's all about the Business Model in '35**

Cloud	HPC
Focus on storage	Focus on computing (flop/s)
Cheap(est) commodity component	High end components (some specialization)
Commodity networks	High performance networks
Pay as you go	Purchased for mission; pay in non-fungible "hours"
< 50% utilization	> 90% utilization
On-demand access	Large jobs wait in queues
Multiple jobs per node	Dedicated set of nodes
On-node disks (air cooled)	Separate storage (compute liquid cooled)

Policy and business model

Rethink software

Transform publishing, research, teaching!

- Higher level
- Different interaction

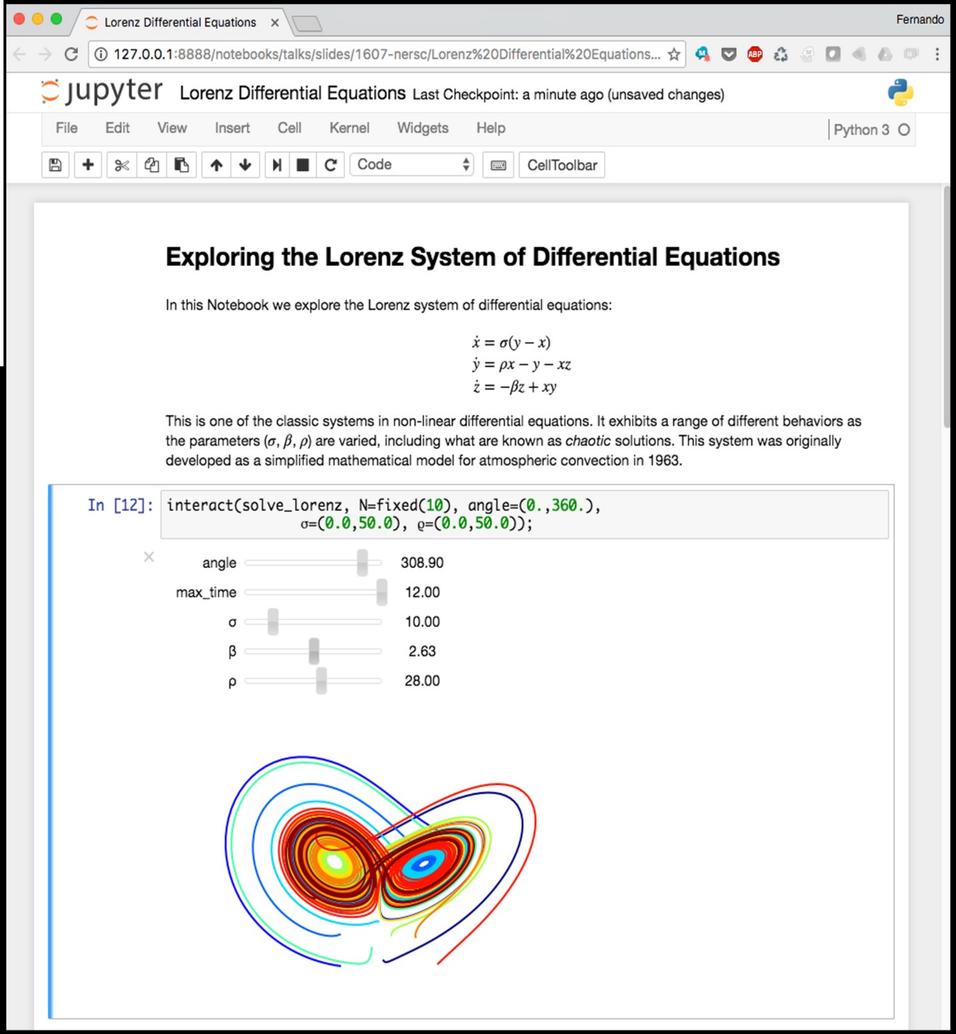

Old programming models never die, they just get buried under layers!

How did we get here?

- **Computing demands** continue to grow
- The benefits of more **weak scaling** are limited
- **Computing technology** has hit several "walls"
- The **computing industry** has changed dramatically
- **AI methods** are having huge impacts elsewhere
- **Quantum computing** potential for science still unknown
- **Cloud computing** is dominating the computing industry
- **Global supply chain** issues present uncertainties

We need a new strategy for post-exascale computing