It's a Multicore World

John Urbanic
Parallel Computing Scientist
Pittsburgh Supercomputing Center

Moore's Law abandoned serial programming around 2004

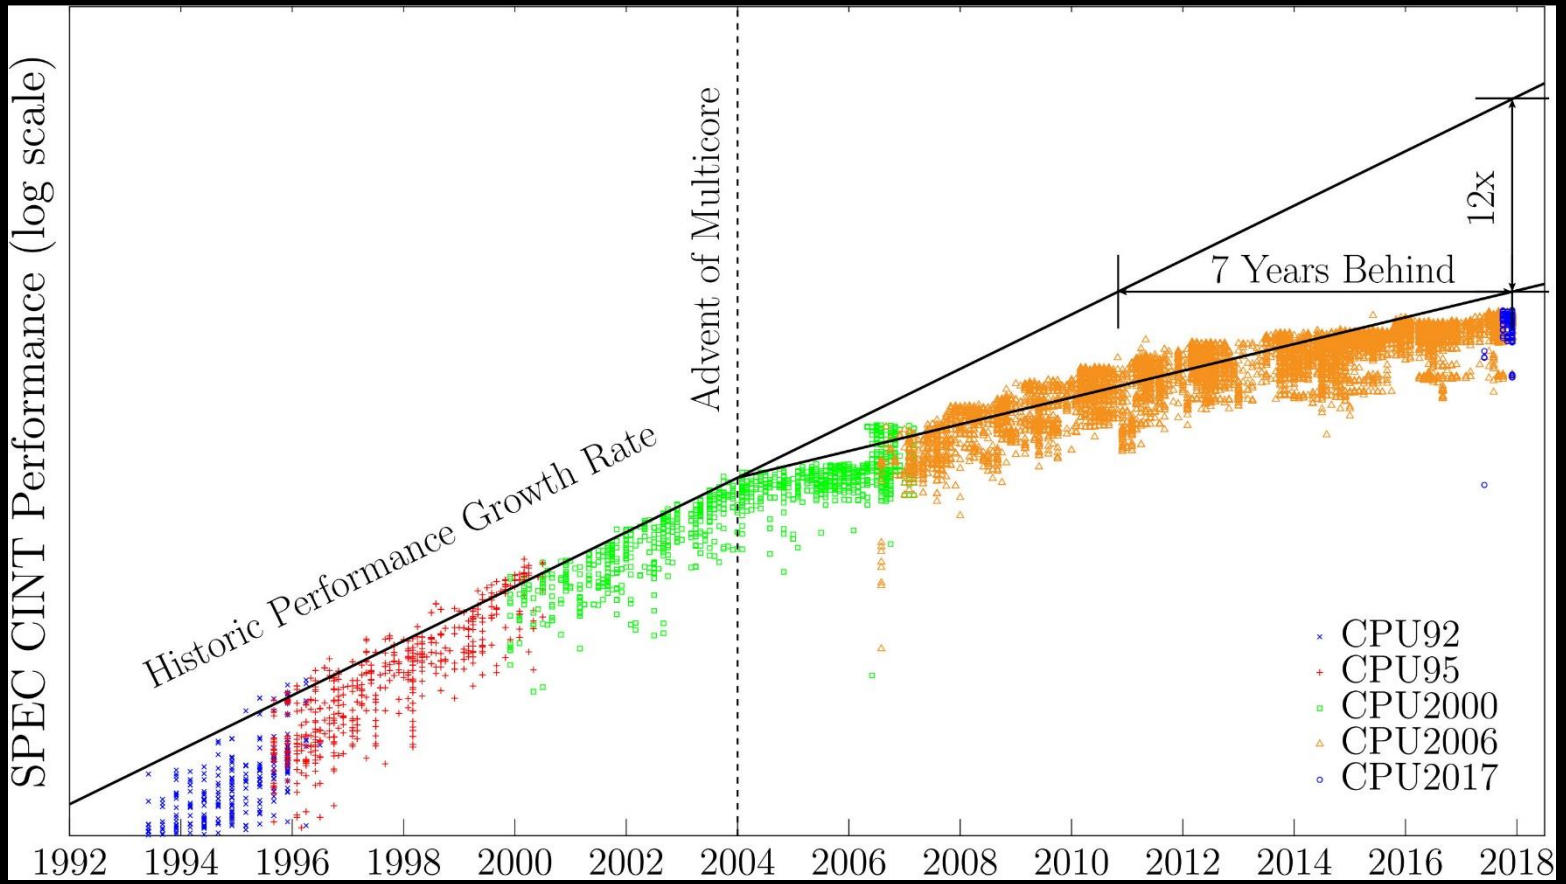

Courtesy Liberty Computer Architecture Research Group

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

Intel process technology capabilities

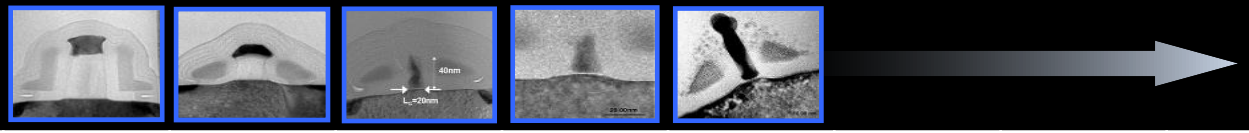

High Volume Manufacturing	2004	2006	2008	2010	2012	2014	2018	2021
Feature Size	90nm	65nm	45nm	32nm	22nm	14nm	10nm	7nm
Integration Capacity (Billions of Transistors)	2	4	8	16	32	64	128	256

