Parallel Computing
&
Accelerators

John Urbanic
Pittsburgh Supercomputing Center
Parallel Computing Scientist

Purpose of this talk

This is the 50,000 ft. view of the parallel computing landscape. We want to orient you a bit before parachuting you down into the trenches to deal with OpenACC. The plan is that you walk away with a knowledge of not just OpenACC, but also where it fits into the world of High Performance Computing.

FLOPS we need: Climate change analysis

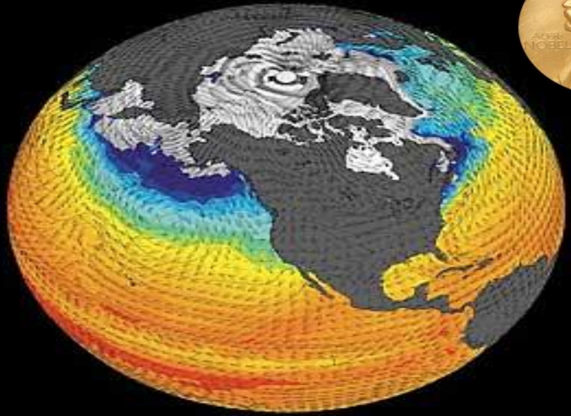

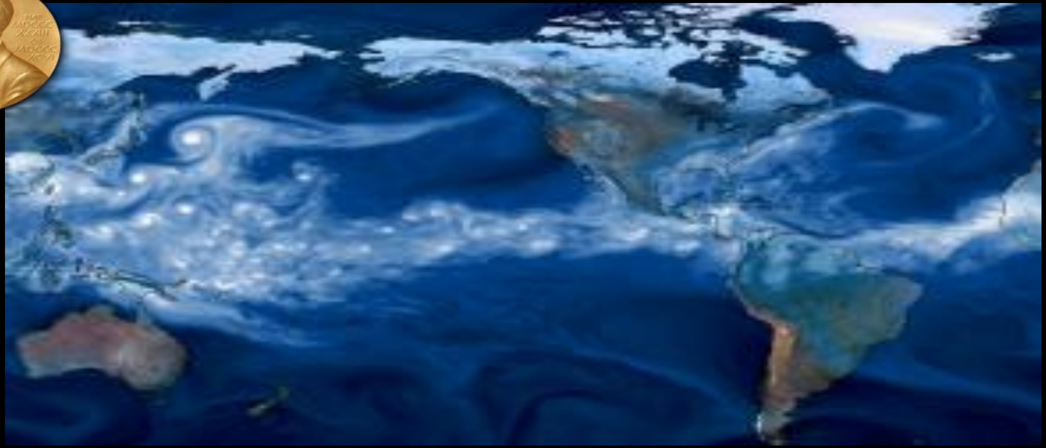

Simulations

- Cloud resolution, quantifying uncertainty, understanding tipping points, etc., will drive climate to exascale platforms
- New math, models, and systems support will be needed

Extreme data

- “Reanalysis” projects need $100\times$ more computing to analyze observations
- Machine learning and other analytics are needed today for petabyte data sets
- Combined simulation/observation will empower policy makers and scientists

Courtesy Horst Simon, LBNL

Exascale combustion simulations

- Goal: 50% improvement in engine efficiency
- Center for Exascale Simulation of Combustion in Turbulence (ExaCT)
  - Combines M&S and experimentation
  - Uses new algorithms, programming models, and computer science

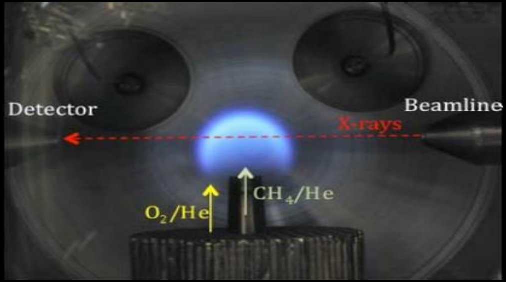

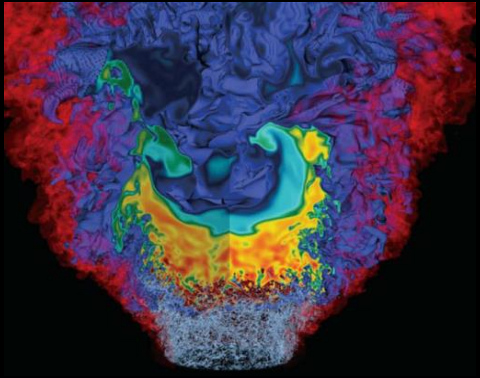

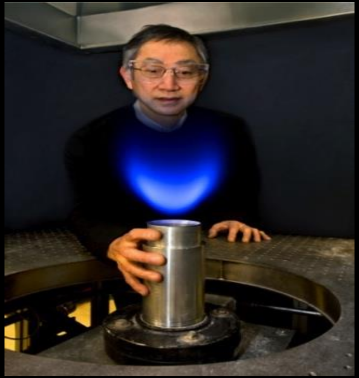

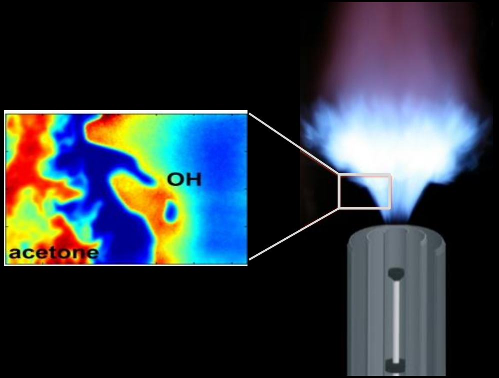

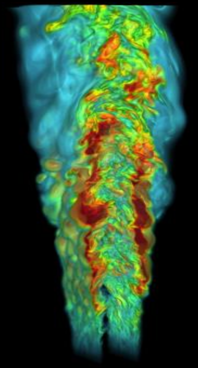

Courtesy Horst Simon, LBNL

The list is long, and growing.

