Australia eResearch and HPC

October 20-24, 2025

Next Generation System Integration

Technology for Sustaining AI/HPC Growth

Josh Fryman, PhD

Intel Fellow

CTO Intel Government Technologies

Director of IGT R&D

What is the industry doing? Pushing on **hard** limits

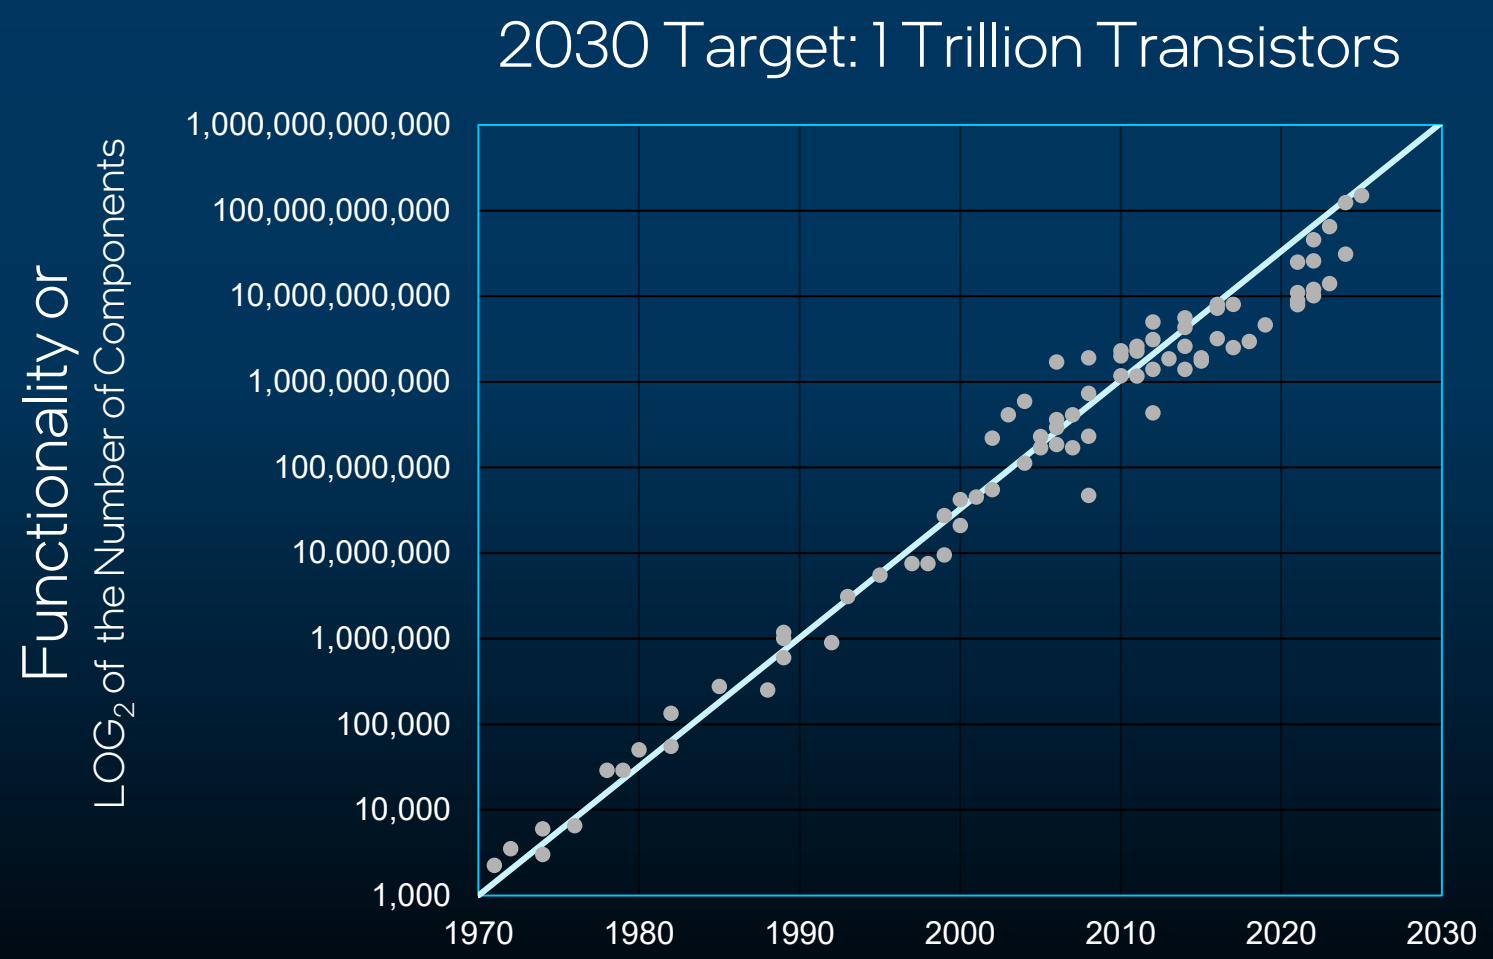

- **Enable multi-core CPU & XPU**
- **Enable Accelerators**
- **Lower Latency & Energy**
- **Increase Memory Capacity & Bandwidth**

“Planet-scale AI infrastructure” shifted the ecosystem

Silicon Valley’s New Strategy: Move Slow and Build Things

Big tech companies are becoming infrastructure companies—just like the steel and railroad giants of old

