# New hardware architecture underpins ambitious global project

Bob Bishop

ICES Foundation, Geneva

SCFE2018 Warsaw Poland, 13 March 2018

# Building an open, integrated holistic model of Planet Earth for decision support, disaster reduction and public good

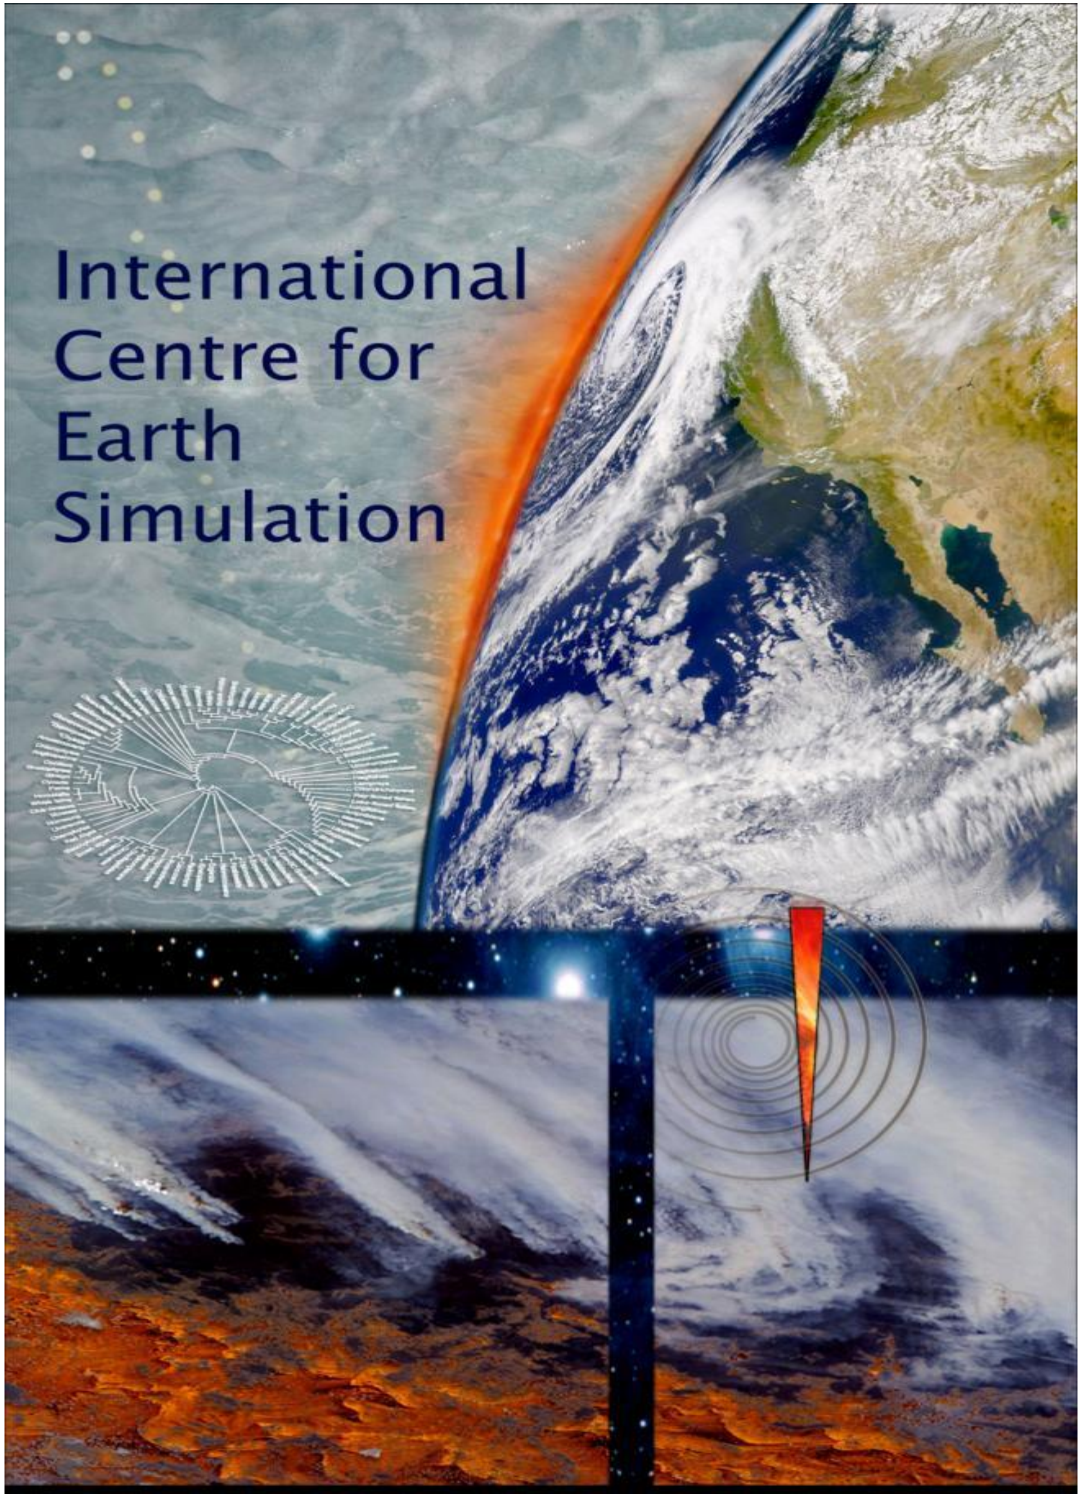
International
Centre for
Earth
Simulation

# www.icesfoundation.org

- The Foundation focuses on the holistic totality of our Earth, the Solar System, Cosmology, and Human Society
- The Foundation's website contains 20,000 URLs of which 5,000 point to new science discoveries over the last 20 years
- We believe now is the time to bring this new science, with new models and new hardware together in one fell swoop!

# Here’s what new models & algorithms can do! QLARM

Combining data available from the UN and satellite info:

- 1.9 million settlements, population, building stock profile
- soil type, landscape & crustal structure

- Calculate shaking level from an actual earthquake epicentre
- Estimate degree of destruction, fatalities & injured
- Distribute to first responders within 30 minutes via Mail Chimp
- QLARM alerts – a major help to earthquake response planning
- Alternatively, calculate such numbers as a precautionary measure and take mitigating measures to protect society ahead of time

# COMMENT



OMAR HAVANA/GETTY

**Rescue workers in Kathmandu, where a magnitude-7.8 earthquake killed 10,000 people in April 2015.**

# Report estimated quake death tolls to save lives

Earthquake survivors could be rescued more quickly if the media communicated the number of likely fatalities from the outset, argues **Max Wyss**.

For a decade, seismologists have been able to generate fast, reliable estimates of the number of people likely to have been killed in an earthquake, to within a factor of two or three[1]. But these valuable tools are still not being used to save lives. Knowing whether 10 or 10,000 people might have died tells governments how much effort they should direct to rescuing people buried under rubble. Time is short — few individuals survive for more than three days.

Fatality predictions are sent by e-mail within half an hour of a harmful quake anywhere in the world, for free, by the International Centre for Earth Simulation (ICES) Foundation and the US Geological Survey (USGS). Yet most officials, first responders and journalists are unaware of this. Instead, decisions are based on information that trickles in from the scene. The death toll is generally underestimated. First accounts come from areas where communications networks still function — far from the epicentre. No information flows from the most devastated areas. Rescue efforts are too little, too late. Many people die needlessly.

I have seen this happen many times, as a seismologist who forecasts earthquake losses for ICES using its QLARM fatality-prediction model. On 24 February 2004, a magnitude-6.4 earthquake struck Morocco at 02:27 local time. Before dawn, the Swiss government offered to send help — its disaster-response team had received my alert indicating that up to 1,000 fatalities were likely. A Moroccan official turned them down. ▶



© 2017 Macmillan Publishers Limited, part of Springer Nature. All rights reserved.

# Kaoshiung M6.4 earthquake on 06 February 2016

# We have arrived in a data intensive era

| Multiplying Factor | SI Prefix | Scientific Notation | Name |
|---|---|---|---|
| 1 000 000 000 000 000 000 000 000 | Yotta (Y) | $10^{24}$ | 1 septillion |
| 1 000 000 000 000 000 000 000 | Zetta (Z) | $10^{21}$ | 1 sextillion |
| 1 000 000 000 000 000 000 | Exa (E) | $10^{18}$ | 1 quintillion |
| 1 000 000 000 000 000 | Peta (P) | $10^{15}$ | 1 quadrillion |
| 1 000 000 000 000 | Tera (T) | $10^{12}$ | 1 trillion |
| 1 000 000 000 | Giga (G) | $10^{9}$ | 1 billion |
| 1 000 000 | Mega (M) | $10^{6}$ | 1 million |
| 1 000 | kilo (k) | $10^{3}$ | 1 thousand |
| 0 001 | milli (m) | $10^{-3}$ | 1 thousandth |
| 0 000 001 | micro (u) | $10^{-6}$ | 1 millionth |
| 0 000 000 001 | nano (n) | $10^{-9}$ | 1 billionth |
| 0 000 000 000 001 | pico (p) | $10^{-12}$ | 1 trillionth |
| 0 000 000 000 000 001 | femto (f) | $10^{-15}$ | 1 quadrillionth |
| 0 000 000 000 000 000 001 | atto (a) | $10^{-18}$ | 1 quintillionth |
| 0 000 000 000 000 000 000 001 | zepto (z) | $10^{-21}$ | 1 sextillionth |
| 0 000 000 000 000 000 000 000 001 | yocto (y) | $10^{-24}$ | 1 septillionth |

*Chart courtesy 19th General Conference on Weights & Measures 1991*

# Data is growing faster than Compute Power

Megabytes, Gigabytes, Terabytes, Petabytes, Exabytes, Zetabytes, Yottabytes

$10^6$ $10^9$ $10^{12}$ $10^{15}$ $10^{18}$ $10^{21}$ $10^{24}$

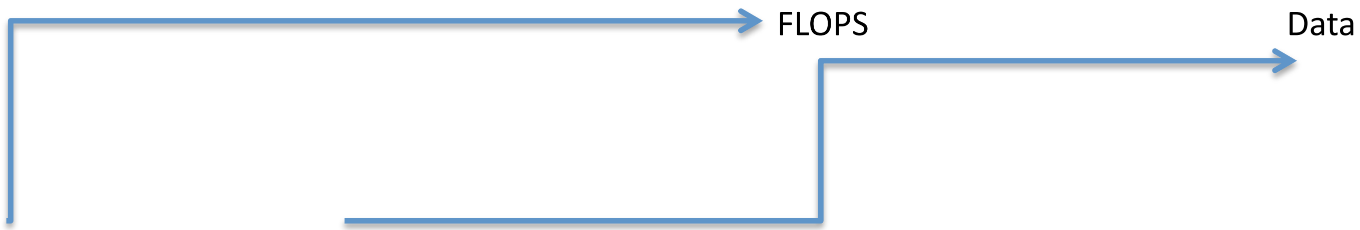


Exaflops vs. Exabytes

Data doubles every one year

FLOPS double every two years (Moore's Law)