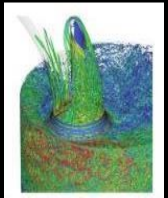

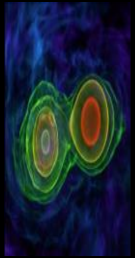

- Molecular-scale Processes: atmospheric aerosol simulations
- AI-Enhanced Science: predicting disruptions in tokomak fusion reactors

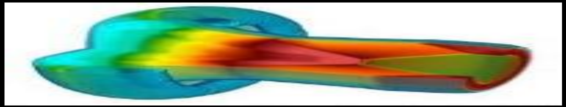

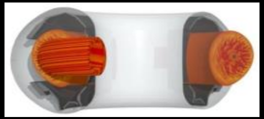

- Hypersonic Flight
- Modeling Thermonuclear X-ray Bursts: 3D simulations of a neutron star surface or supernovae

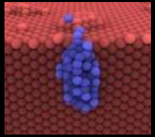

- Quantum Materials Engineering: electrical conductivity photovoltaic and plasmonic devices

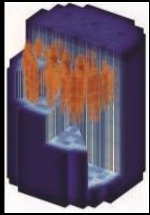

- Physics of Fundamental Particles: mass estimates of the bottom quark
- Digital Cells

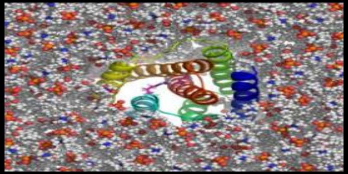

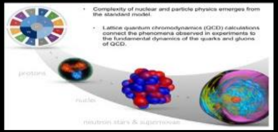

And many of you doubtless brought your own immediate research concerns. Great!

Welcome to The Exascale Era!

exa = 10^{18} = 1,000,000,000,000,000,000 = quintillion

64-bit precision floating point operations per second

23,800
Cray Red Storms
2004 (42 Tflops)

133,33
NVIDIA V100
(7.5 Tflops)

There may also be a Chinese machine, OceanLight, or 3-letter-agency machines on the scene.

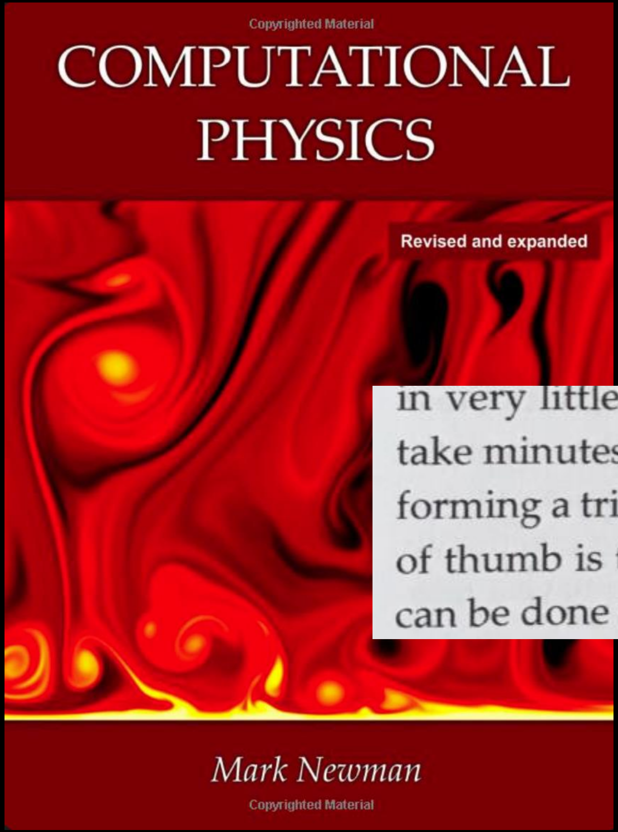

in very little time. Performing a billion operations, on the other hand, could take minutes or hours, though it's still possible provided you are patient. Performing a trillion operations, however, will basically take forever. So a fair rule of thumb is that the calculations we can perform on a computer are ones that can be done with *about a billion operations or less.*

Where are those 10 or 12 orders of magnitude?

How do we get there from here?

BTW, that's a bigger gap than

vs.