https://www.wsj.com/tech/ai/silicon-valley-ai-infrastructure-capex

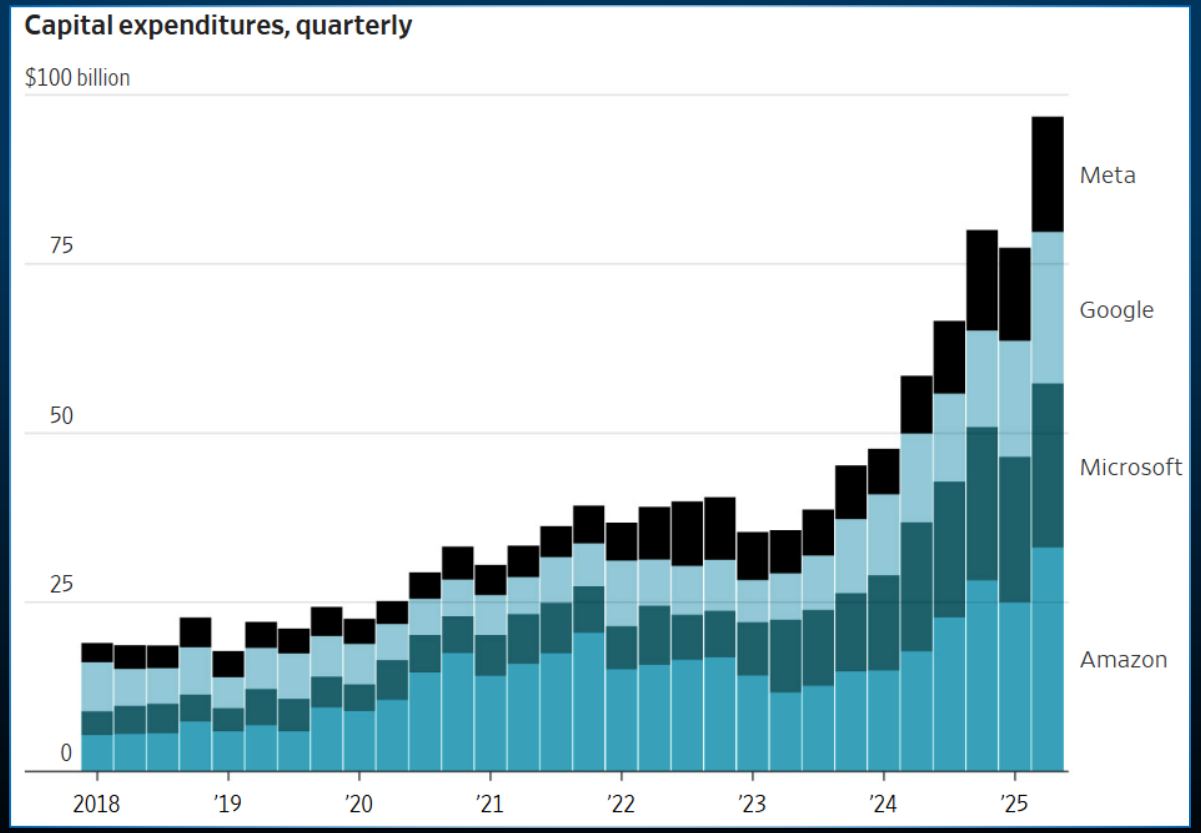

Budgets for building	%
HPC only	3.8%
AI only	10.2%
Blended	86%

HPC-only segment continues a steady decline YoY

Source: Intersect360 Research

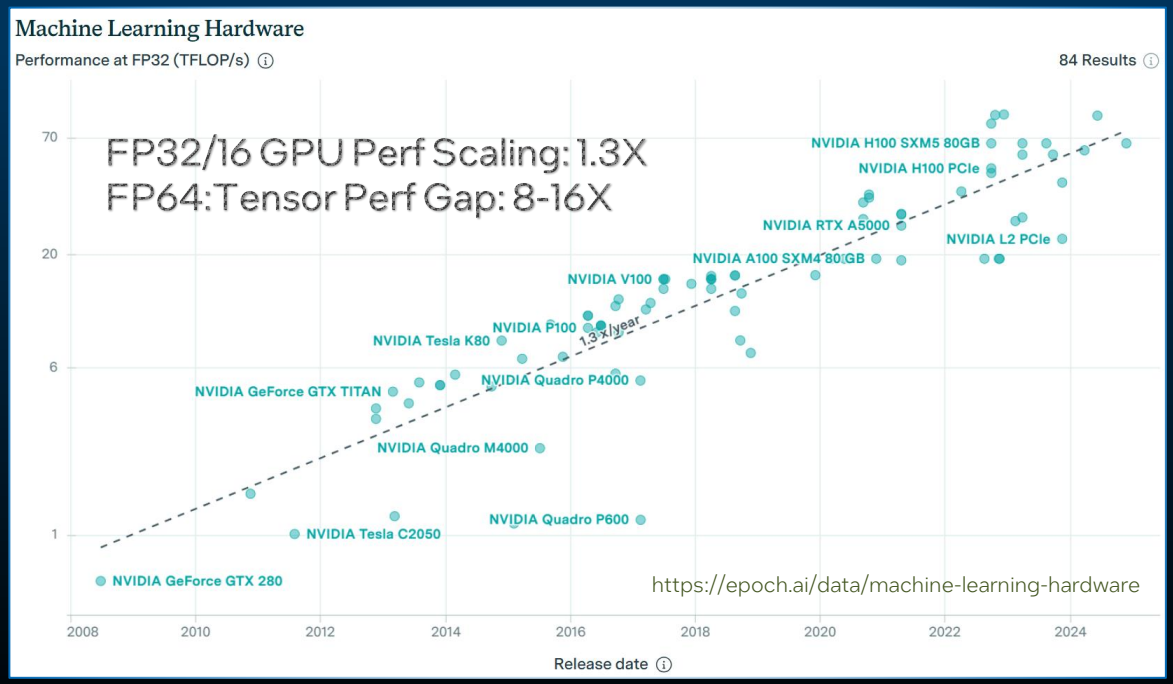

Memory Wall: Why there's really a cliff

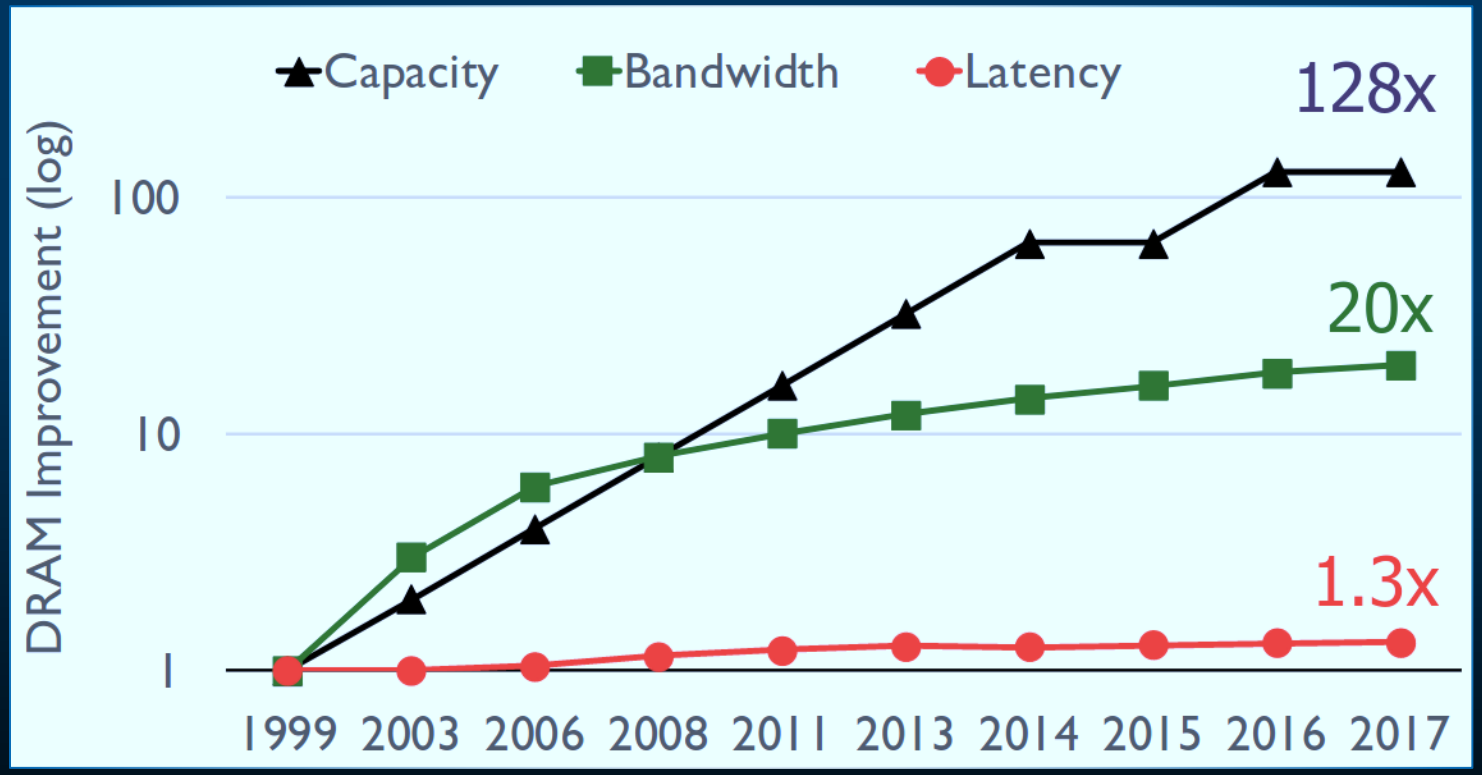

Source: "A modern primer on processing in memory," by Onur Mutlu et al, arxiv.org, Dec 2020.

Recent Historical Prices

DDR5 memory ~ $5/GB

LPDDR5 ~ $3/GB

HBM memory ~ $10-25/GB

Capacity:BW scaling is off by >6x

Customers buy for bandwidth, but pay for (unused) capacity – creating a TCO problem at a global scale

Volume customers drive market solutions, dragging everyone along for the ride – whether or not they have the same problems