Compute power will not keep up with growth of data

Data is overwhelming compute power in many applications

Much of this data is ***unstructured*** (image, A-V, analog, real-time IoT)

# Unstructured vs. Structured Data

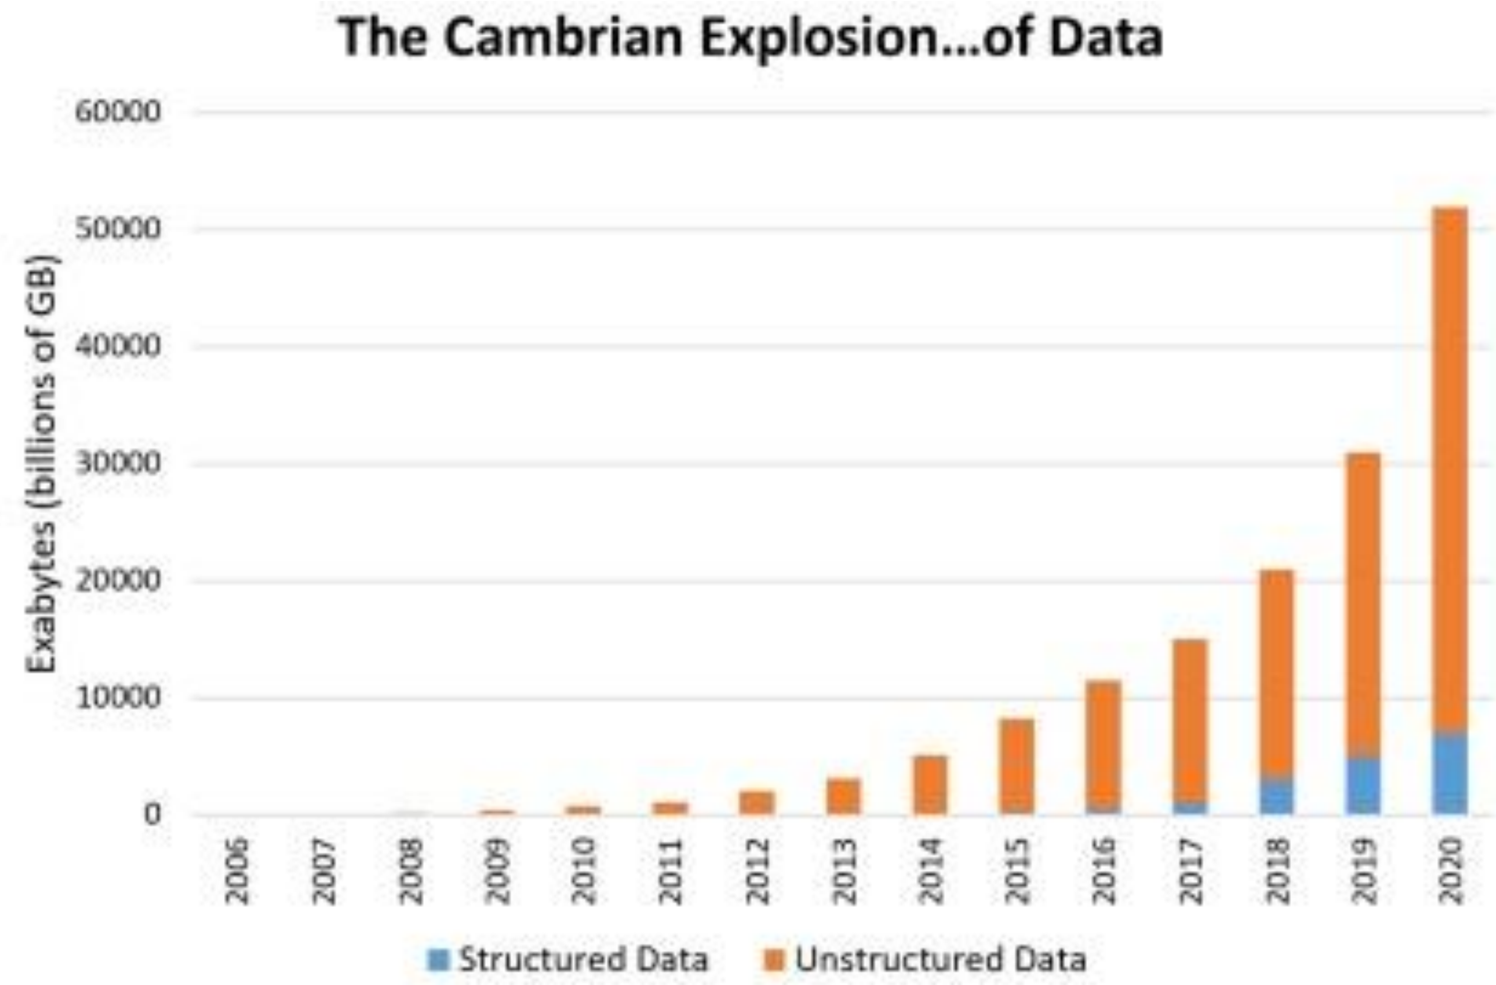


*Chart from Synopsis*

# The problems of the Big Data era are many:

Data Movement – electrical energy to move the data
Data Latency - finding the data
Bandwidth - transfer speed
Data File Format – usability
Data Ontologies
Data Quality
Data Access Privileges – authority
Real-time data update mechanisms
Data in sparse matrices vs. data in dense matrices

Such problems can grind Von Neumann machines to a halt!