IBM 709
12 kiloflops

Moore's Law abandoned serial programming around 2004

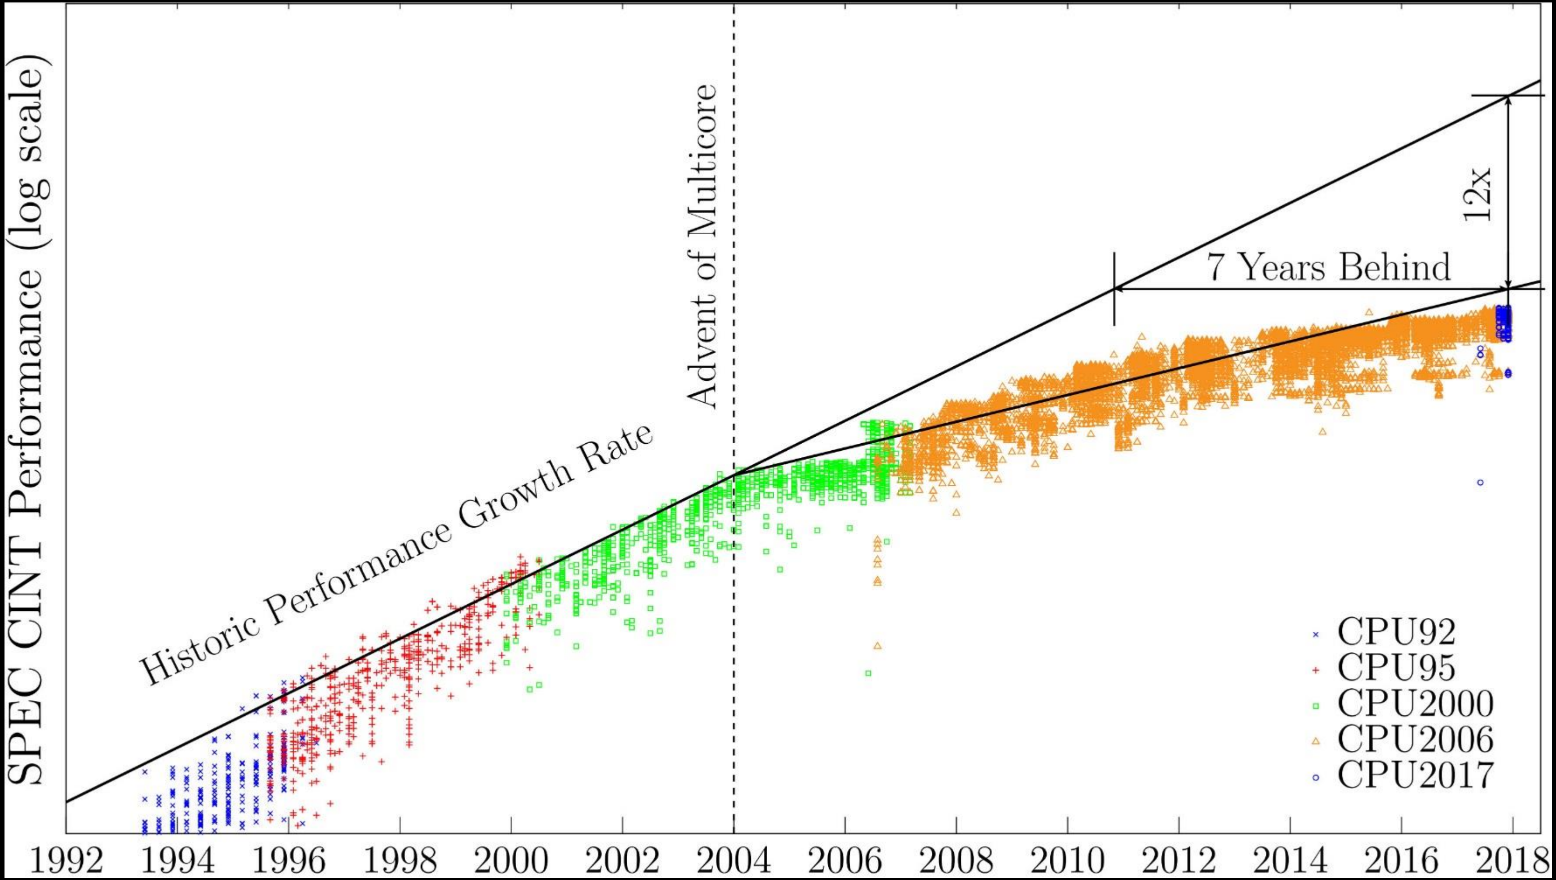

Courtesy Liberty Computer Architecture Research Group

But Moore's Law is only beginning to stumble now.

Intel process technology capabilities

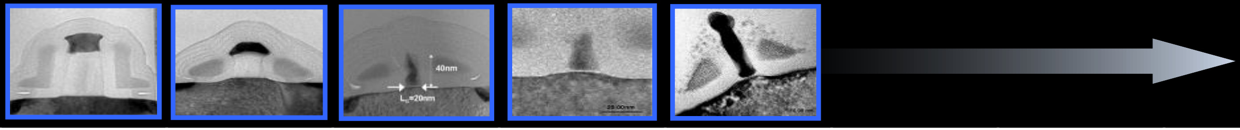

High Volume Manufacturing	2004	2006	2008	2010	2012	2014	2018	2021
Feature Size	90nm	65nm	45nm	32nm	22nm	14nm	10nm	7nm
Integration Capacity (Billions of Transistors)	2	4	8	16	32	64	128	256

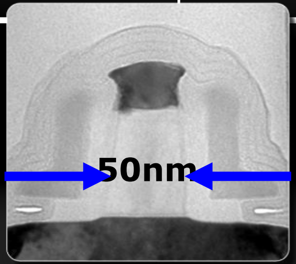

Transistor for 90nm Process

Source: Intel

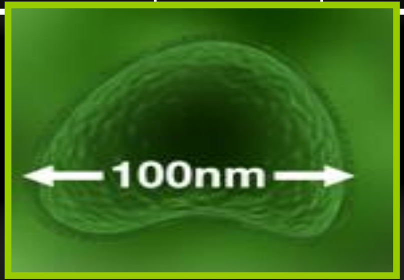

Influenza Virus

Source: CDC

And at end of day we keep using getting more transistors.