The Energy Cost of Data movement

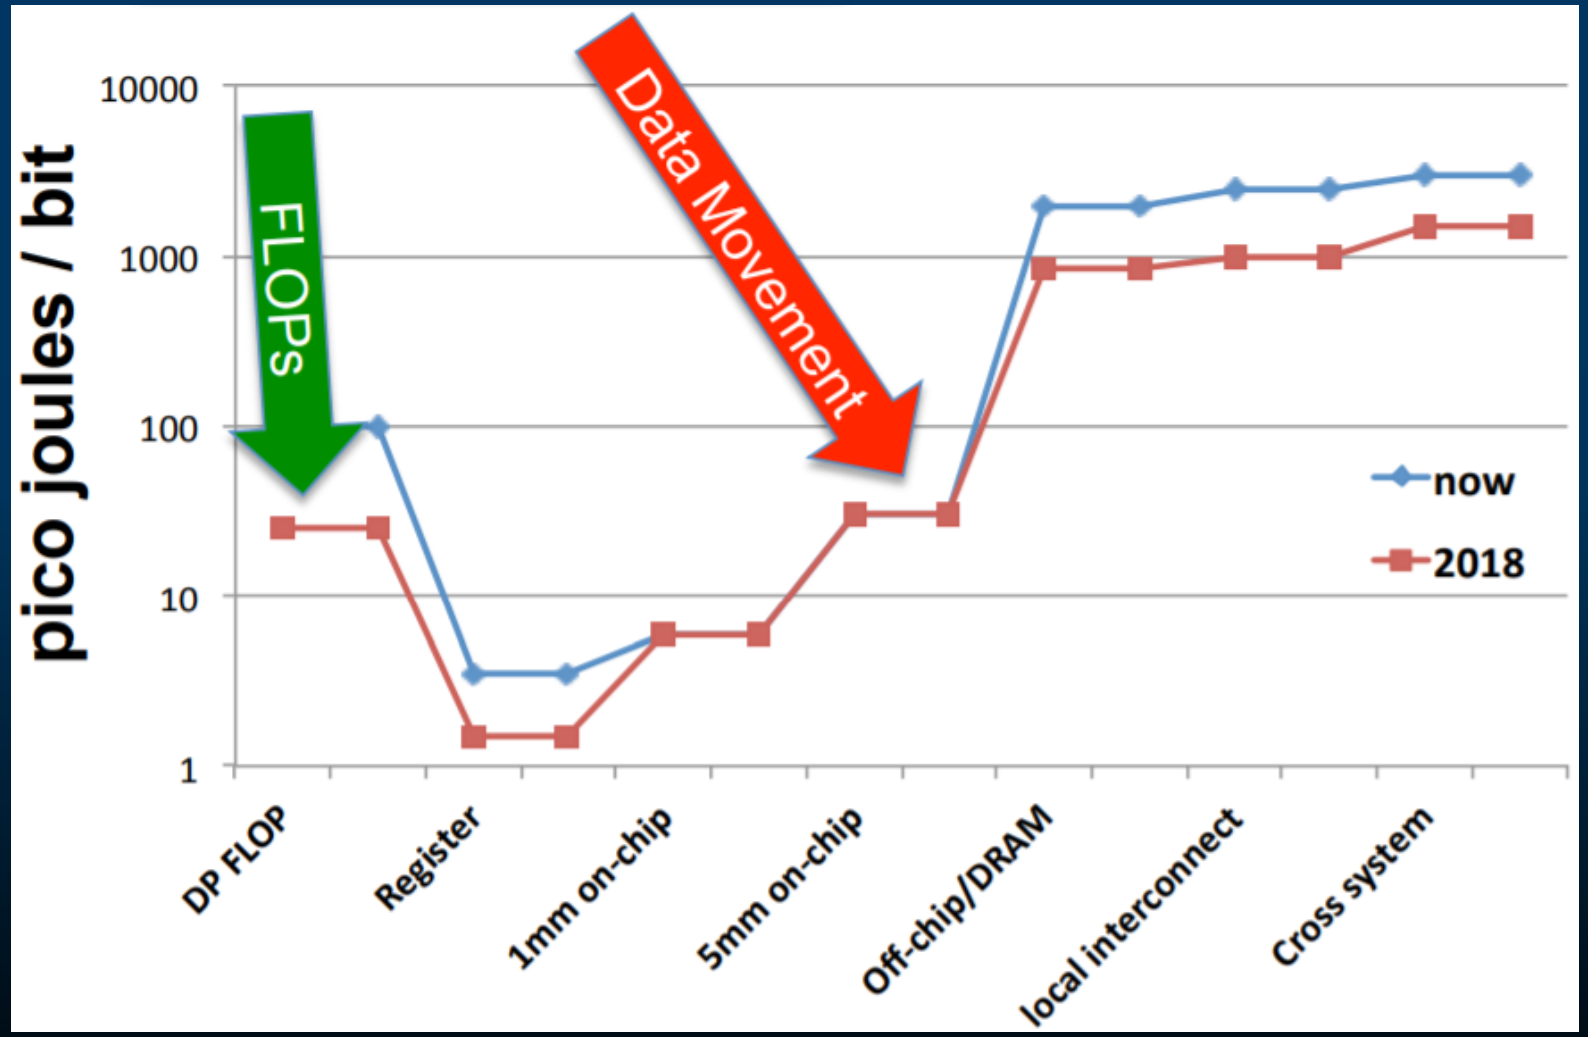

2014 projection of 2018 energy per bit moved*
(by Horst Simon, adapted from John Shalf)
* H. D. Simon, "Why we need Exascale and why we won't get there by 2020," 2014.

"Ten Lessons from Three Generations Shaped Google's TPUv4i"

Operation		Picojoules per Operation		
		45 nm	7 nm	45 / 7
+	Int 8	0.03	0.007	4.3
	Int 32	0.1	0.03	3.3
	BFloat 16	--	0.11	--
	IEEE FP 16	0.4	0.16	2.5
	IEEE FP 32	0.9	0.38	2.4
×	Int 8	0.2	0.07	2.9
	Int 32	3.1	1.48	2.1
	BFloat 16	--	0.21	--
	IEEE FP 16	1.1	0.34	3.2
	IEEE FP 32	3.7	1.31	2.8
SRAM	8 KB SRAM	10	7.5	1.3
	32 KB SRAM	20	8.5	2.4
	1 MB SRAM[1]	100	14	7.1
GeoMean[1]		--	--	2.6
DRAM		Circa 45 nm	Circa 7 nm	
	DDR3/4	1300[2]	1300[2]	1.0
	HBM2	--	250-450[2]	--
	GDDR6	--	350-480[2]	--

Jouppi et al, ISCA 2021

Tackling the root cause of the memory wall cliff

Uncle Sam needs novel memory for nuke sims. So why did it choose Intel?

Didn't the x86 giant just blow up its data storage biz?

Tobias Mann Tue 13 Dec 2022 // 11:30 UTC

The US Department of Energy's Sandia National Labs believes that novel memory tech may be the secret to faster, more accurate nuclear weapon simulations.

https://www.theregister.com/2022/12/13/intel_doe_nukes_mem/

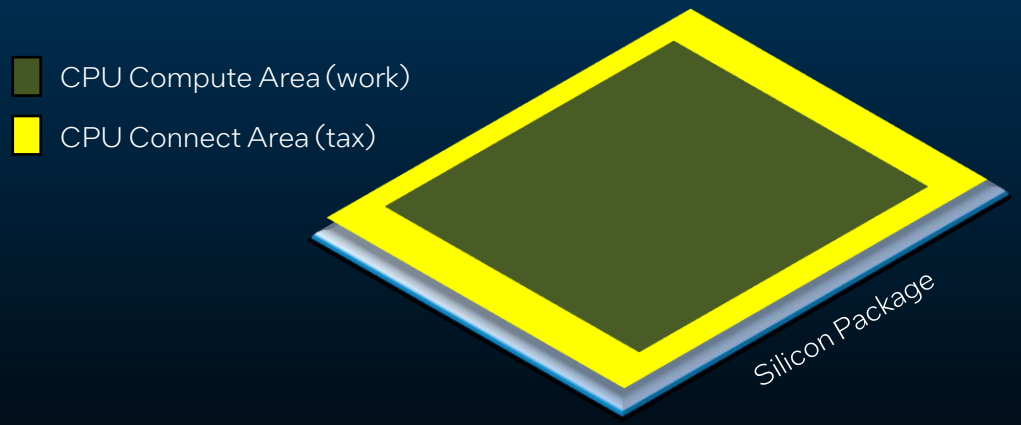

Impact to area for "work" by 2030?

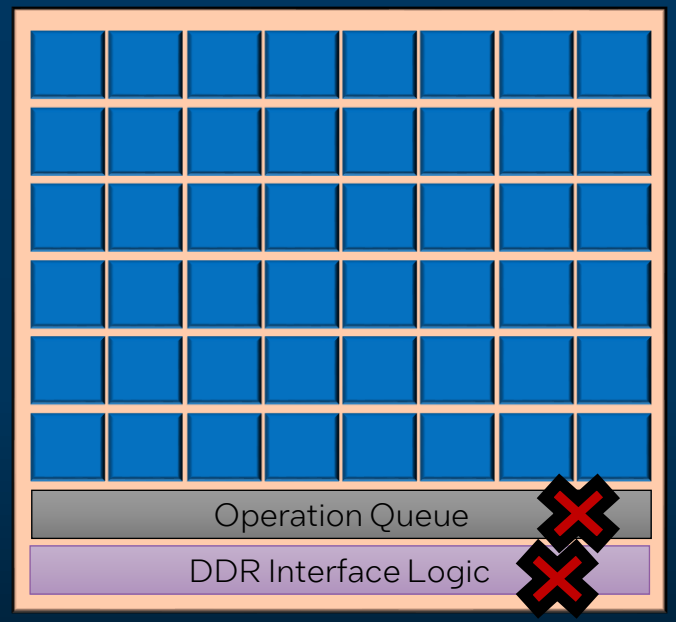

Breaking the perimeter-area problem means delivering O(1) TB/s of bandwidth into O(10) GB of capacity with both dense and sparse access performance