# Evolution of total Top500 Supercomputing FLOPS

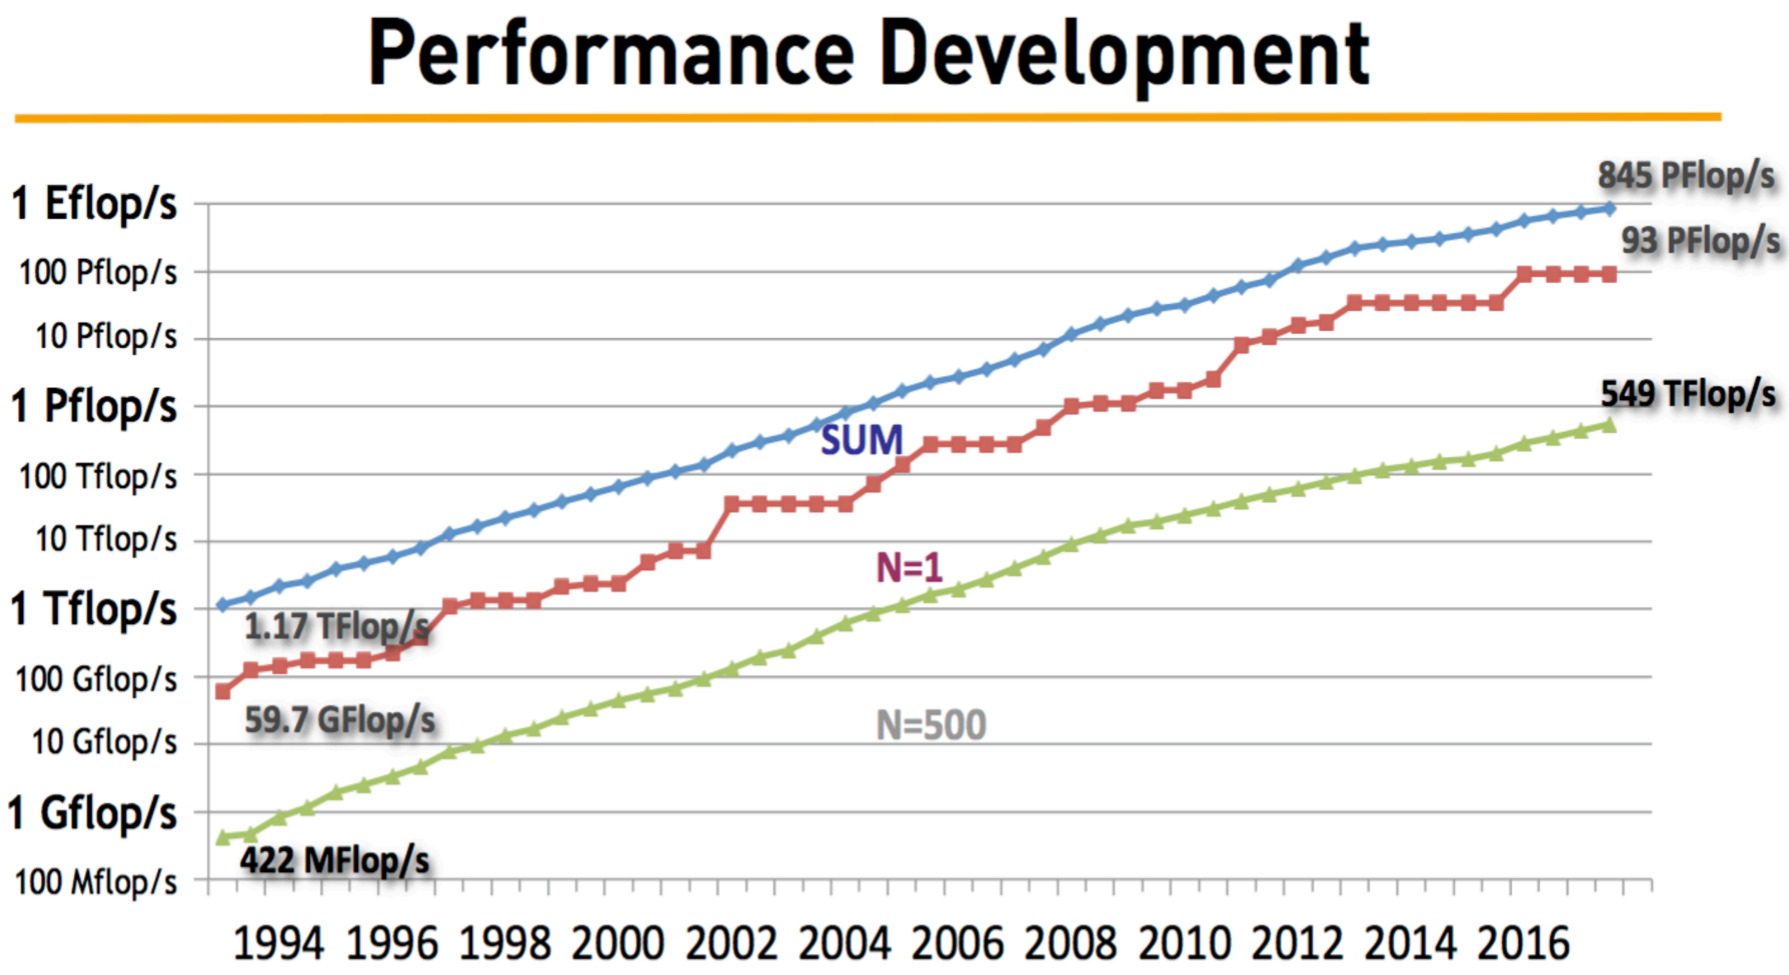


*Chart from Top500 list*

# FLOPS no longer tell the whole story of computing

- Top500 Supercomputers – dense matrices - Linpack
- HPCG – Conjugate Gradient – sparse matrix multiply
- spMV – sparse matrix & vector multiply
- HPC Challenge – suite of 7 tests – memory access
- STREAM – sustainable memory bandwidth
- SPEC/HPG – MPI, OMP, OCL comparisons
- BLAS – basic linear algebra subprograms
- Firehose – find events in streaming data
- Graph 500 – data intensity tests
- Green 500 – energy efficiency – FLOPS/Watt
- Green Graph 500 – energy consumption of big data apps

# Benchmark comparison of HPC systems

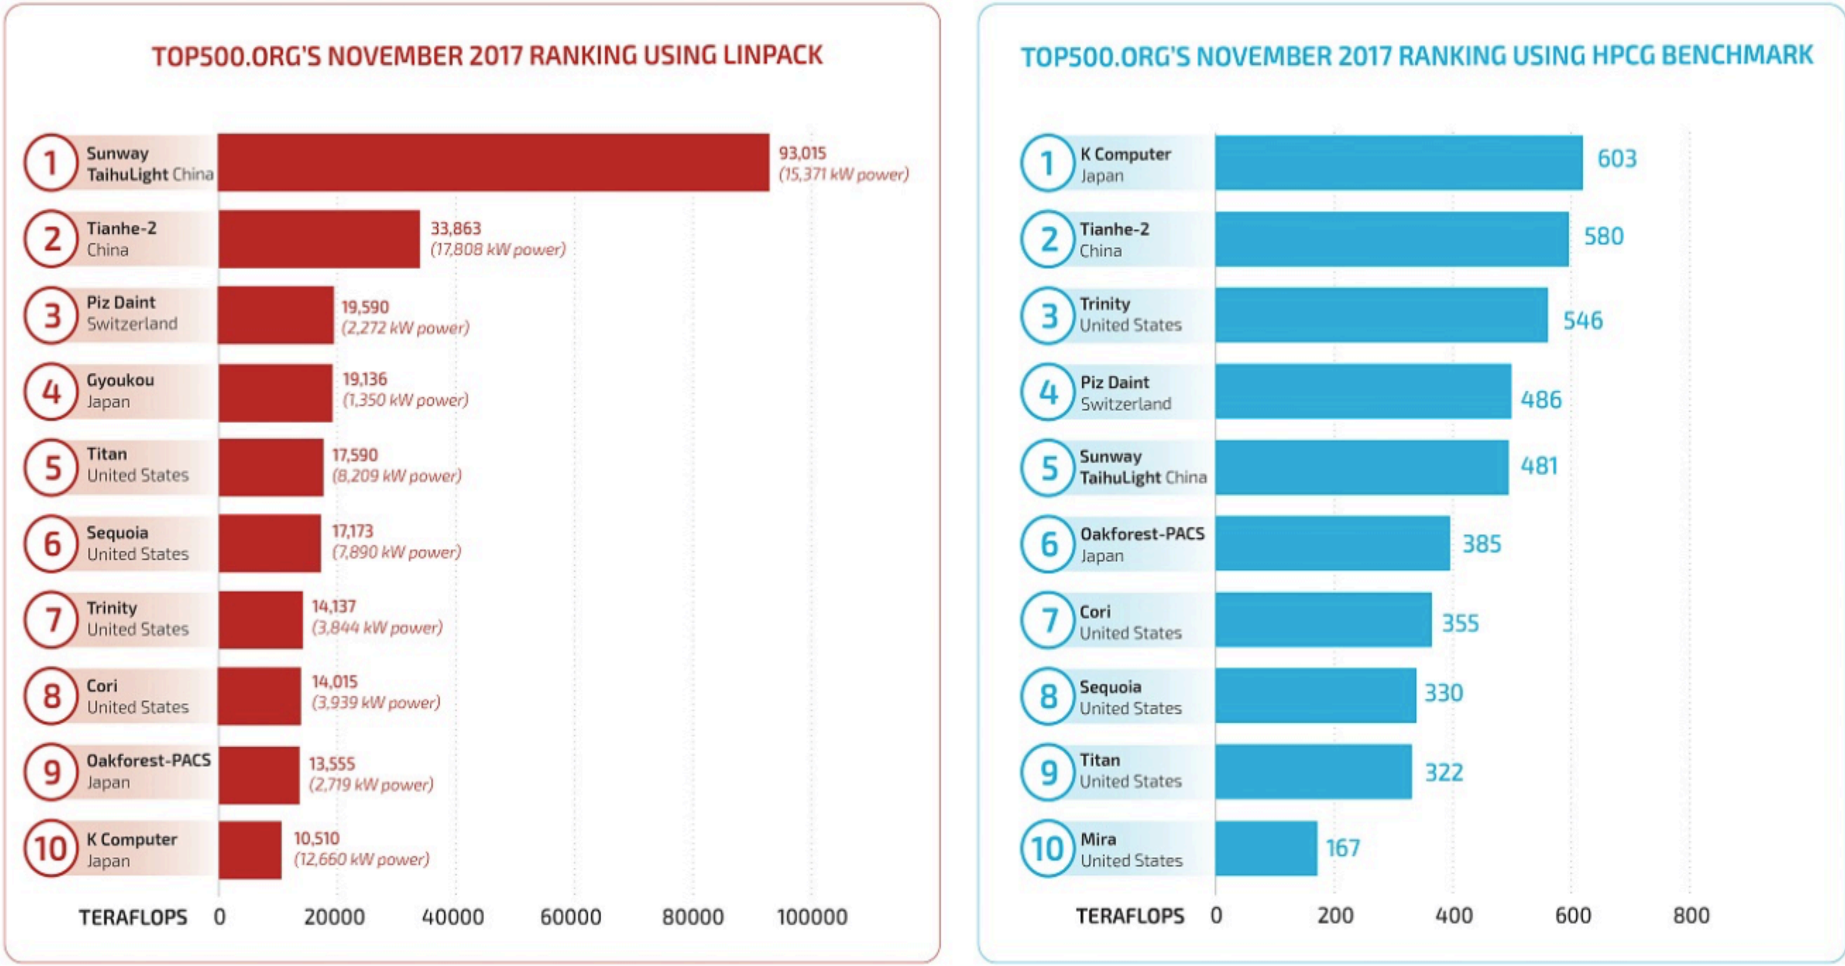


*Chart from Sandia Labs*

Benchmarking HPC systems demands knowledge of future workloads!

# Future Workloads = Wicked Problems

- Climate Change
- Global Warming
- Sea Level Rise
- Ocean Acidification
- Extreme Weather
- Flood, Drought, Heatwave
- Environmental Disaster
- Desertification
- Earthquake
- Geomagnetic Storm
- Volcanic Eruption
- Asteroid Impact
- Biodiversity reduction

Political Collapse
Social Instability
Mass Migration
Planetary Health
Pandemics
Pollution
Deforestation
Soil Depletion
Resource Depletion
Famine, Food Security
Energy Security
Nuclear Accident
Military intervention

***Such wicked problems involve Big Data and real-time updates!***

# When will usable Quantum Computing arrive?

- Most likely, 10 years from now to have some useful apps
- Transformative Quantum may be 20 years away
- Quantum Advantage & Quantum Supremacy

  ~ 50 qubits + error detection
- A qubit is neither 1 or 0, but in some hybrid state of uncertainty
- As soon as a qubit is observed, it immediately locks down to 1 or 0
- Error correction alone takes much more than 100 qubits today!