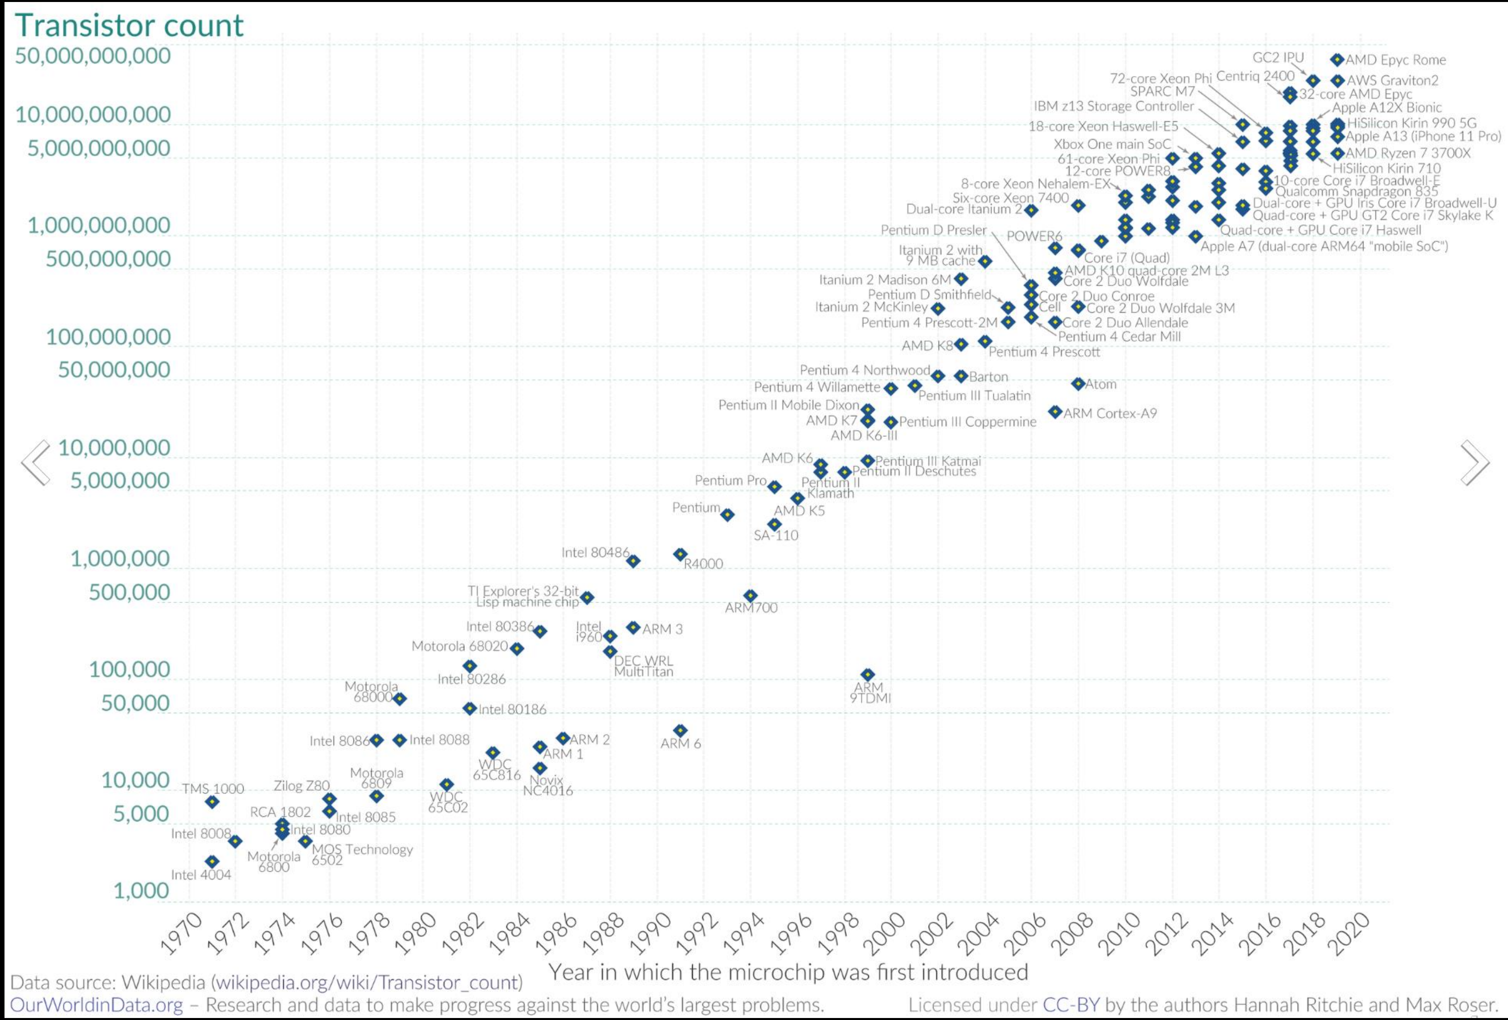

That Power and Clock Inflection Point in 2004… didn't get better.

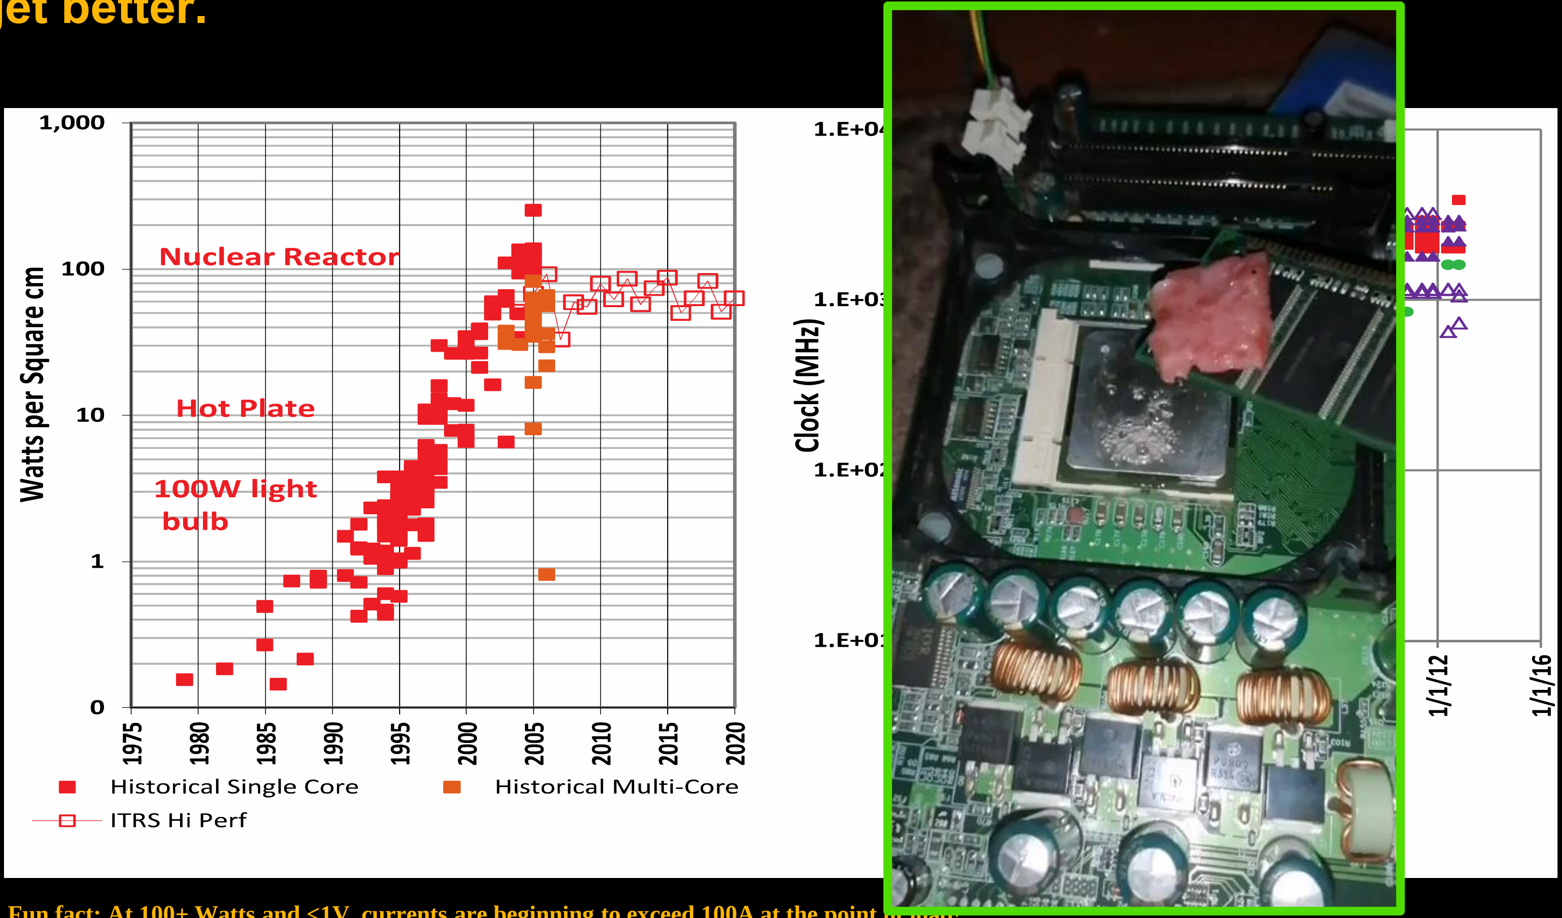

Fun fact: At 100+ Watts and <1V, currents are beginning to exceed 100A at the point of load.