Source: Intel Corporation

Bandwidth density, energy efficiency, and reach

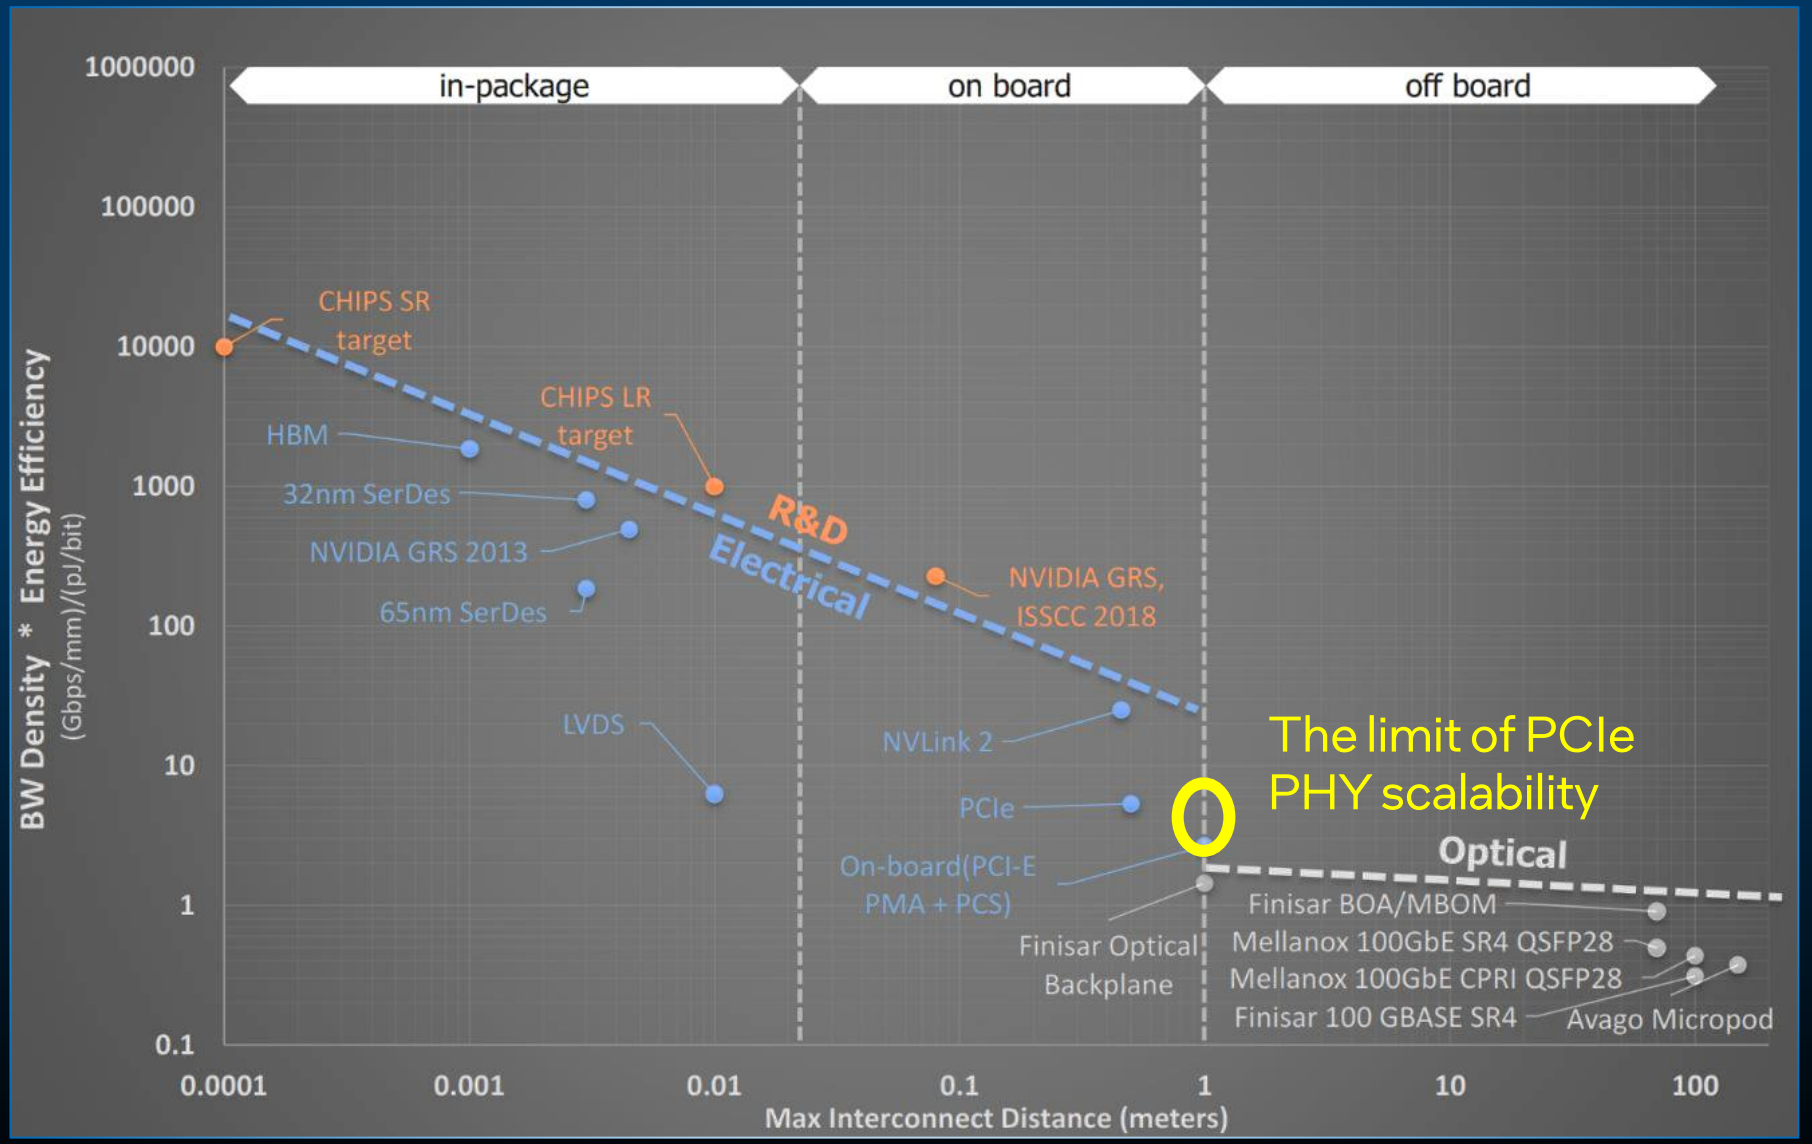

Credit: Gordon Keeler, DARPA PIPES, Distribution A: Public Release

Let's Back Up .. where are we really going next?

Moore's "Page 3" refocused upon by DARPA's ERI programs:

The total cost of making a particular system function must be minimized. To do so, we could amortize the engineering over several identical items, or evolve flexible techniques for the engineering of large functions so that **no disproportionate expense need be borne by a particular array** . . .