# Reducing error correction qubits is key for QC

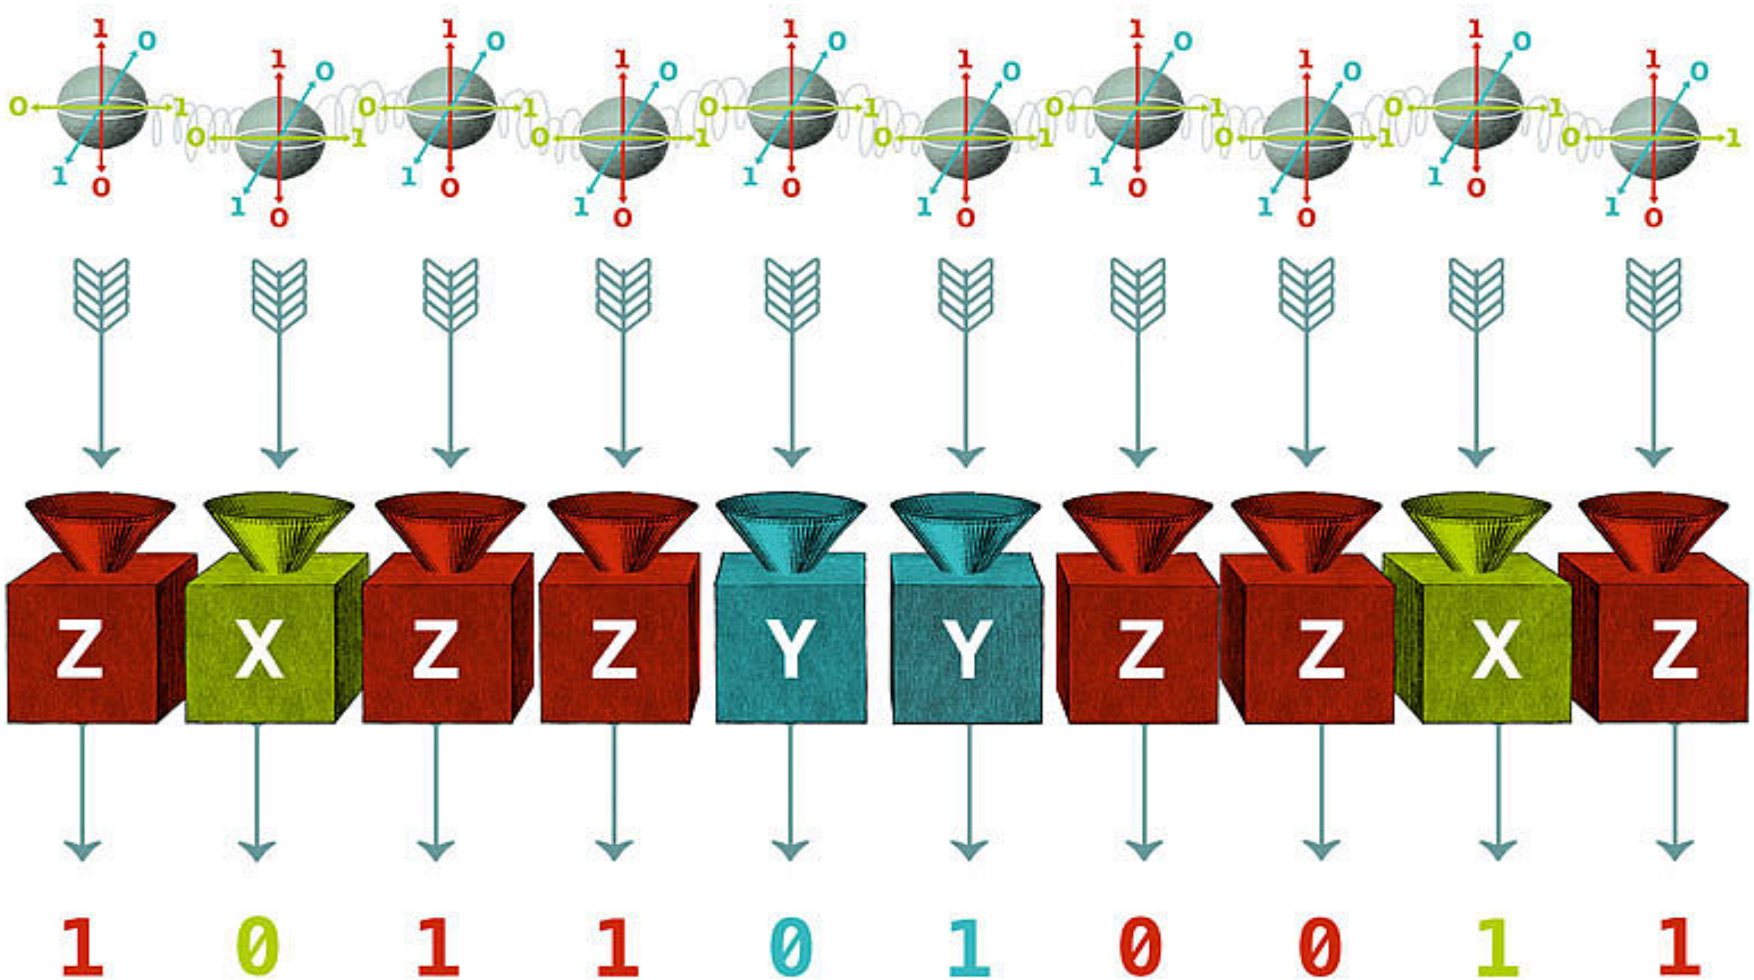


*Chart © Juan Palomino*

# Google's error correction goals

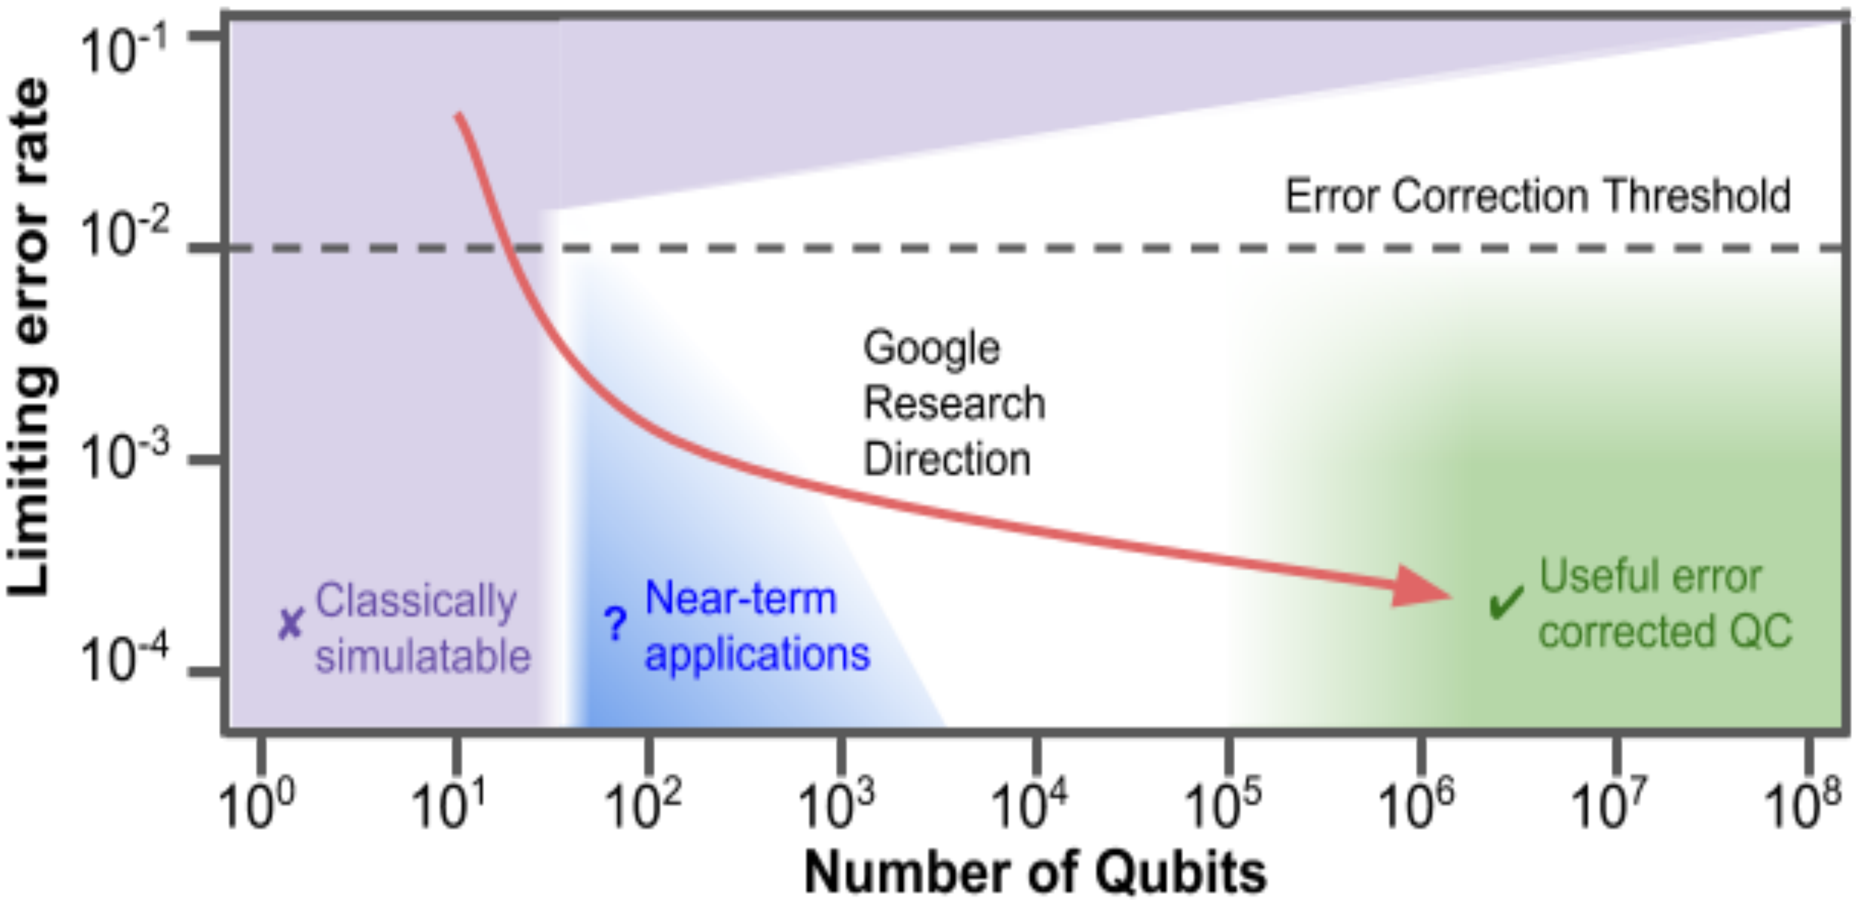

# When will Neuromorphic Computing arrive?

- Most likely, 10 years from now to have some useful apps
- Transformative Neuromorphics may be 20 years away
- Neuromorphic Computing is an attempt to mimic human brain architecture:

  - IBM TrueNorth CMOS chip with 4096 cores, and simulating over 1 milliion neurons, and 46 billion synaptic ops/second
  - SpiNNaker spiking neural network SoC, 0.5M ARM cores (Manchester U)
  - BrainScaleS direct silicon-based image of neuronal network (Heidelberg U)

***The universal problem is power consumption, since our brain only uses 25 watts!***

# Overall structure of the human brain

- 80 billion neurons
- 500 trillion synaptic junctions
- 100 km of fibre
- 25 watts of power
- 3 meals/day