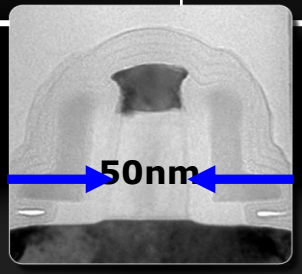

Transistor for 90nm Process

Source: Intel

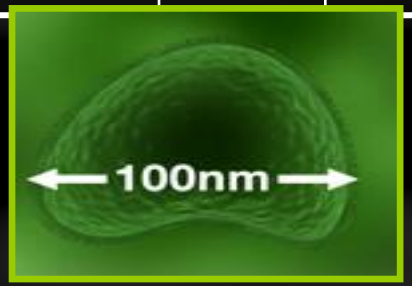

Influenza Virus

Source: CDC

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

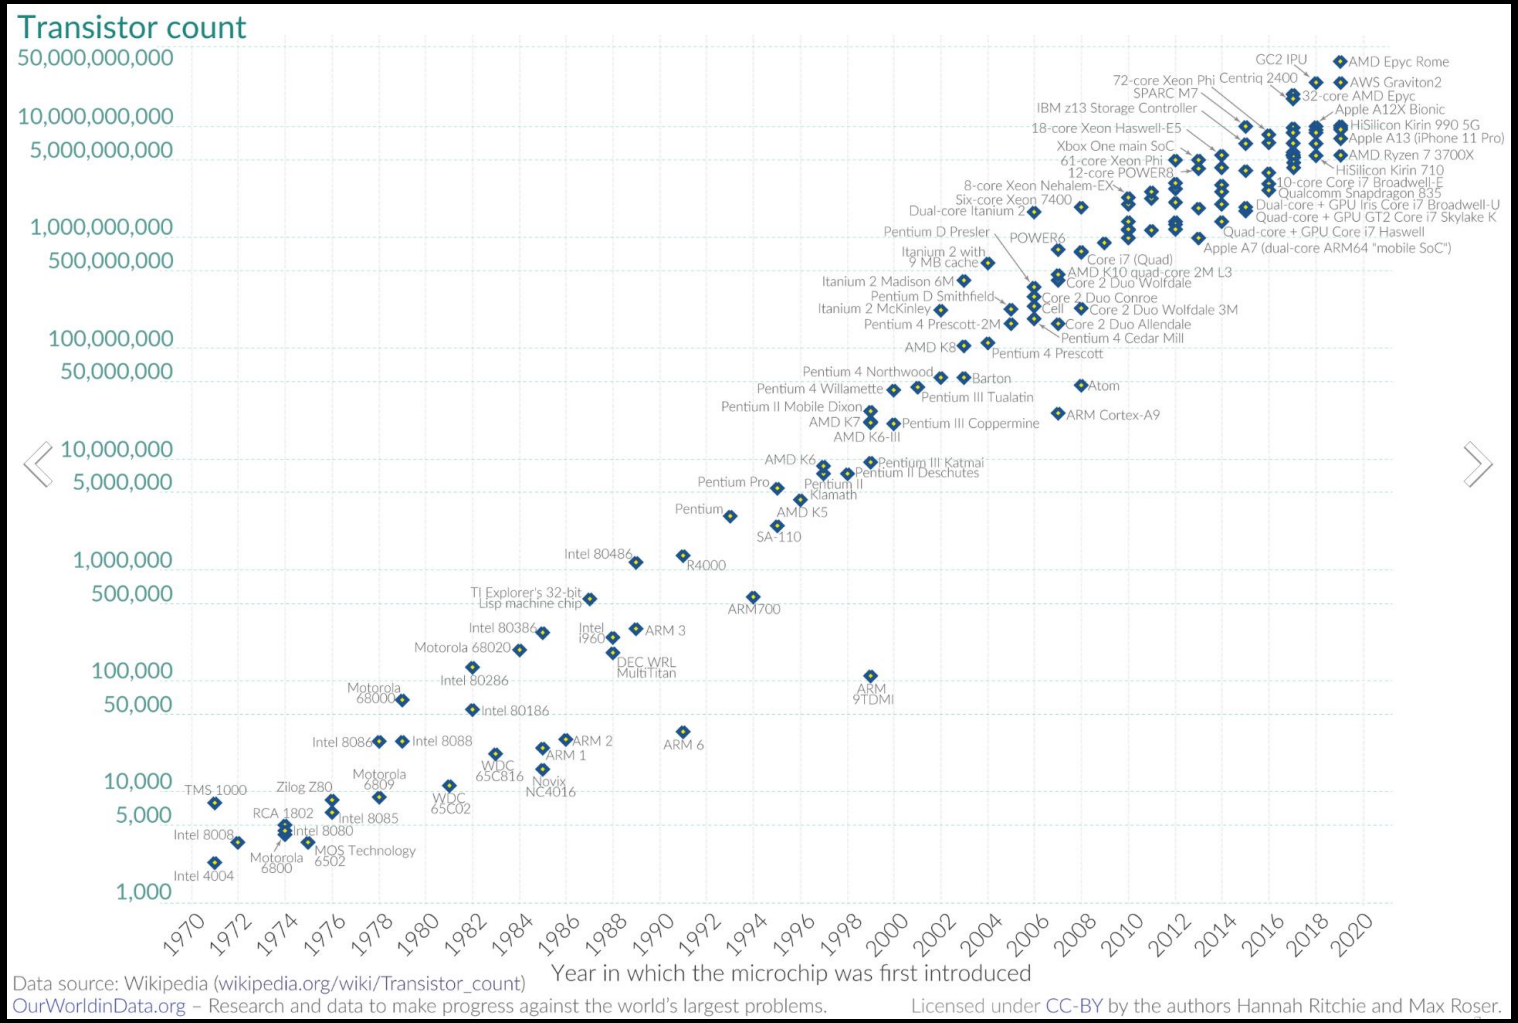

Data source: Wikipedia (wikipedia.org/wiki/Transistor_count)
OurWorldinData.org – Research and data to make progress against the world's largest problems.