. . . it may prove to be **more economical to build large systems out of smaller functions, which are separately packaged and interconnected.** The availability of large functions, combined with functional design and construction, should allow the manufacturer of large systems to design and construct a considerable variety of equipment both rapidly and economically.

– Gordon Moore, Electronics, No. 38, Vol. 8, April 19, 1965

Source: Intel Corporation

Package Sprawl vs. Package Scrapers

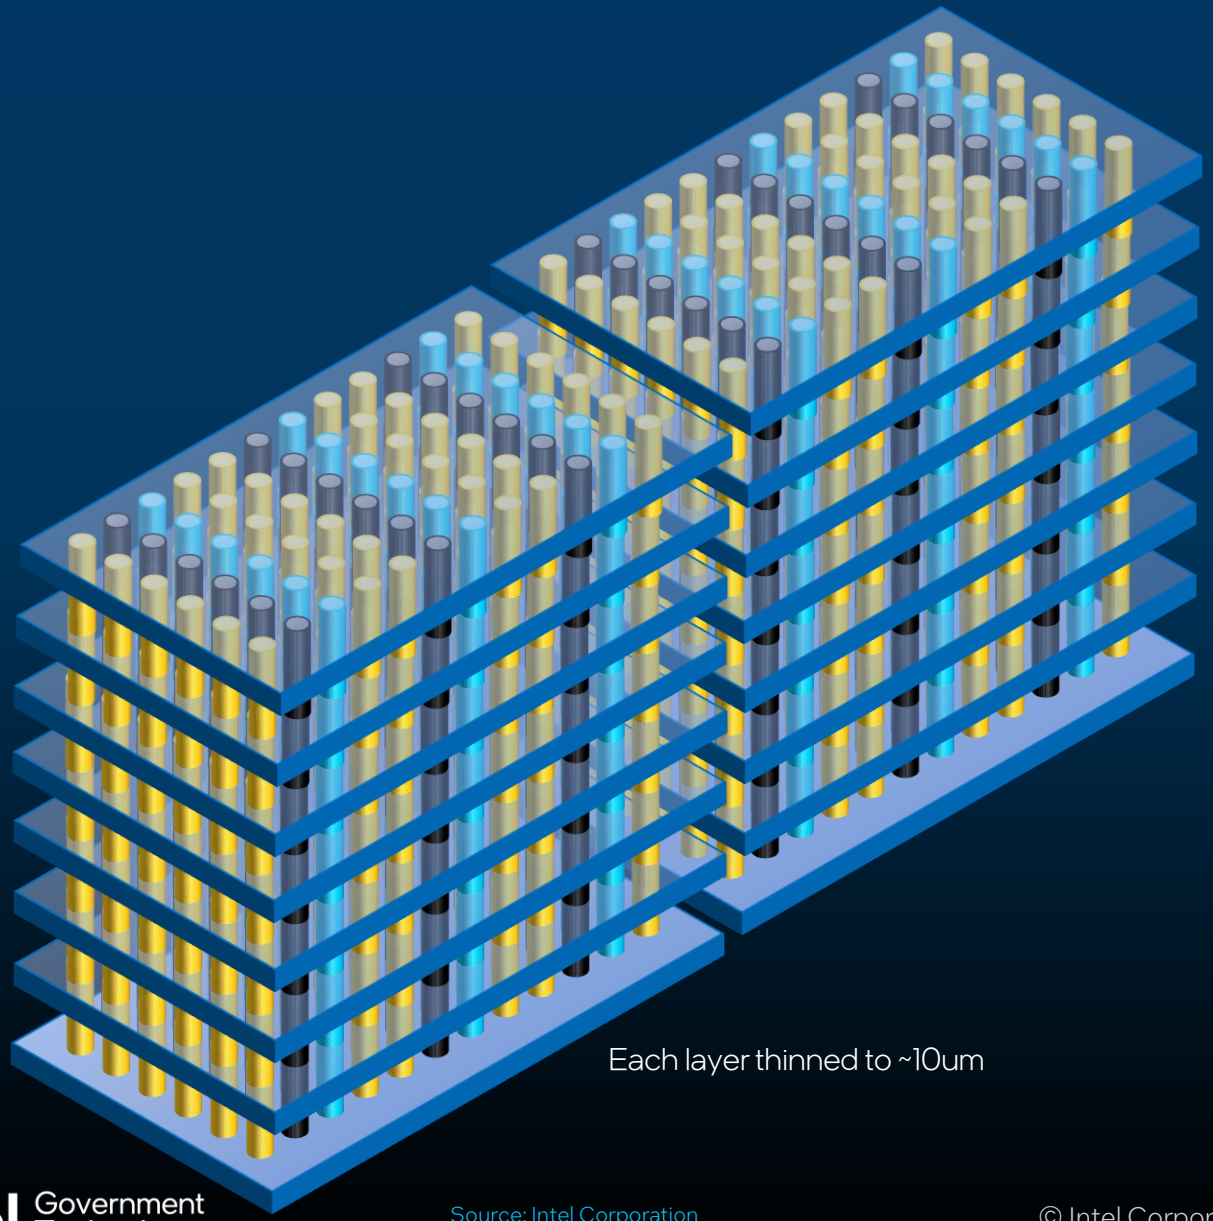

Each layer thinned to ~10um

Assume ~8mm x ~8mm die size

- 800 layers at 10um each makes it ~8mm tall
- Provides ~50,000mm^2 of design space

HBI-TSV density drives bandwidth by pitch

- 9um → 12k/mm2 → 800K total
- 3um → 111k/mm2 → 7.2M total

Assume TSVs are ~67% power/gnd, ~33% I/O

- ~0.27-2.4M I/O signals per 1 GHz
- ~0.260 – 2.4 Pbps total bidir BW per GHz

16 TB/s/dir – 150 TB/s/dir IO capable per GHz

- Lots of floorplanning challenges
- Future delivery of ≤1 um HBI-TSV resolves wiring challenges
- Bottom die PHY area limits IO sustainable

Multiple hard R&D problems

- Thermals, Power, Ground
- Time on tools, **warpage**
- EDA, DFX, DFT, FA, etc.
- Lack of lateral connectivity between high-rises

Source: Intel Corporation

"A New Golden Age of System Architecture"