# The synaptic junction is the main diode

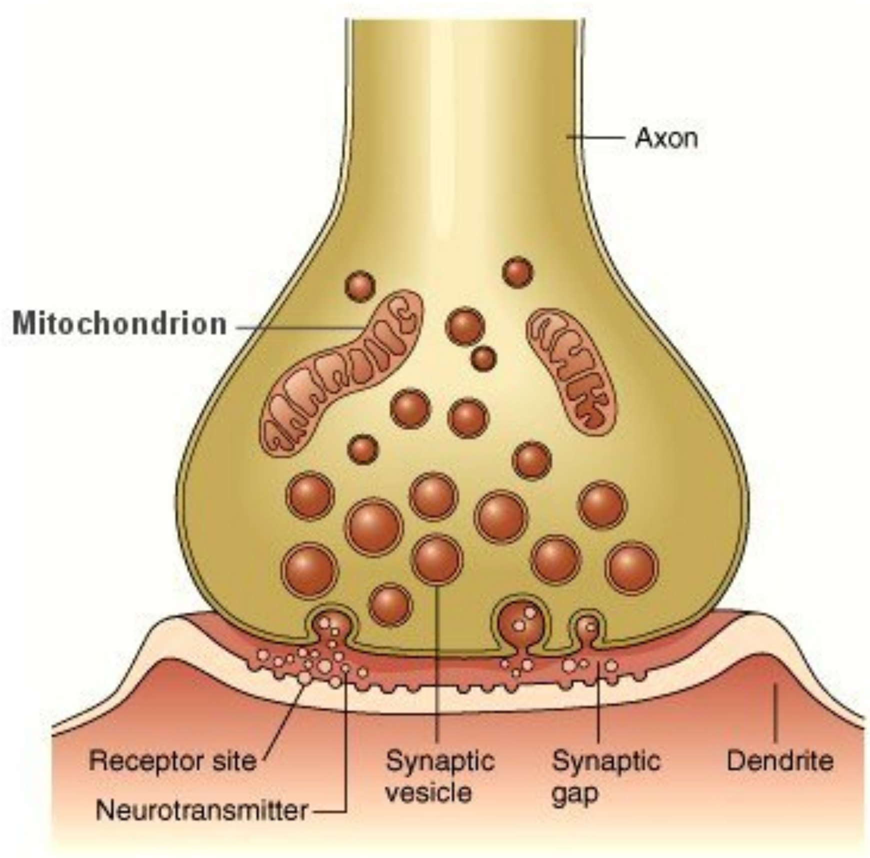

# Cortical column organisation level

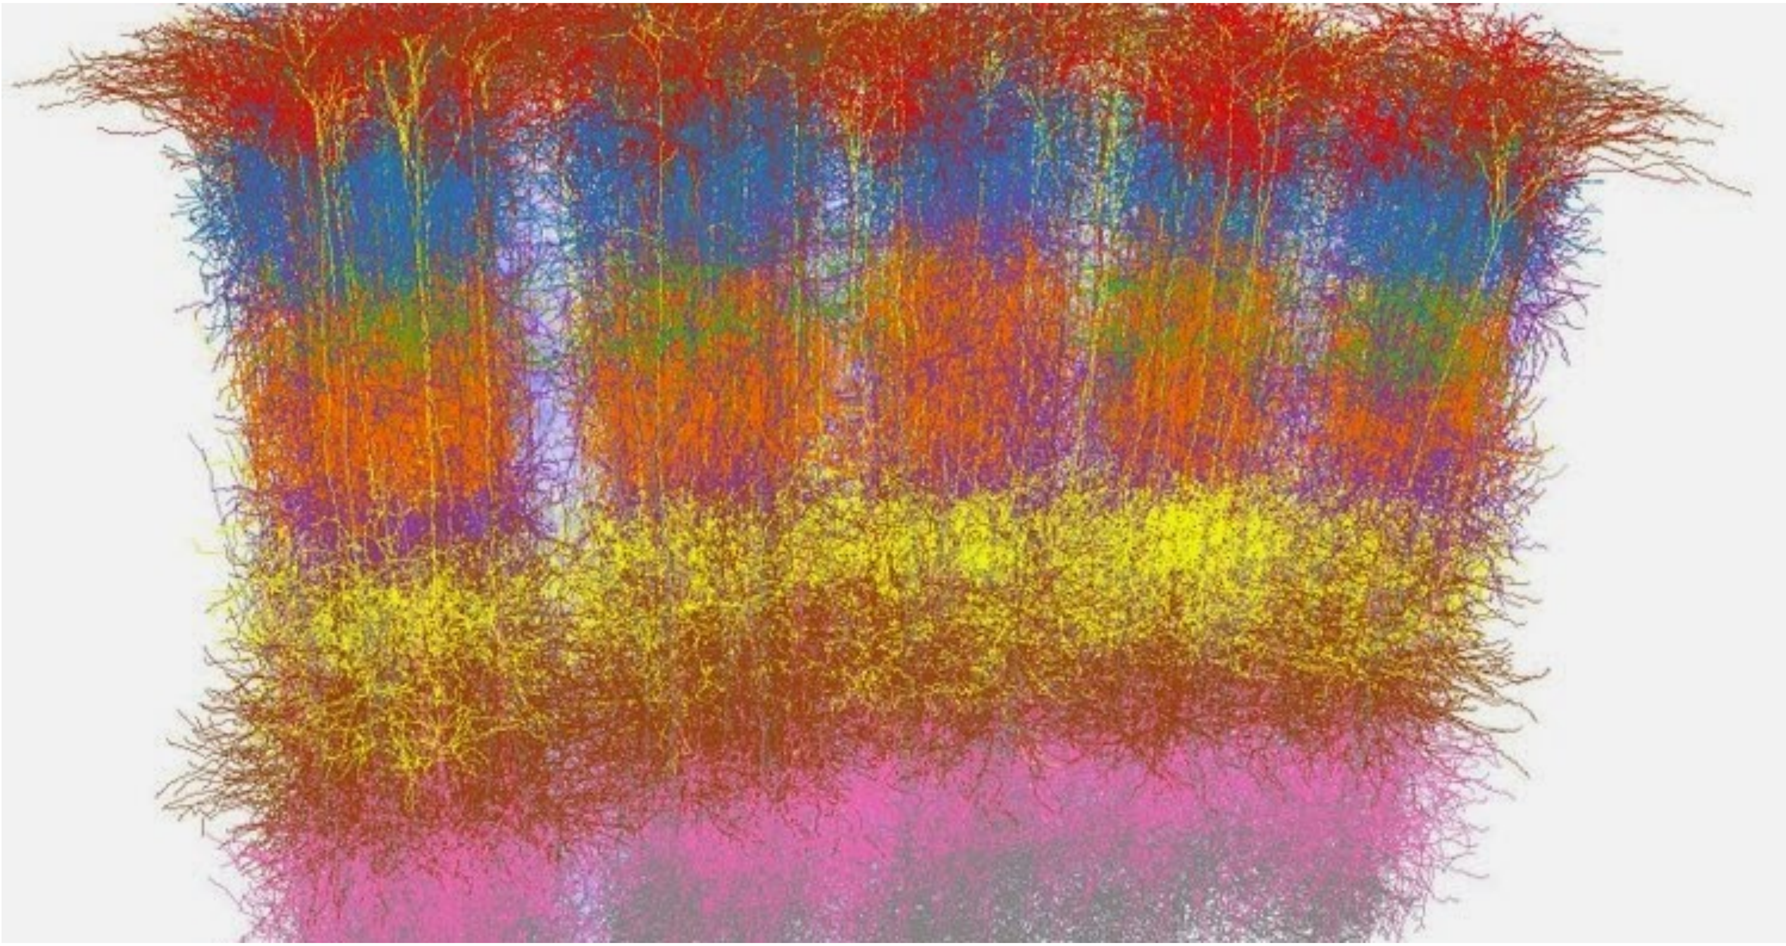

# Contrasting the Brain to Digital Computing

- Analog – Digital
- Wet - Dry
- Plastic – Rigid
- Resilient - Fragile
- Intuitive – Logical
- Creative – Deterministic
- Self organised - Designed
- Self assembled - Manufactured
- Asynchronous – Mostly synchronous
- Embedded memory – Segregated memory
- Variable neural network – Fixed instruction set