Courtesy Horst Simon, LBNL

Not a new problem, just a new scale…

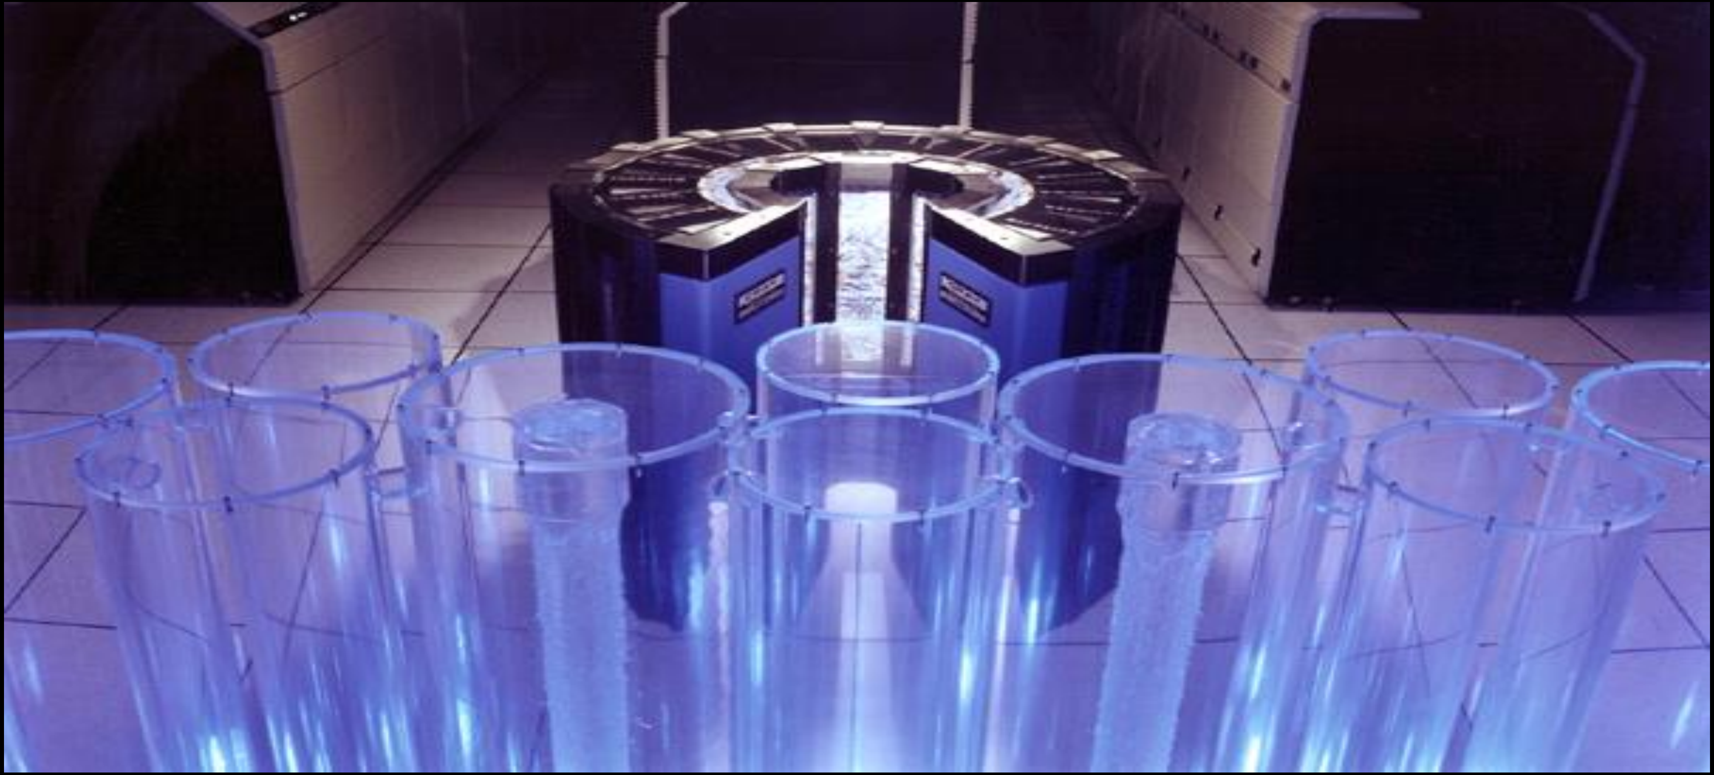

Cray-2 with cooling tower in foreground, circa 1985

And how to get more performance from more transistors with the same power.

A 15% Reduction In Voltage Yields

RULE OF THUMB

Frequency Reduction	Power Reduction	Performance Reduction
15%	45%	10%

SINGLE CORE

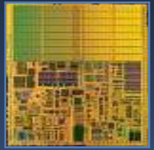

- Area = 1
- Voltage = 1
- Freq = 1
- Power = 1
- Perf = 1

DUAL CORE

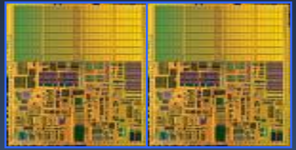

- Area = 2
- Voltage = 0.85
- Freq = 0.85
- Power = 1
- Perf = ~1.8

Parallel Computing

One woman can make a baby in 9 months.

Can 9 women make a baby in 1 month?

But 9 women can make 9 babies in 9 months.

First two bullets are Brook's Law. From *The Mythical Man-Month.*

A must-read for serious project programmers that includes many other classics such as:
"What one programmer can do in one month, two programmers can do in two months."