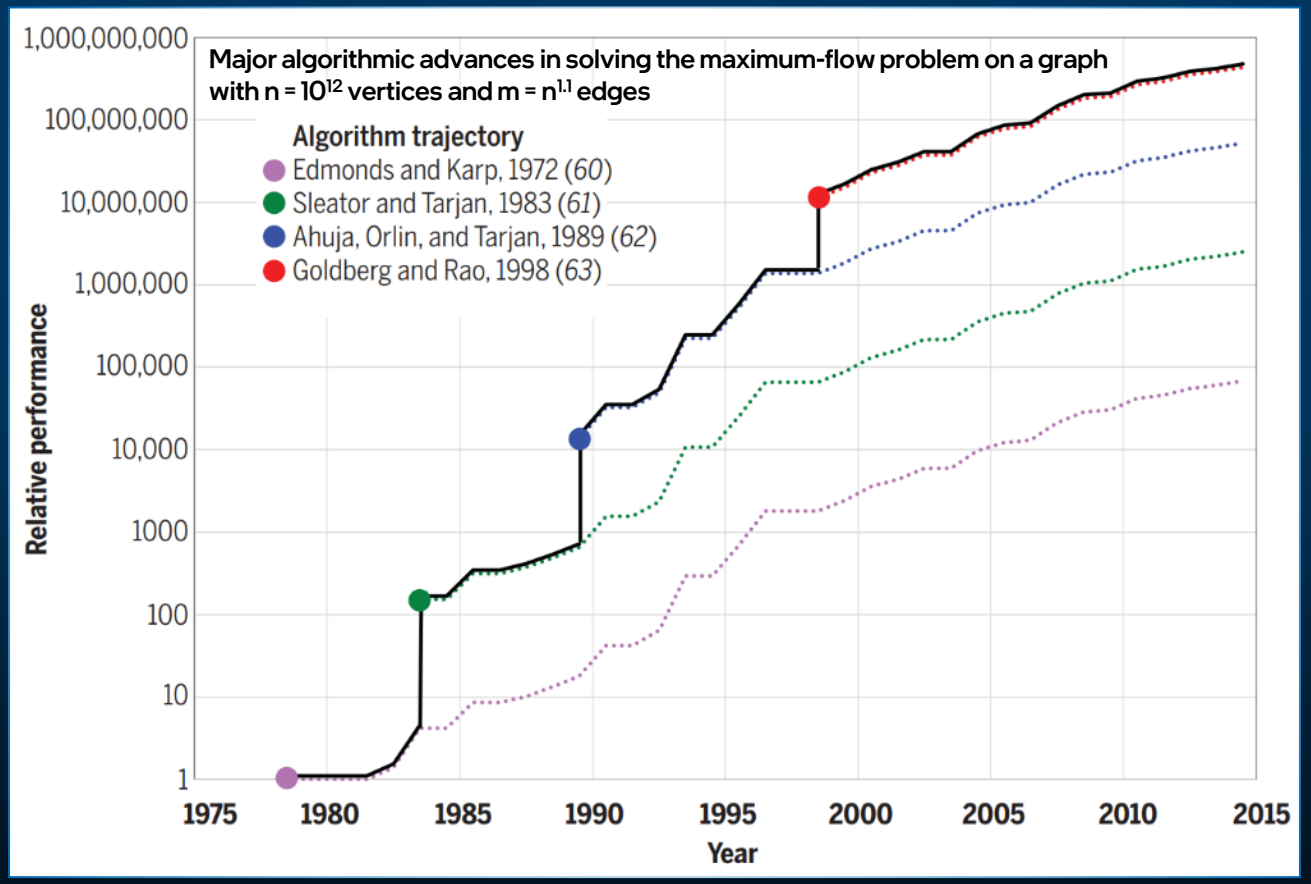

Credit: There's Plenty of Room at the Top, Leiserson et al, *Science, June 2020, Vol 368*

Ozaki Scheme II: A GEMM-oriented emulation of floating-point matrix multiplication using an integer modular technique

Katsuhisa Ozaki[1] Yuki Uchino[2]

and Toshiyuki Imamura[2]

[1] Department of Mathematical Sciences, Shibaura Institute of Technology, Japan

[2] RIKEN Center for Computational Science, Japan

ozaki@sic.shibaura-it.ac.jp

HPC Wire, April 2025

eResearch Challenge: Big-O Do-Over

Fourier transform

$$\widehat{f}(\xi) = \int_{-\infty}^{\infty} f(x)\, e^{-i2\pi\xi x}\, dx, \quad \forall \xi \in \mathbb{R}. \qquad \textbf{(Eq.1)}$$

DFT Sequence O(n^2)

$$X_k = \sum_{m=0}^{n-1} x_m e^{-i2\pi km/n} \qquad k = 0, \ldots, n-1$$

Partial FFT Sequence

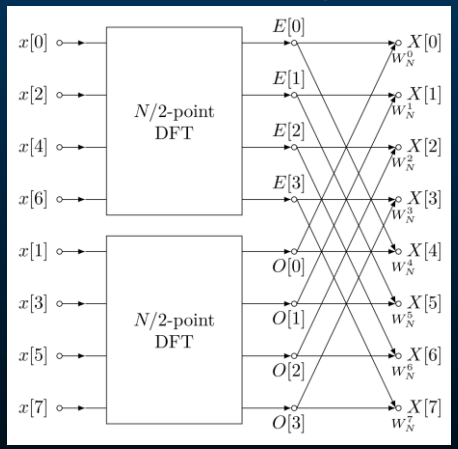

O(n log n)

. . . but O(-) analysis is based on precious compute and free communications
. . . sometimes with memory capacity but not bandwidth factor
. . . and no sensitivity to latency or pressure on any IO interface
. . . O(right) in 1894-1909 is O(wrong) by 2020