# Does our brain act like a quantum device?

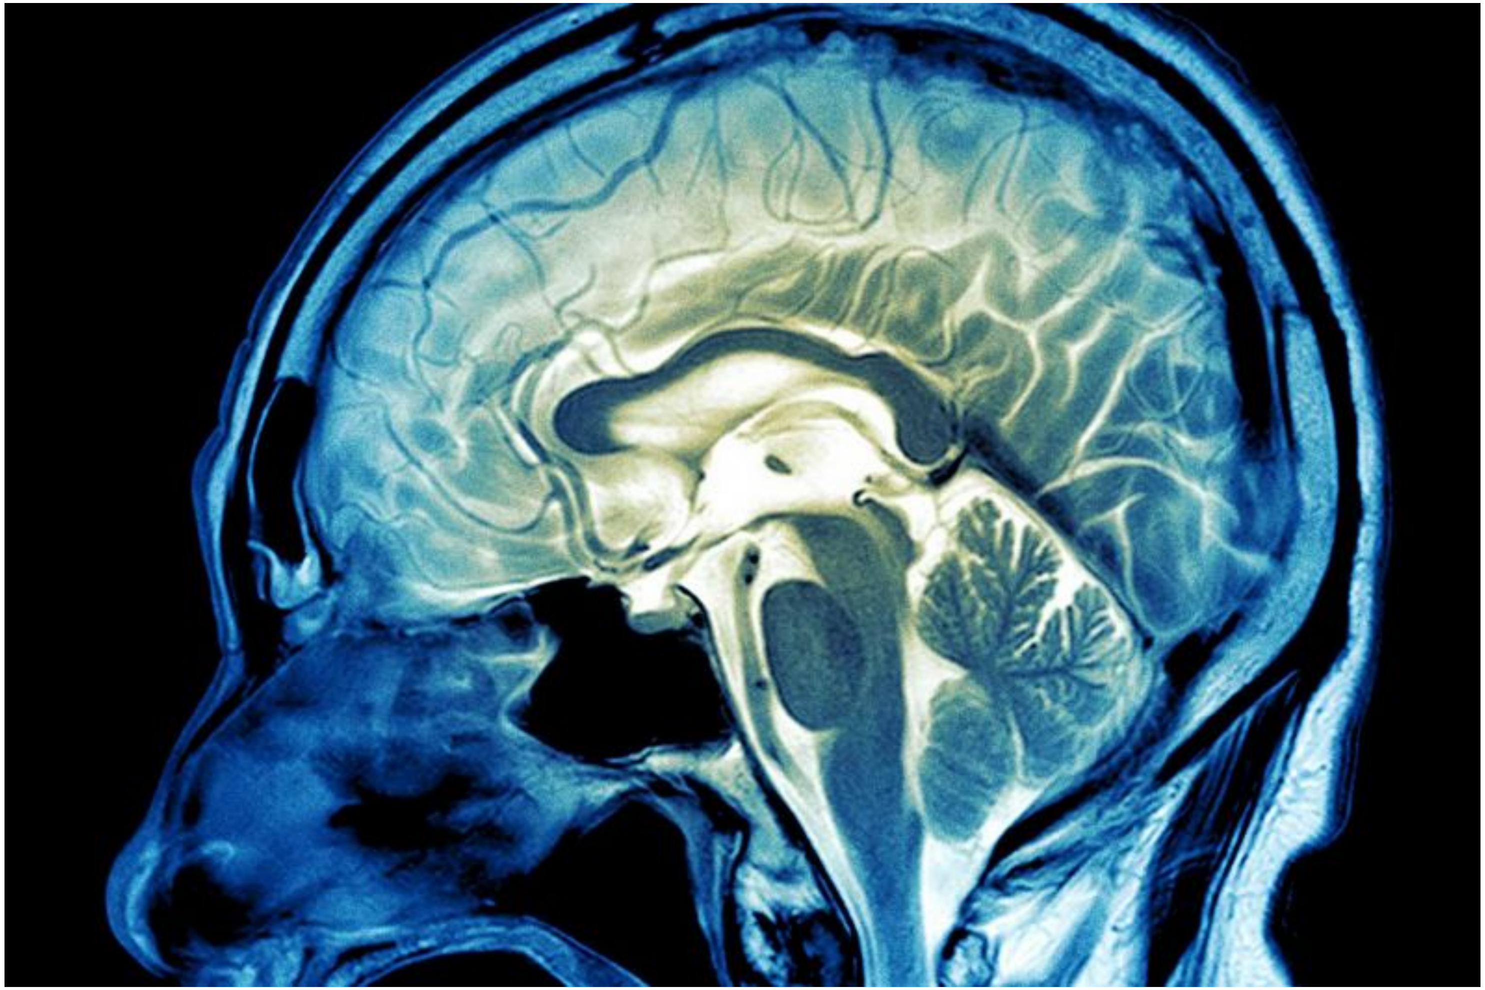

# We can take 8 steps before QC or NC arrive!

- Embed AI/Machine Learning with conventional HPC apps
- Improve data ingestion from IoT devices
- Move intelligence to the network's edge
- Minaturization and 3D memory
- Deploy Processor-in-Memory
- Deploy silicon-photonics
- Deploy 4D VR & AR visualization
- Establish global-local downscaling to city level

# Step 1: Embed AI/ML with HPC apps

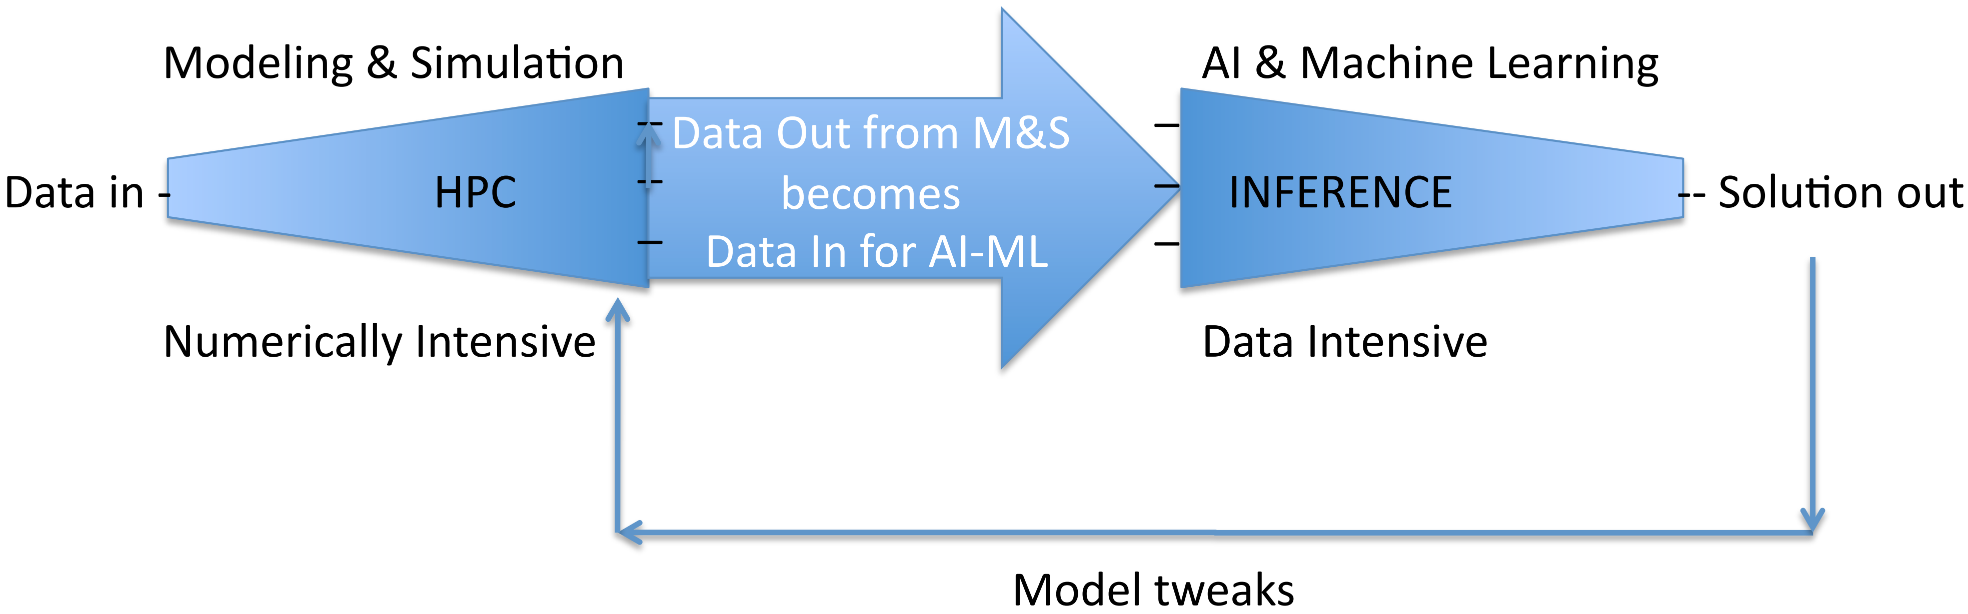


*Chart derivative from Goh Eng Lim*

# Step 2: Improve data ingestion from IoT devices

- 150 billion IoT devices by 2030
- we need to install big pipes @ 100 Gbps
- and we need to install big buffer data storage units
- driverless cars will each produce 4 TB of data per day!

# Step 3: Intelligence at the Edge of the network

- local decision making instead of via remote cloud
- overcome latency issues
- space based platforms
- aerial based platforms
- mobile intelligence
- pocket devices
- driverless cars

# Step 4: Minaturization and 3D memory

- Bring logic fab down from 11-7-5-3 nanometers
- Stack 3D memory planes & logic
- Smart memory controllers
- Processor-in-Memory (PIM)

# Step 5: PIM resolves many Big Data issues:

Here's how:

- Reduce data motion by an order of magnitude
- Move less data and over shorter distances
- Break up big jobs to many small jobs (*threads*)
- Break these threads into even smaller threads
- Asynchronously migrate these small *threads to the data locale* instead of moving *big data to the jobs*

# EMU Technology, Inc.

Peter Kogge
Founder

Jay Brockman
Founder

Ed Upchurch
Founder

Jim Hunt, Chairman & CEO

Richard Sheroff, VP Sales

Marty Deneroff, COO

Anne Vincenti, VP Marketing

# EMU deploys PIM with 3 levels of processing

- Smart memory controllers (Atomics)
- Thin cores near data in memory (FPGA or ASIC)
- and standard cores for overall sync (X86, Power + Linux OS)