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

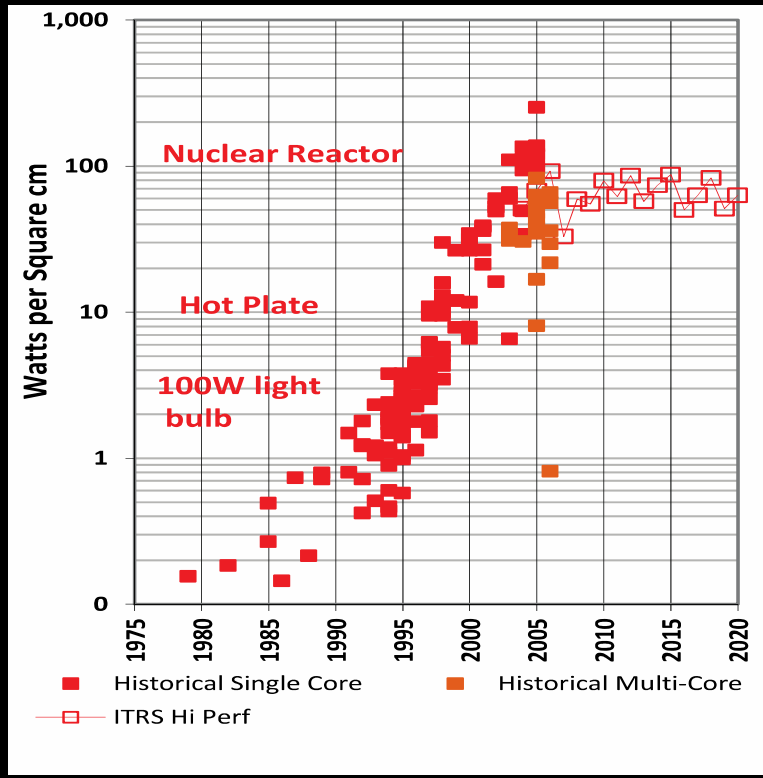

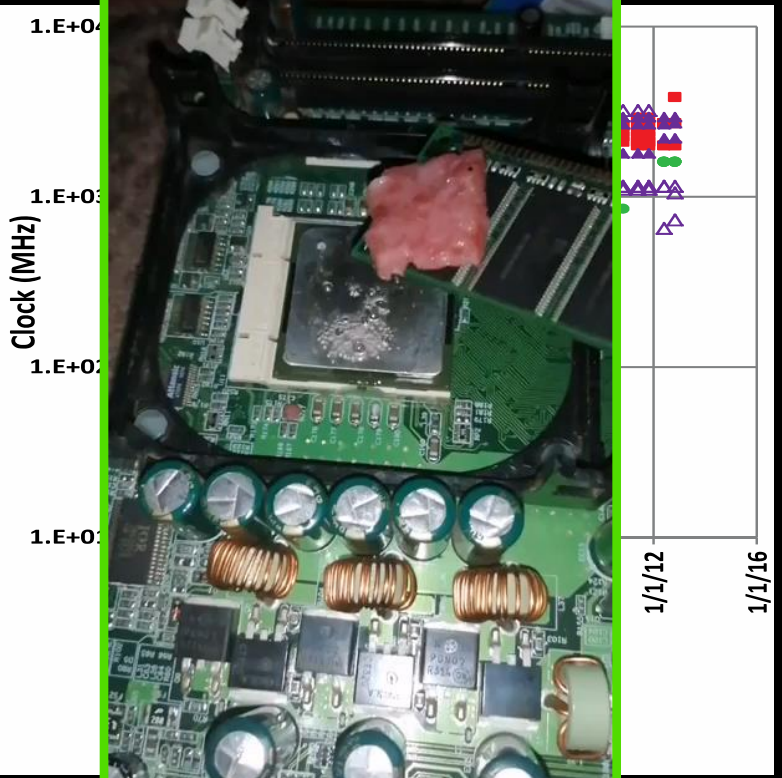

Fun fact: At 100+ Watts and <1V, currents are beginning to exceed 100A at the point of load.