Image credit: Wikipedia


intel

But compute performance hit other limits

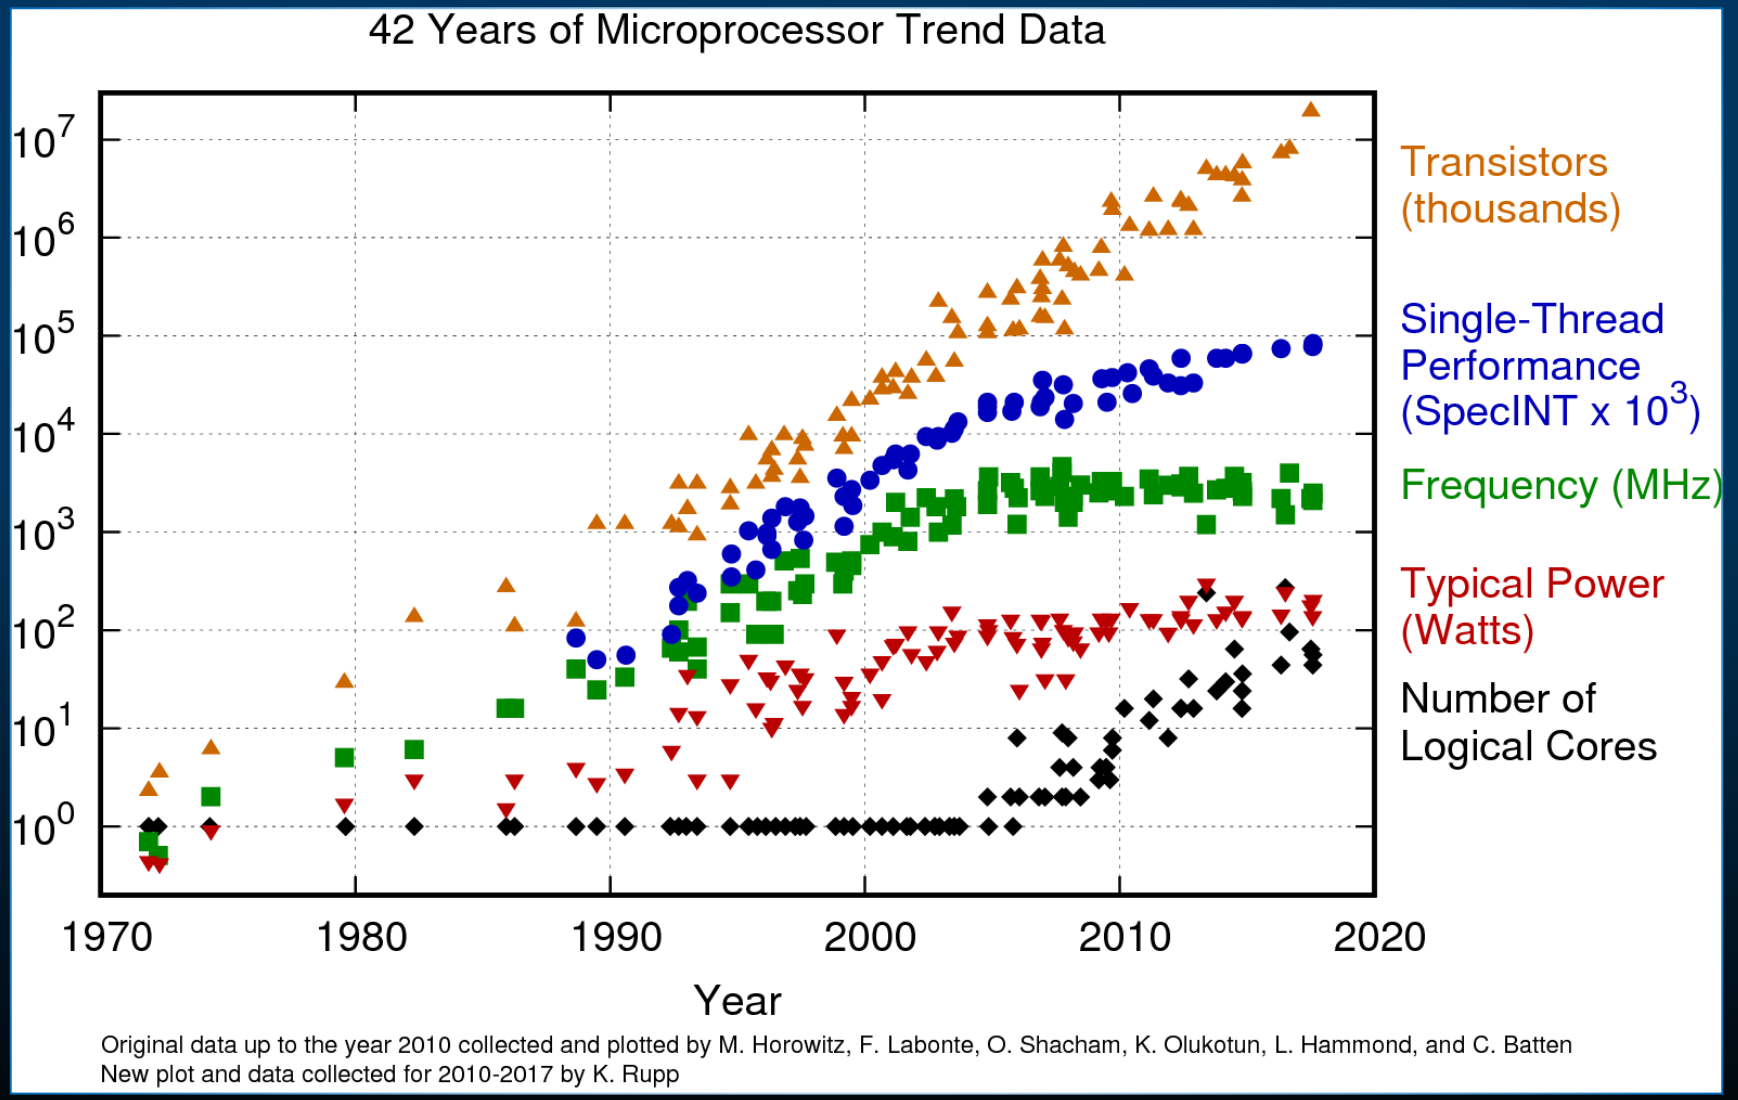

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

https://www.nextplatform.com/2019/06/18/dennard-scaling-demise-puts-permanent-dent-in-supercomputing/

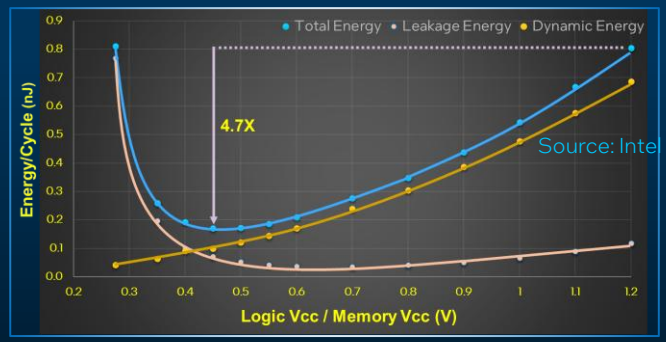

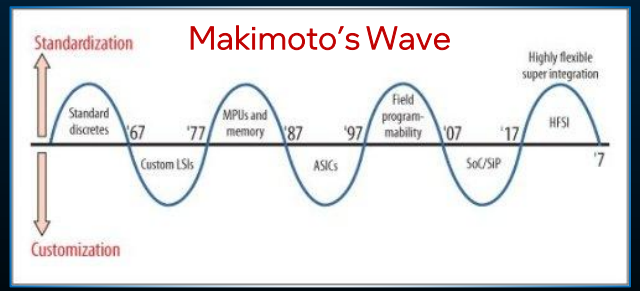

https://semiengineering.com/knowledge_centers/standards-laws/laws/makimotos-wave/