# PIM architectural example

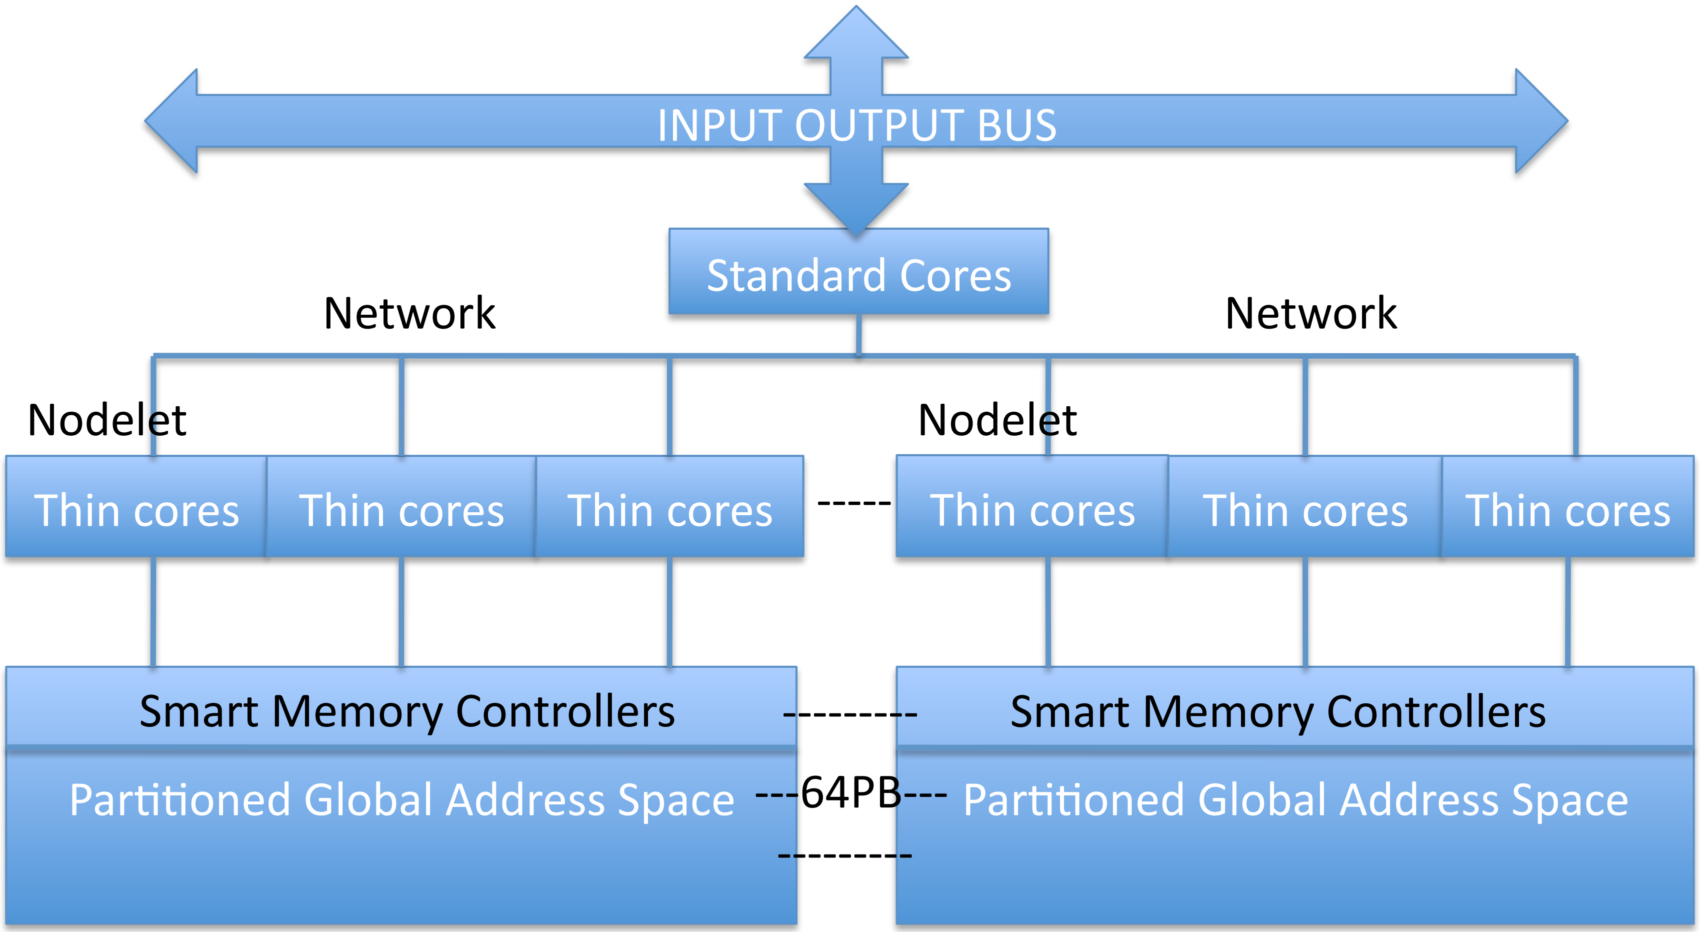


*Chart derivative from EMU Technology*

# EMU uses a large Partitioned Global Address Space

- 56-bit address space = 64 PB of memory
- EMU does not deploy L1, L2, L3 cache logic
- EMU does not deploy cache coherent logic

# EMU Programming Model

- CilkPlus is an ideal programming language to spawn such threads
- CilkPlus supports thread migration to data locality
- CilkPlus supports C++11 and is a familiar environment

Jeffrey Young, Researcher at Georgia Tech:

*We were completing ports of code and algorithms immediately!*

*We were thrilled with the software environment - we could focus on our research instead of learning new programming models.*