Courtesy Horst Simon, LBNL

Not a new problem, just a new scale…

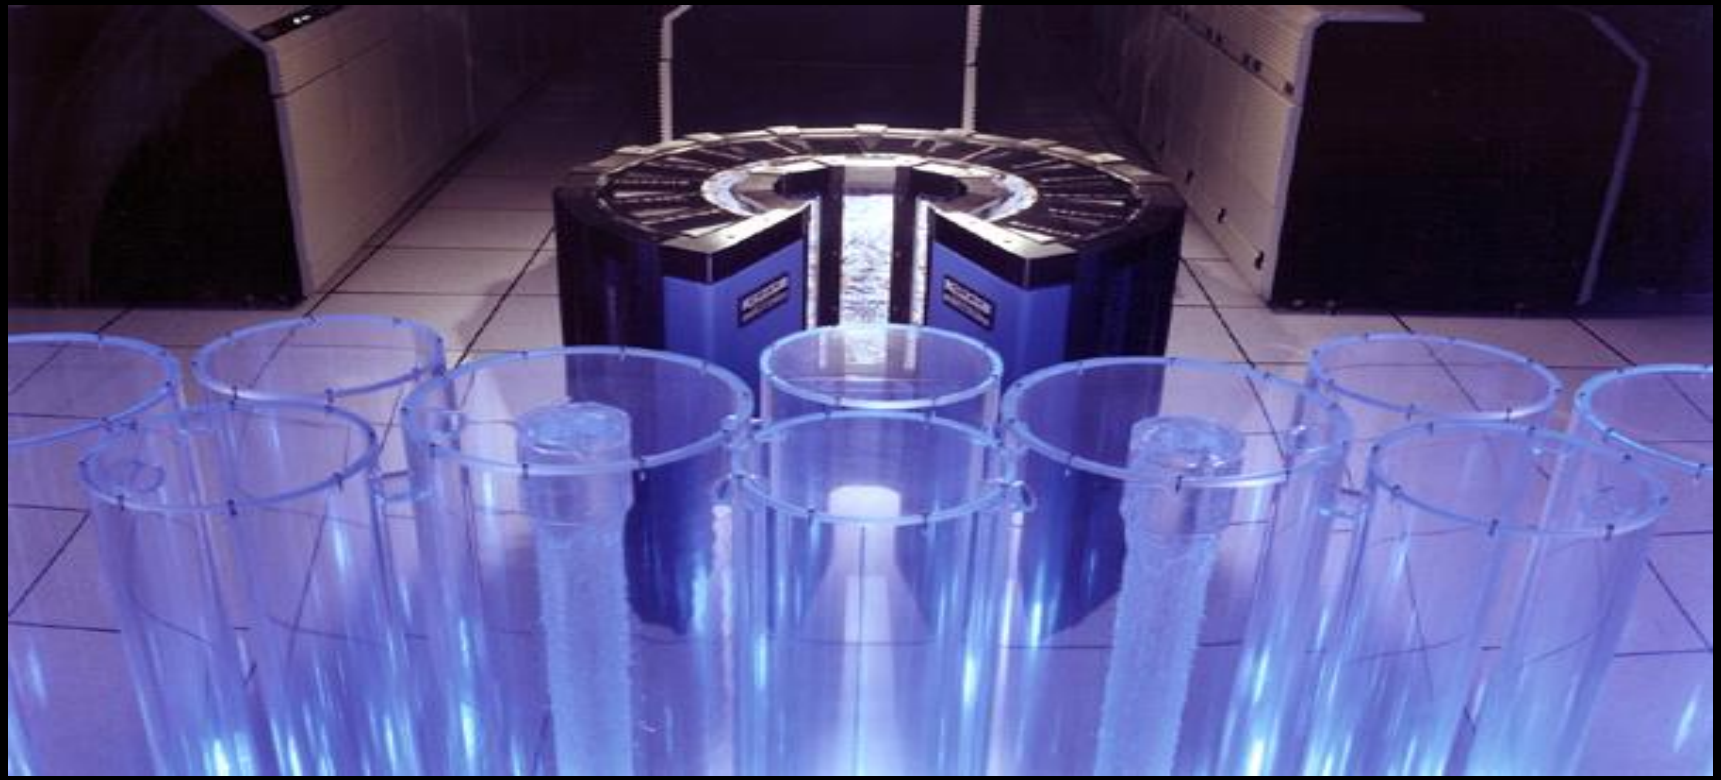

Cray-2 with cooling tower in foreground, circa 1985

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

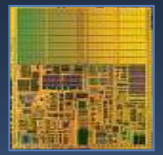

Area = 1
Voltage = 1
Freq = 1
Power = 1
Perf = 1

DUAL CORE

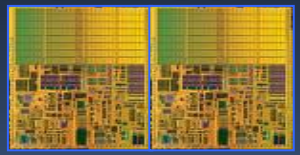

Area = 2
Voltage = 0.85
Freq = 0.85
Power = 1
Perf = ~1.8

Single Socket Parallelism

Processor	Year	Vector	Bits	SP FLOPs / core / cycle	Cores	FLOPs/cycle
Pentium III	1999	SSE	128	3	1	3
Pentium IV	2001	SSE2	128	4	1	4
Core	2006	SSE3	128	8	2	16
Nehalem	2008	SSE4	128	8	10	80
Sandybridge	2011	AVX	256	16	12	192
Haswell	2013	AVX2	256	32	18	576
KNC	2012	AVX512	512	32	64	2048
KNL	2016	AVX512	512	64	72	4608
Skylake	2017	AVX512	512	96	28	2688