Prototypical Application: Serial Weather Model

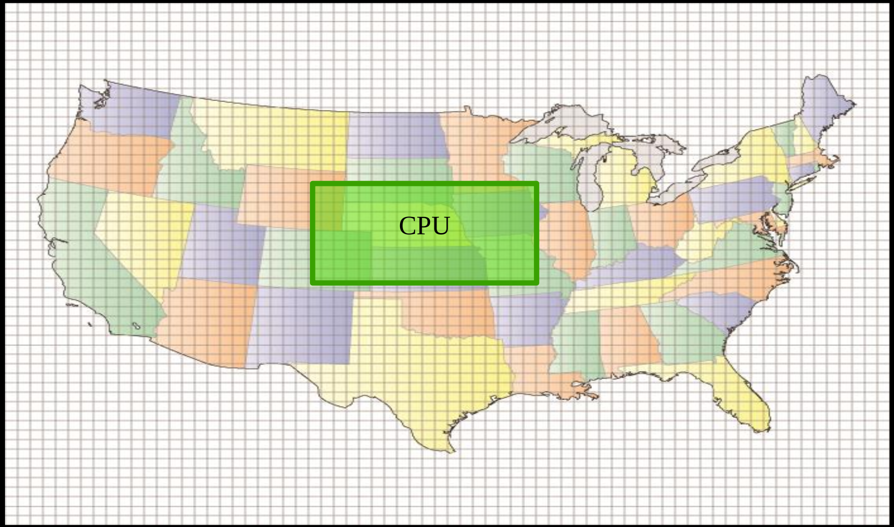

First Parallel Weather Modeling Algorithm: Richardson in 1917

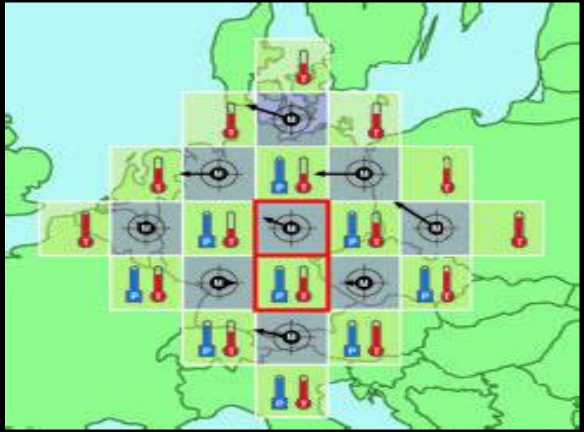

Courtesy John Burkhardt, Virginia Tech

Weather Model: Shared Memory (OpenMP)

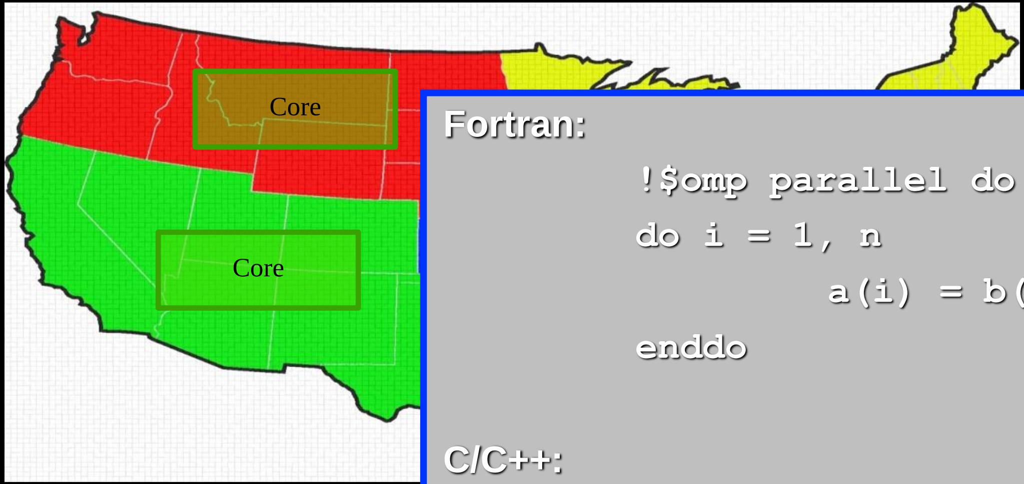

Fortran:

```
!$omp parallel do
do i = 1, n
        a(i) = b(i) + c(i)
enddo
```

C/C++:

```
#pragma omp parallel for
for(i=1; i<=n; i++)
          a[i] = b[i] + c[i];
```

Four meteorologists in th

Weather Model: Distributed Memory (MPI)

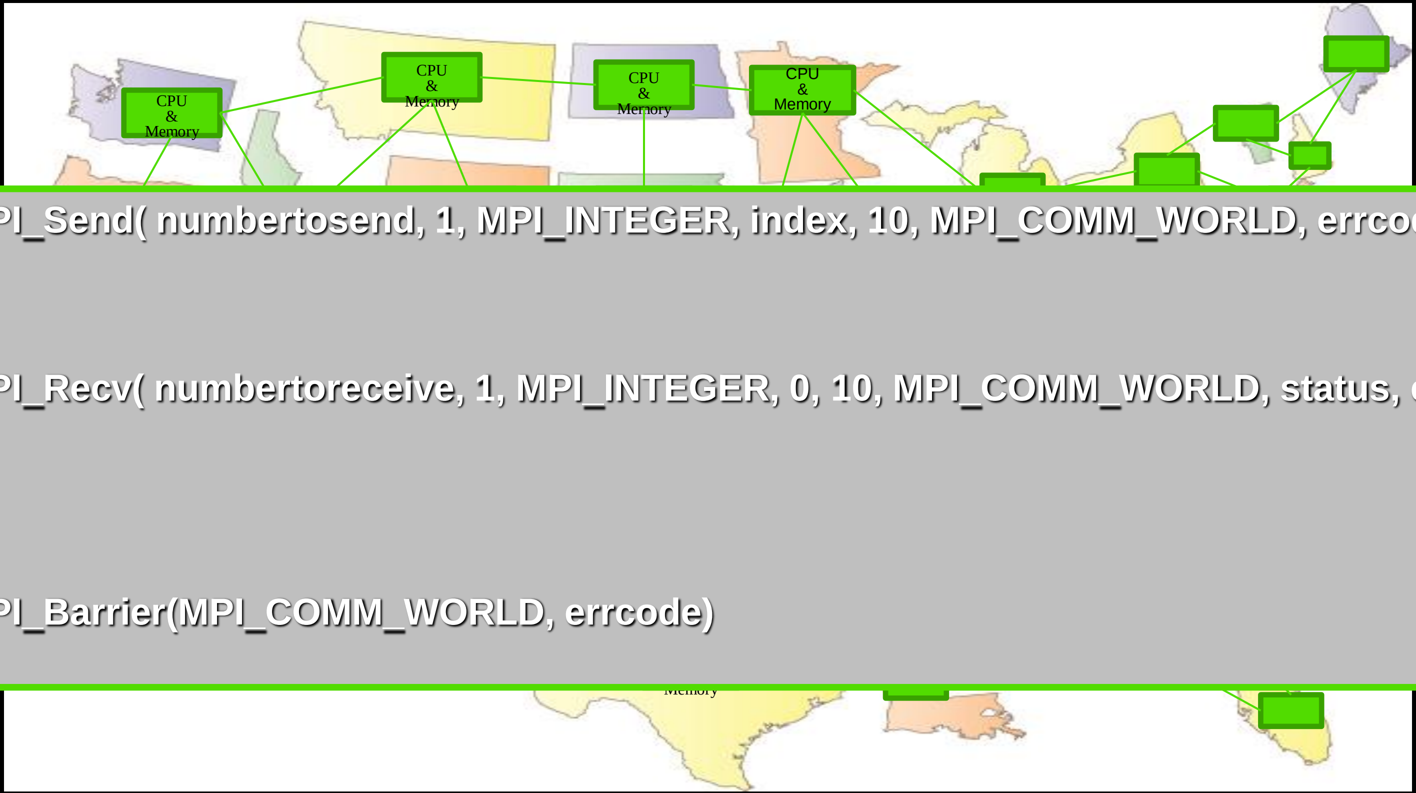


```
call MPI_Send( numbertosend, 1, MPI_INTEGER, index, 10, MPI_COMM_WORLD, errcode)
.
.
call MPI_Recv( numbertoreceive, 1, MPI_INTEGER, 0, 10, MPI_COMM_WORLD, status, errcode)
.
.
.
call MPI_Barrier(MPI_COMM_WORLD, errcode)
.
```

50 meteorologists using a telegraph.