# SpMV EMU desktop Performance

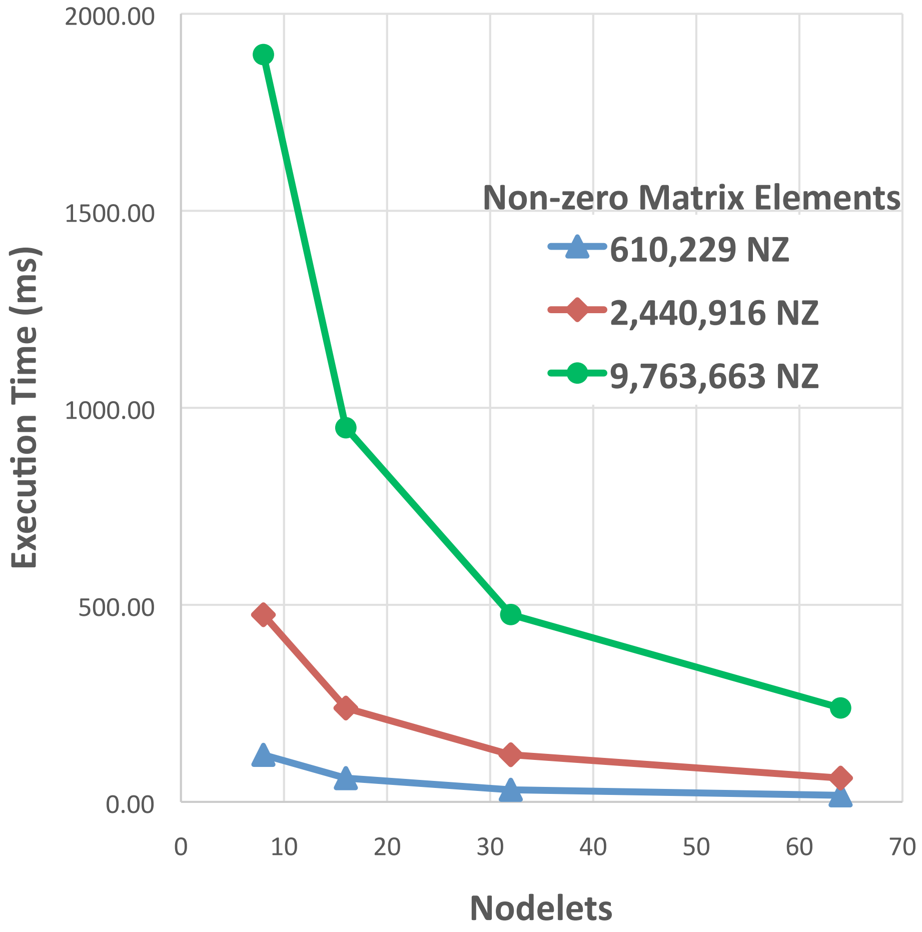


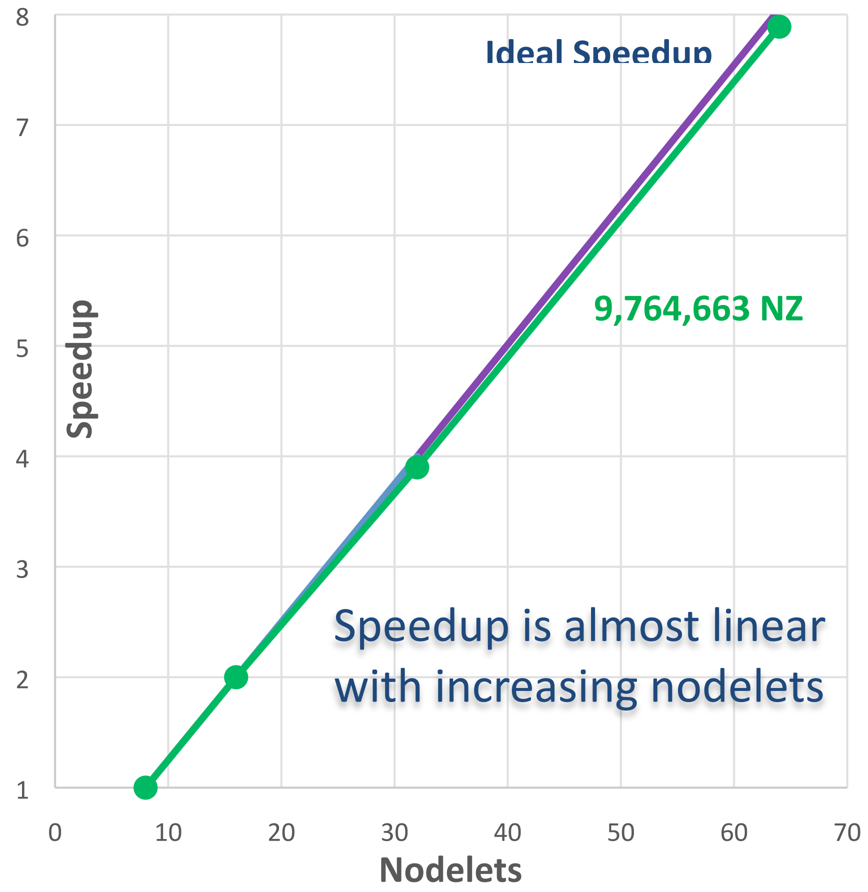




EMU Technology | EMU PROPRIETARY | COPYRIGHT 2018

# Conventional Von Neumann architecture

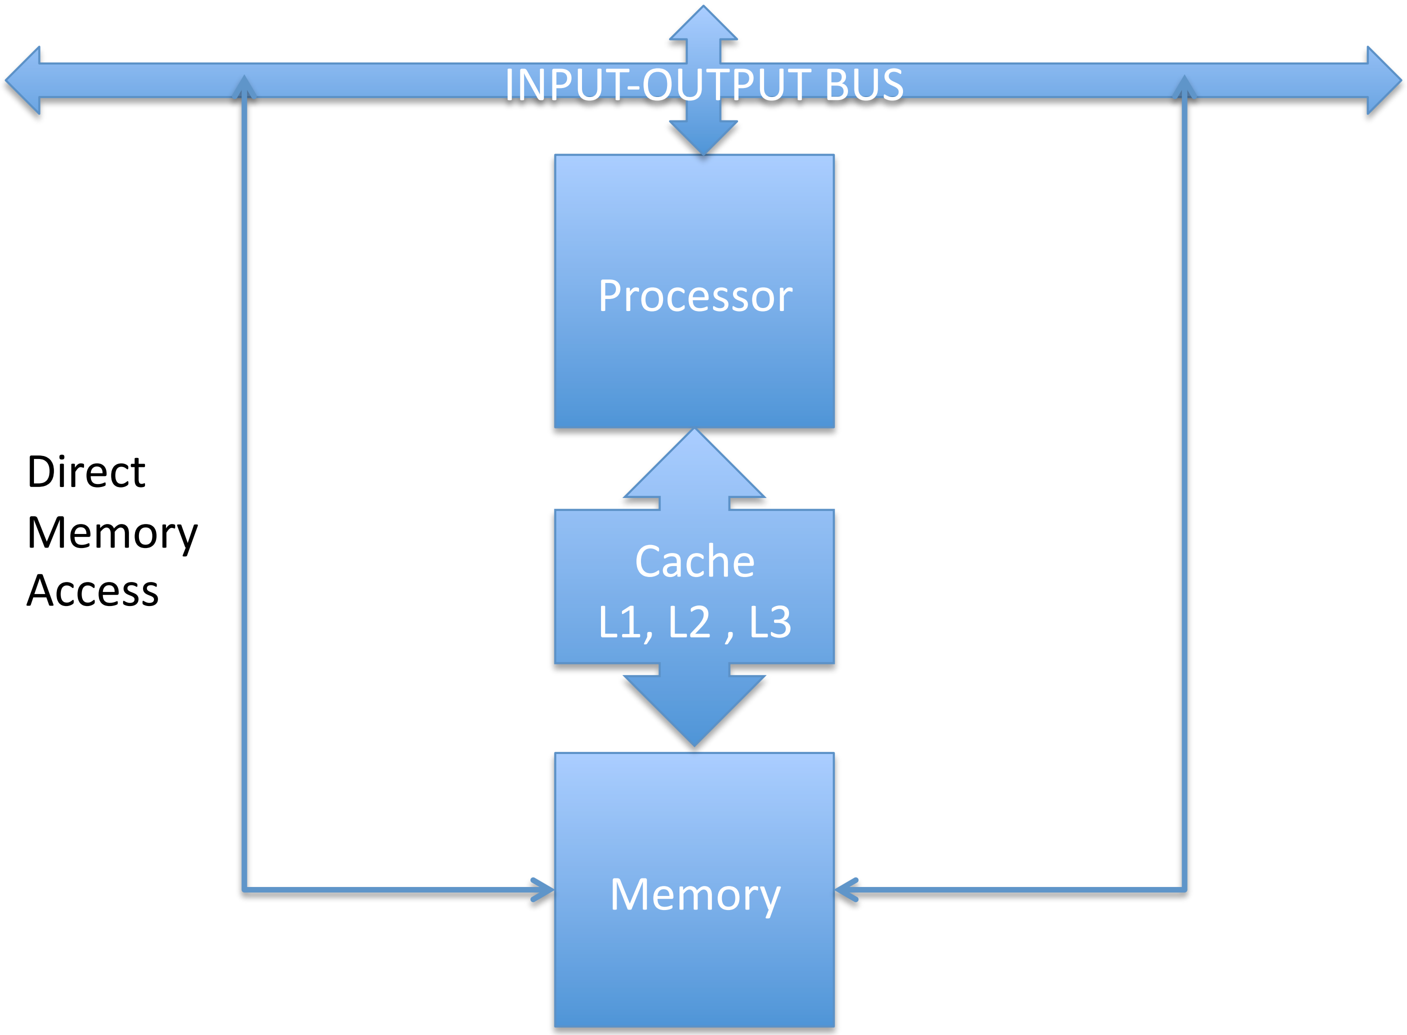

# Extended Von Neumann architectures

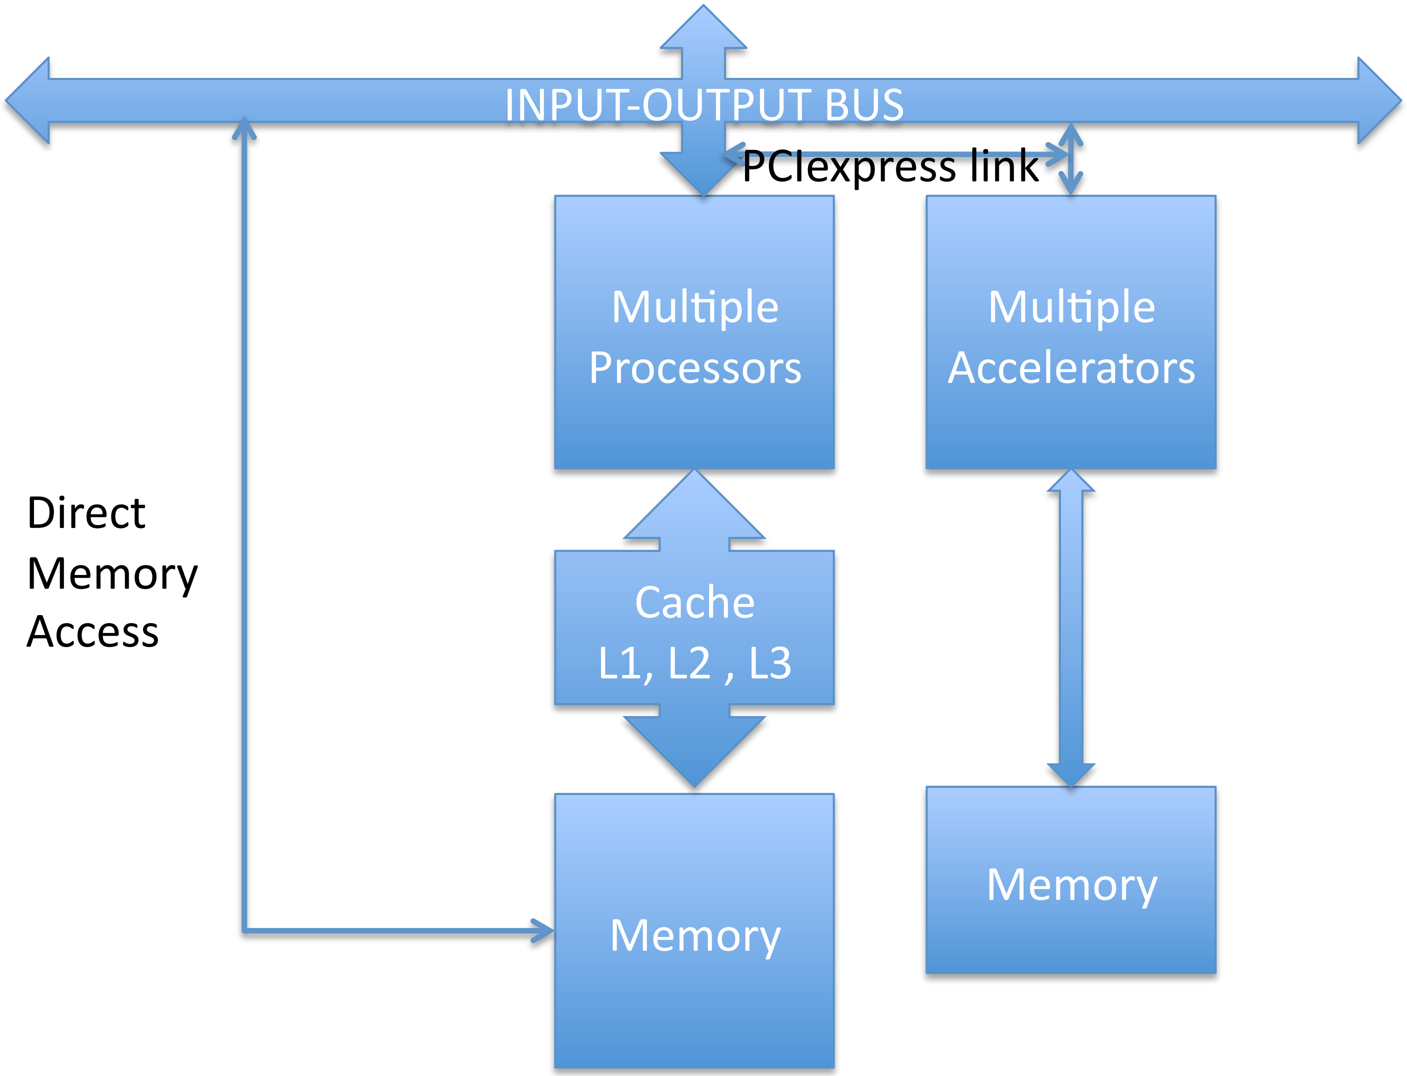

# Where von Neumann architectures burn energy

| Operation | Energy (pJ) |
|---|---|
| 64-bit integer operation | 1 |
| 64-bit floating-point operation | 20 |
| 256 bit on-die SRAM access | 50 |
| 256 bit bus transfer (short) | 26 |
| 256 bit bus transfer (1/2 die) | 256 |
| Off-die link (efficient) | 500 |
| 256 bit bus transfer(across die) | 1,000 |
| DRAM read/write (512 bits) | 16,000 |
| HDD read/write | $\mathcal{O}(10^6)$ |

28nm CMOS, DDR3

*Chart by Greg Astfalk, HPE*

# Step 6: Deploy silicon-photonics

- Silicon based lasers
- 1200nm ultra-high-speed connectivity between chips
- electrically-pumped lasers grown directly onto Si substrate
- deploying optoelectronics: light-driven transistor switching

(this allows for petahertz vs. terahertz transistor switching)

$10^{15}$ $10^{12}$

# Step 7: Deploy 4D VR & AR visualization

- VR & AR have become commodity priced
- Excellent way to diffuse new knowledge
- Interactive & immersive, 3D + time
- Causation not just correlation
- Navigable time-evolution of physical processes
- Planetarium-auditorium-laptop-tablet-cell phone access
- The ultimate game, entertainment and knowledge centre
- Accelerate the speed to insight in research and monitoring

# Step 8: Global-Local downscaling to city level

- ICES-TEST partnership
- The Ecological Sequestration Trust (UK)
- ***Resilience.io*** simulator for hyperlocal
- Capturing city level human activity
- Sanitation systems, water security
- Food security, energy security
- Transportation, Economy
- Public health & safety
- Quality of life

# Summary

- Global alert system for critical infrastructure before it's too late
- Full-on integration, massive paradigm shift & code rewrite
- International collaboration and teamwork
- Fully supportive to the 17 UN SDGs
- Requiring a full decade to accomplish: 2020-2030
- From global to local, from soil to sky, from space to face
- Open science, source code & data; open access publishing
- We need to surmount geopolitical barriers to achieve this vision!

ICES

FOUNDATION