Putting It All Together

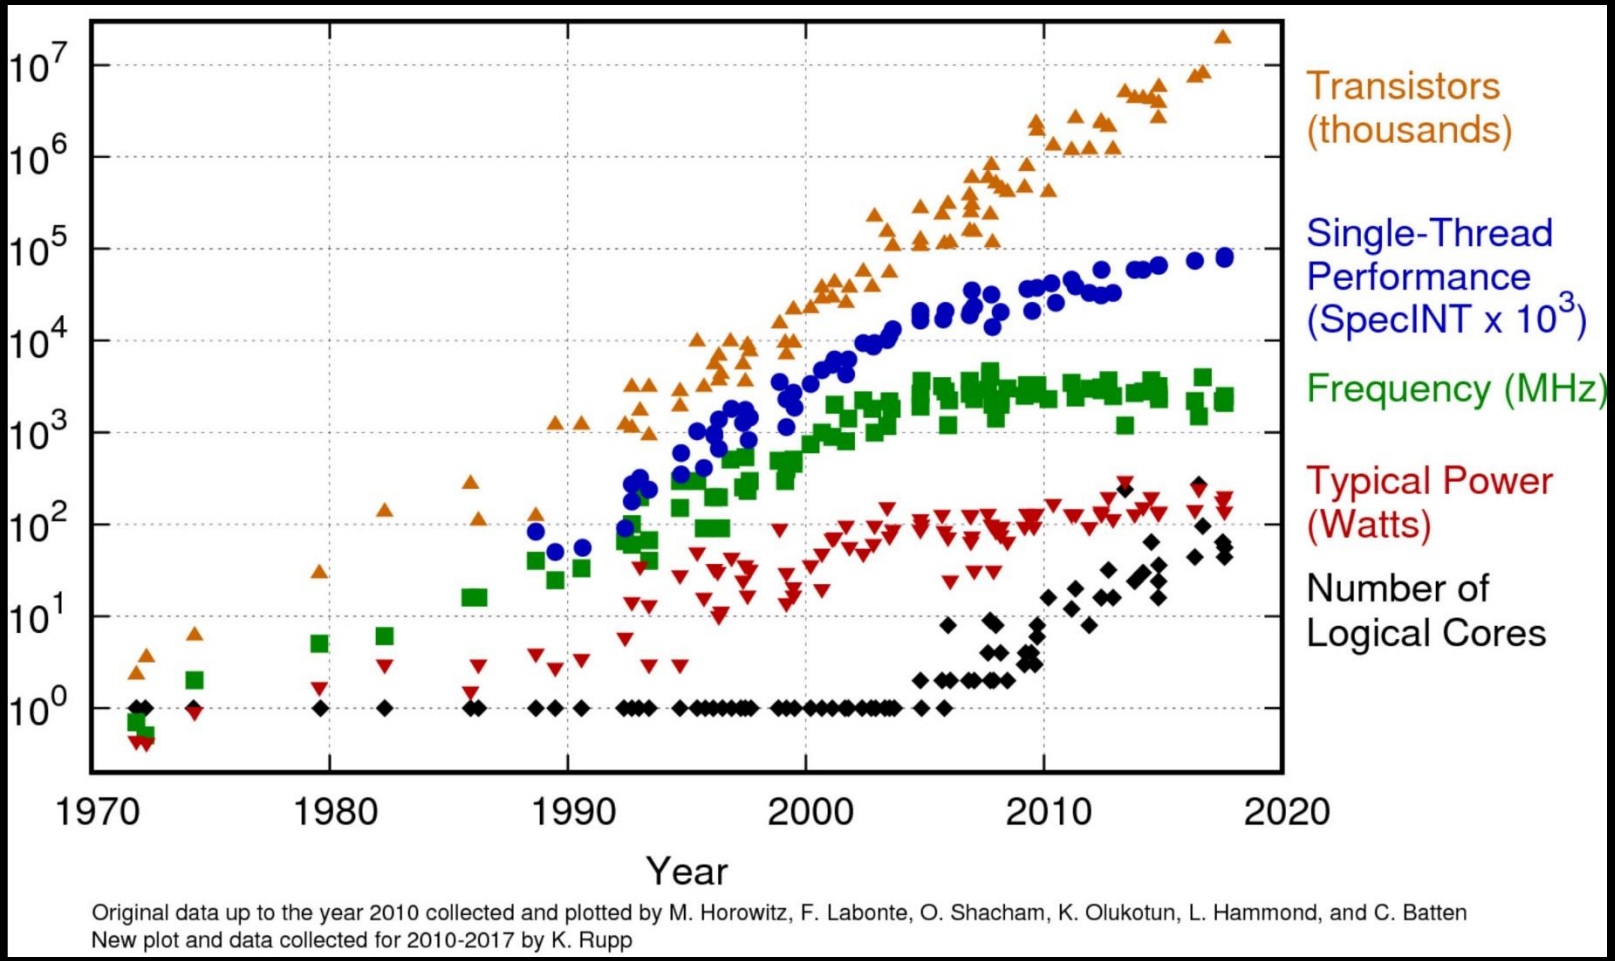

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Prototypical Application: Serial Weather Model

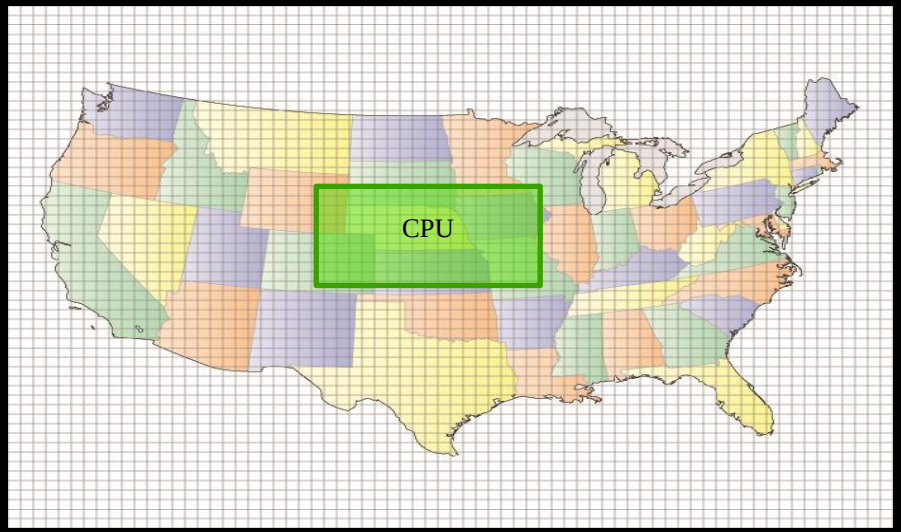

First Parallel Weather Modeling Algorithm: Richardson in 1917

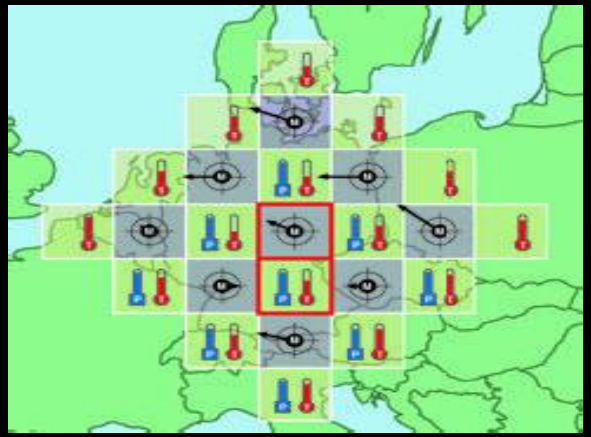

Courtesy John Burkhardt, Virginia Tech

Weather Model: Shared Memory (OpenMP)