Weather Model: Accelerator (OpenACC)

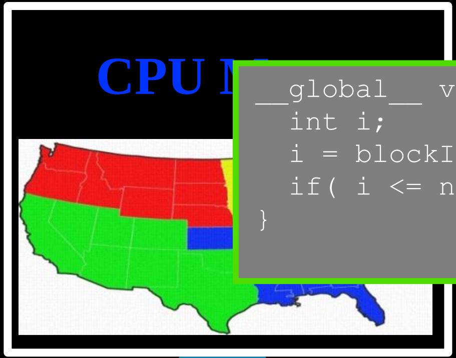

```
__global__ void saxpy_kernel( float a, float* x, float* y, int n ){
  int i;
  i = blockIdx.x*blockDim.x + threadIdx.x;
  if( i <= n ) x[i] = a*x[i] + y[i];
}
```

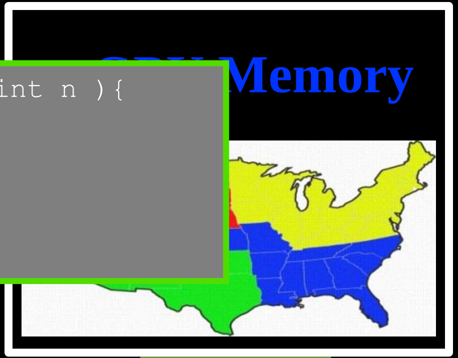

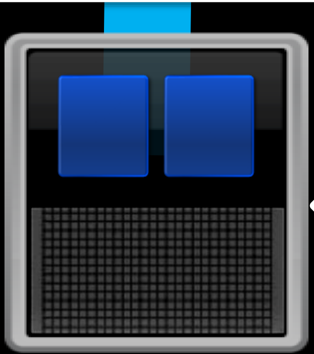

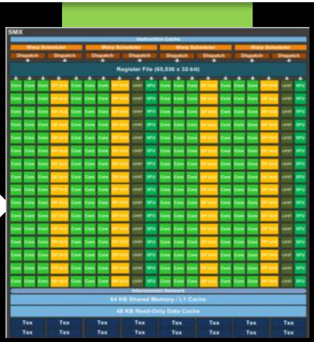

CPU

GPU

1 meteorologists coordinating 1000 math savants using tin cans and a string.

The pieces fit like this…

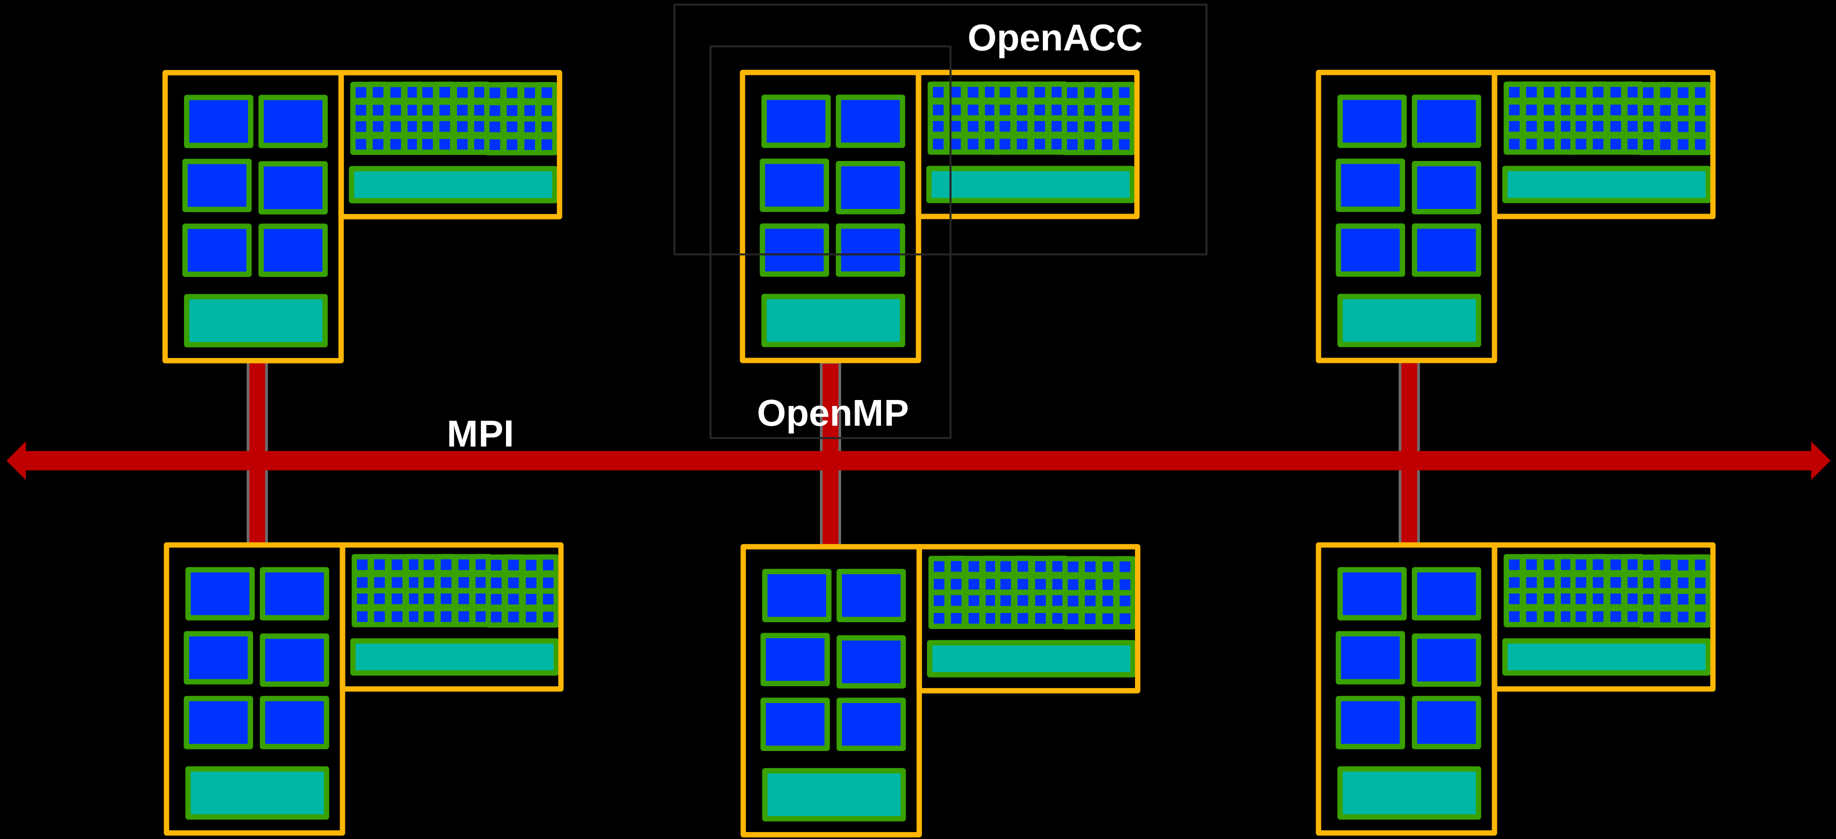

Top 10 Systems as of June 2024

#	Computer	Site	Manufacturer	CPU Interconnect [*Accelerator*]	Cores	Rmax (Pflops)	Rpeak (Pflops)	Power (MW)
1	**Frontier**	Oak Ridge National Laboratory **United States**	HPE	AMD EPYC 64C 2GHz Slingshot-11 AMD Instinct MI250X	**8,699,904**	**1194**	**1692**	**22.7**
2	**Aurora**	Argonne National Laboratory **United States**	HPE	Intel Xeon Max 9470 52C 2.4GHz Slingshot-11 Intel Data Center GPU Max	**4,742,808**	**585**	**1059**	**24.6**
3	**Eagle**	Microsoft **United States**	Microsoft	Intel Xeon 8480C 48C 2GHz Infiniband NDR NVIDIA H100	**1,123,200**	**561**	**846**	
4	**Fugaku**	RIKEN Center for Computational Science **Japan**	Fujitsu	ARM 8.2A+ 48C 2.2GHz Torus Fusion Interconnect	**7,630,072**	**442**	**537**	**29.9**
5	**LUMI**	EuroHPC **Finland**	HPE	AMD EPYC 64C 2GHz Slingshot-11 AMD Instinct MI250X	**2,752,704**	**379**	**531**	**7.1**
6	**Leonardo**	EuroHPC **Italy**	Atos	Intel Xeon 8358 32C 2.6GHz Infiniband HDR NVIDIA A100	**1,824,768**	**238**	**304**	**7.4**
7	**Summit**	Oak Ridge National Laboratory **United States**	IBM	Power9 22C 3.0 GHz Dual-rail Infiniband EDR *NVIDIA V100*	**2,414,592**	**148**	**200**	**10.1**
8	**MareNostrum 5**	EuroHPC/BSC **Spain**	EVIDEN	*Intel Xeon 8460Y+ 40C 2.3GHz,* *Infiniband NDR200* *NVIDIA H100 64GB*	**680,960**	**138**	**256**	**2.5**
9	**Eos NVIDIA DGX SuperPOD**	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	**188**	
10	**Sierra**	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	**125**	**7.4**

500	**TX-Green2** - PowerEdge C6420, Xeon Platinum 8260 24C 2.4GHz, 25G Ethernet, ACTION MIT Lincoln Laboratory Supercomputing Center United States	43,200	2.02	53.08

The word is *Heterogeneous*

And it's not just supercomputers. It's on your desk, and in your phone.

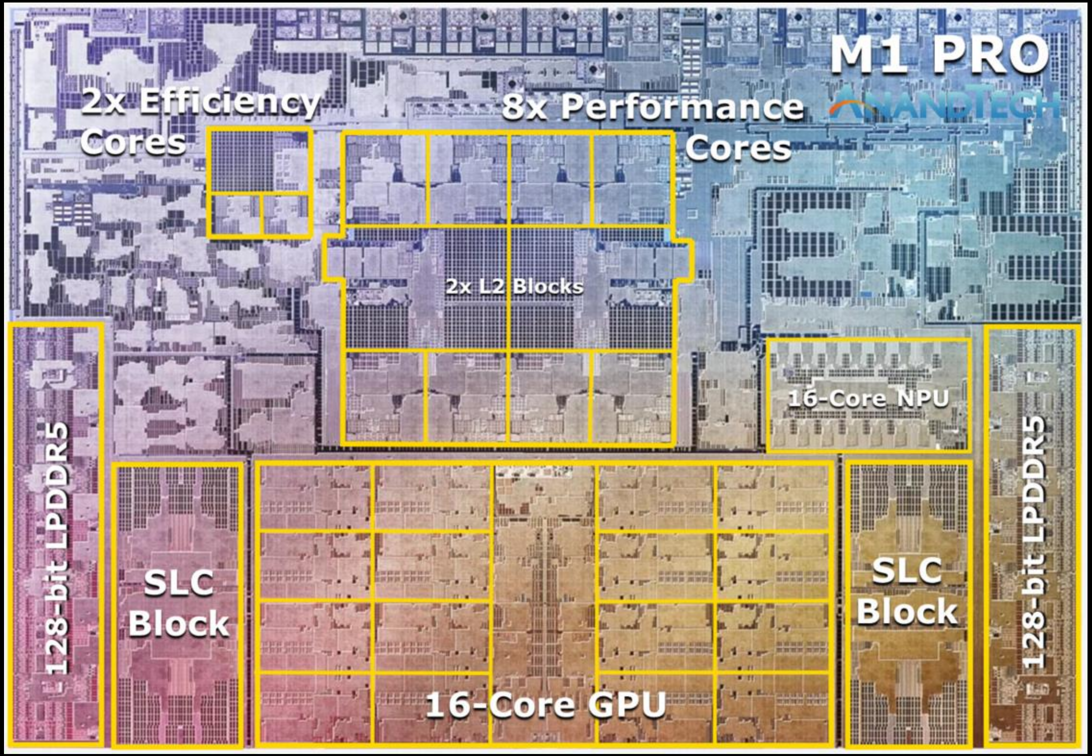

AnandTech

How much of this can you program?

We can do better. We have a role model.

- We hope to "simulate" a human brain in real time on one of these Exascale platforms with about 1 - 10 Exaflop/s and 4 PB of memory
- These newest Exascale computers use 20+ MW
- The human brain runs at 20W
- Our brain is a million times more power efficient!

Why you should be (extra) motivated.

- This parallel computing thing is no fad.
- The laws of physics are drawing this roadmap.
- If you get on board (the right bus), you can ride this trend for a long, exciting trip.

Let’s learn how to use these things!