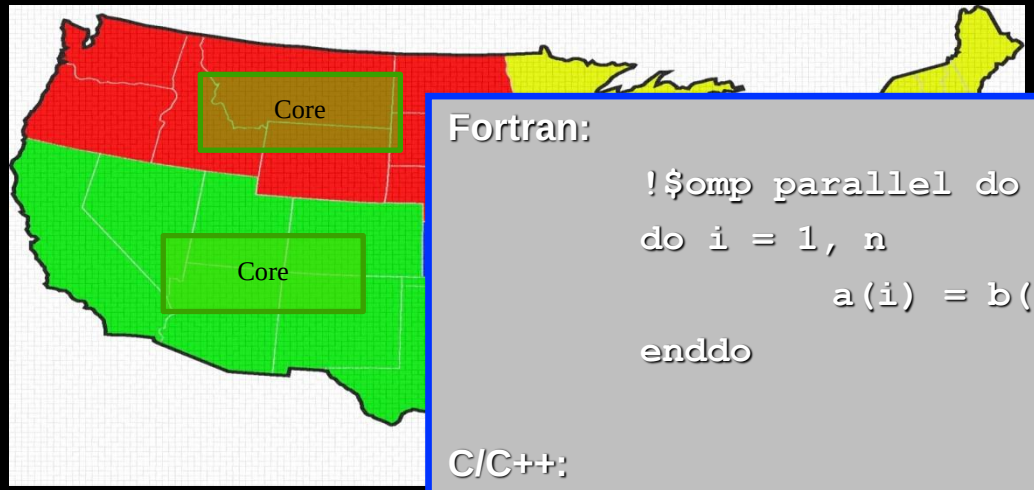

Four meteorologists in t

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];
```

Weather Model: Accelerator (OpenACC)

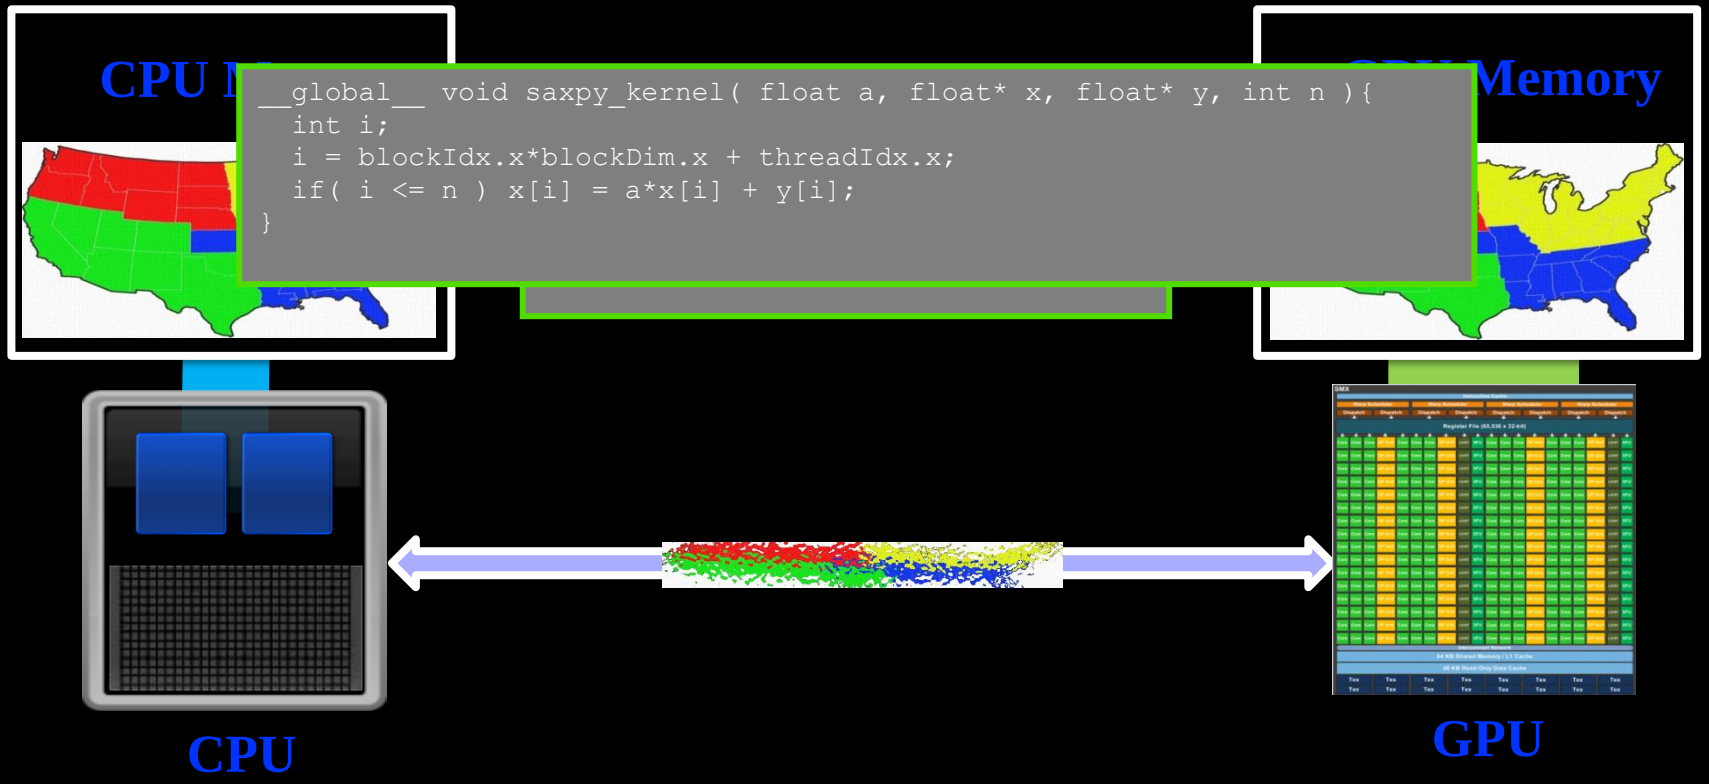

1 meteorologists coordinating 1000 math savants using tin cans and a string.

Weather Model: Distributed Memory (MPI)

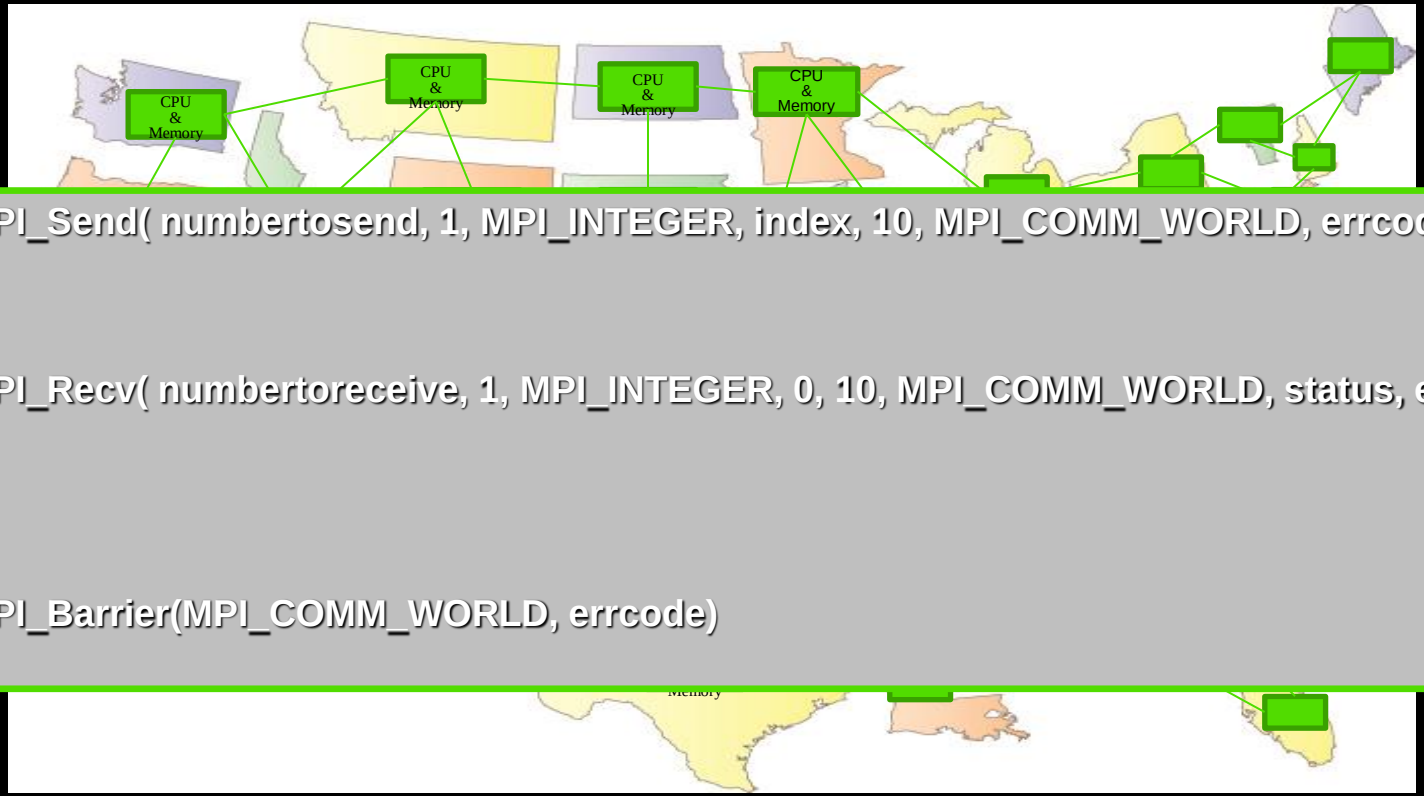


```
call MPI_Send( numbertosend, 1, MPI_INTEGER, index, 10, MPI_COMM_WORLD, errcode)
.
.
call MPI_Recv( numbertoreceive, 1, MPI_INTEGER, 0, 10, MPI_COMM_WORLD, status, errcode)
.
.
.
call MPI_Barrier(MPI_COMM_WORLD, errcode)
.
```

50 meteorologists using a telegraph.

The pieces fit like this…

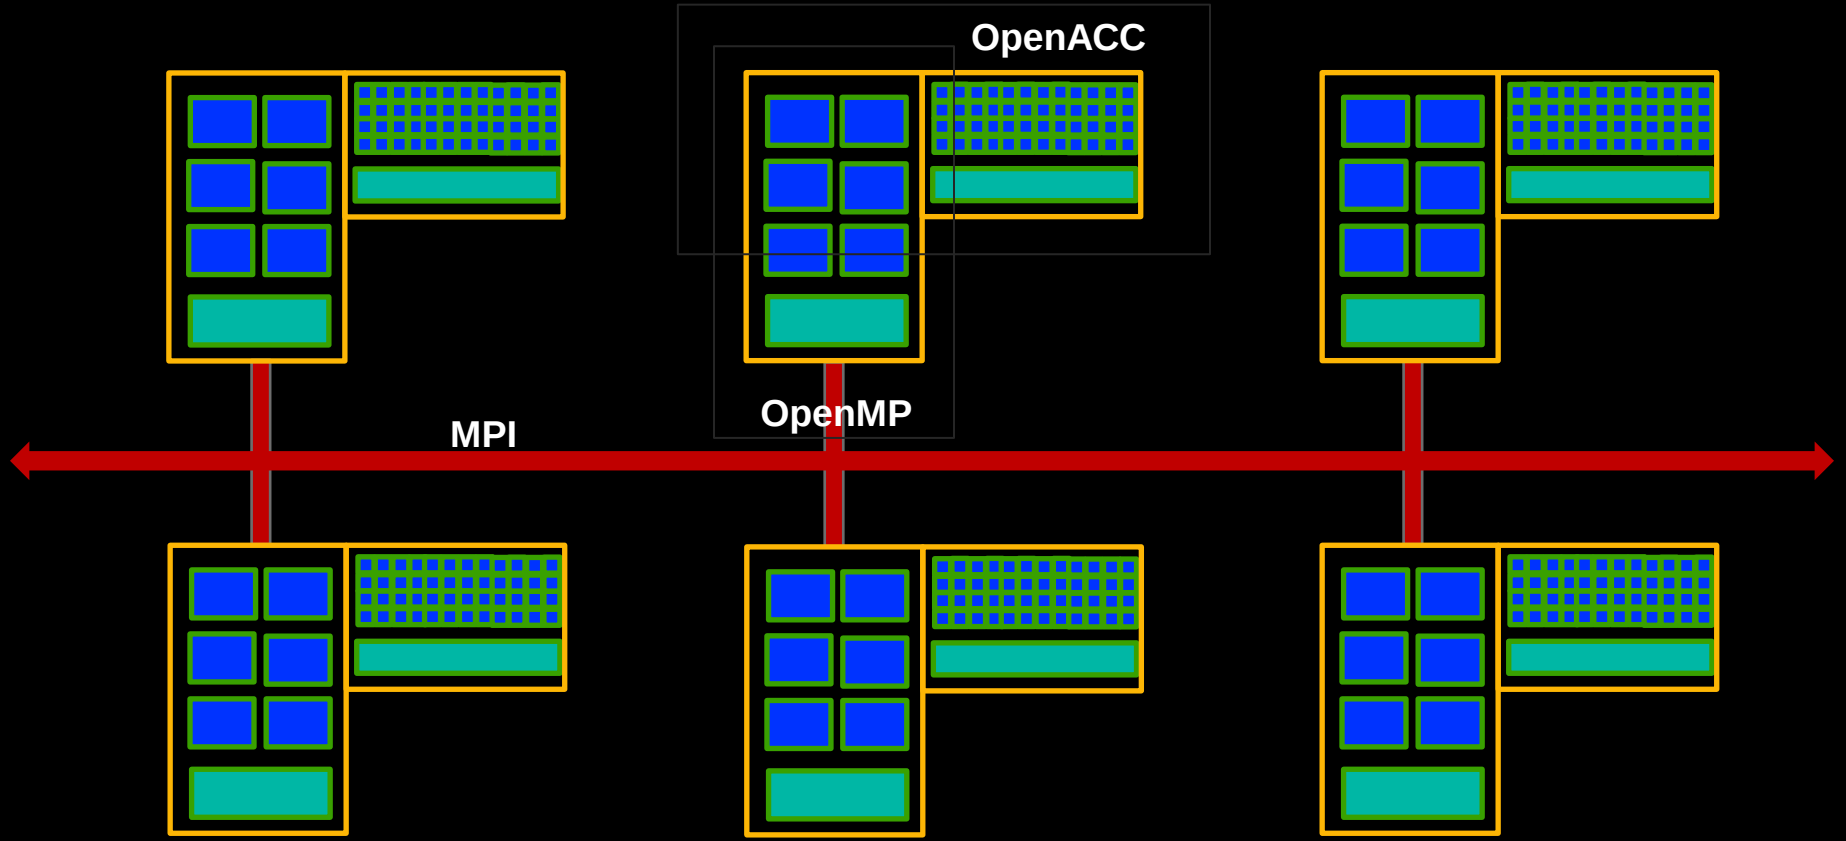

Many Levels and Types of Parallelism

- Vector (SIMD)
- Instruction Level (ILP)
  - Instruction pipelining
  - Superscaler (multiple instruction units)
  - Out-of-order
  - Register renaming
  - Speculative execution
  - Branch prediction

Compiler
(not your problem)

OpenMP

- Multi-Core (Threads)
- SMP/Multi-socket

OpenACC

- Accelerators: GPU & MIC

MPI

- Clusters
- MPPs

Also Important

- ASIC/FPGA/DSP
- RAID/IO

Cores, Nodes, Processors, PEs?

- The most unambiguous way to refer to the smallest useful computing device is as a Processing Element, or PE.
- This is usually the same as a single core.
- “Processors” usually have more than one core – as per the previous list.
- “Nodes” is commonly used to refer to an actual physical unit, most commonly a circuit board or blade with a network connection. These often have multiple processors.

I will try to use the term PE consistently here, but I may slip up myself. Get used to it as you will quite often hear all of the above terms used interchangeably where they shouldn’t be.

MPPs (Massively Parallel Processors)

Distributed memory at largest scale. Shared memory at lower level.

Summit (ORNL)

- 122 PFlops Rmax and 187 PFlops Rpeak
- IBM Power 9, 22 core, 3GHz CPUs
- 2,282,544 cores
- NVIDIA Volta GPUs
- EDR Infiniband

Sunway TaihuLight (NSC, China)

- 93 PFlops Rmax and 125 PFlops Rpeak
- Sunway SW26010 260 core, 1.45GHz CPU
- 10,649,600 cores
- Sunway interconnect

Top 10 Systems as of June 2023

#	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	**Fugaku**	RIKEN Center for Computational Science **Japan**	Fujitsu	ARM 8.2A+ 48C 2.2GHz Torus Fusion Interconnect	**7,630,072**	**442**	**537**	**29.9**
3	**LUMI**	EuroHPC **Finland**	HPE	AMD EPYC 64C 2GHz Slingshot-11 AMD Instinct MI250X	**2,220,288**	**309**	**428**	**6.0**
4	**Leonardo**	EuroHPC **Italy**	Atos	Intel Xeon 8358 32C 2.6GHz Infiniband HDR NVIDIA A100	**1,824,768**	**238**	**304**	**7.4**
5	**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**
6	**Sierra**	Lawrence Livermore National Laboratory **United States**	IBM	Power9 3.1 GHz 22C Infiniband EDR *NVIDIA V100*	**1,572,480**	**95**	**125**	**7.4**
7	**Sunway TaihuLight**	National Super Computer Center in Wuxi **China**	NRCPC	Sunway SW26010 260C 1.45GHz Sunway Interconnect	**10,649,600**	**93**	**125**	**15.3**
8	**Perlmutter**	NERSC **United States**	HPE	EPYC 64C 2.45 GHz Slingshot-10 *NVIDIA A100*	**761,304**	**70**	**93**	**2.6**
9	**Selene**						**79**	**2.6**
10	**Tiahne-2A**						**101**	**18.4**

500	Inspur TS10000, Xeon Gold 6130 16C 2.1GHz, NVIDIA Tesla V100, 25G Ethernet, Inspur Internet Service P China	40,320	1.87	3.52

The word is *Heterogeneous*

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

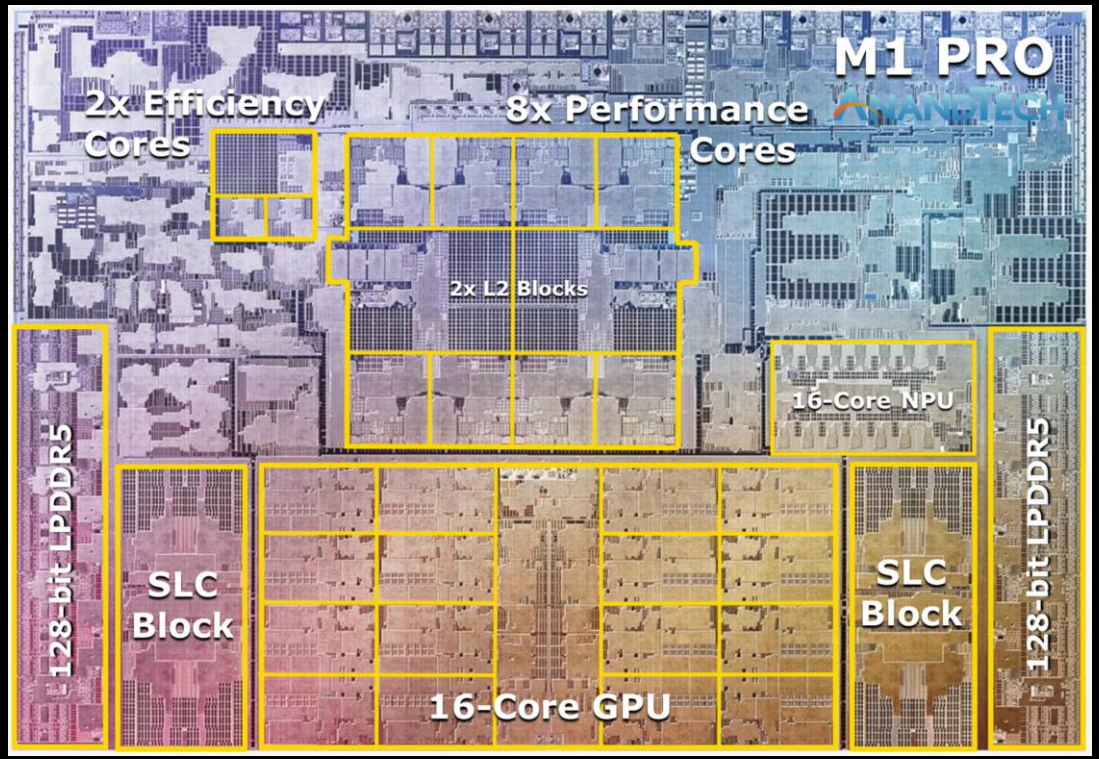

AnandTech

How much of this can you program?

In Conclusion…