GPU Computing

Hatem Ltaief
Principal Research Scientist, KAUST

Advanced School on High-Performance Computing and Applied AI for High-Resolution Regional Climate Modeling

September 8-19 2025

UM6P | College of Computing

The Abdus Salam International Centre for Theoretical Physics

Advances in Computer Technology

“The rising tide of change shows no respect for the established order. Those who are unwilling or unable to adapt in response to this profound movement not only lose access to the opportunities that the information technology revolution is creating, they risk being rendered obsolete by smarter, more agile, or more daring competitors.”

Jack J. Dongarra, Turing Award 2022
University of Tennessee

Technology Scaling Trends
Exascale in 2022… and then what?

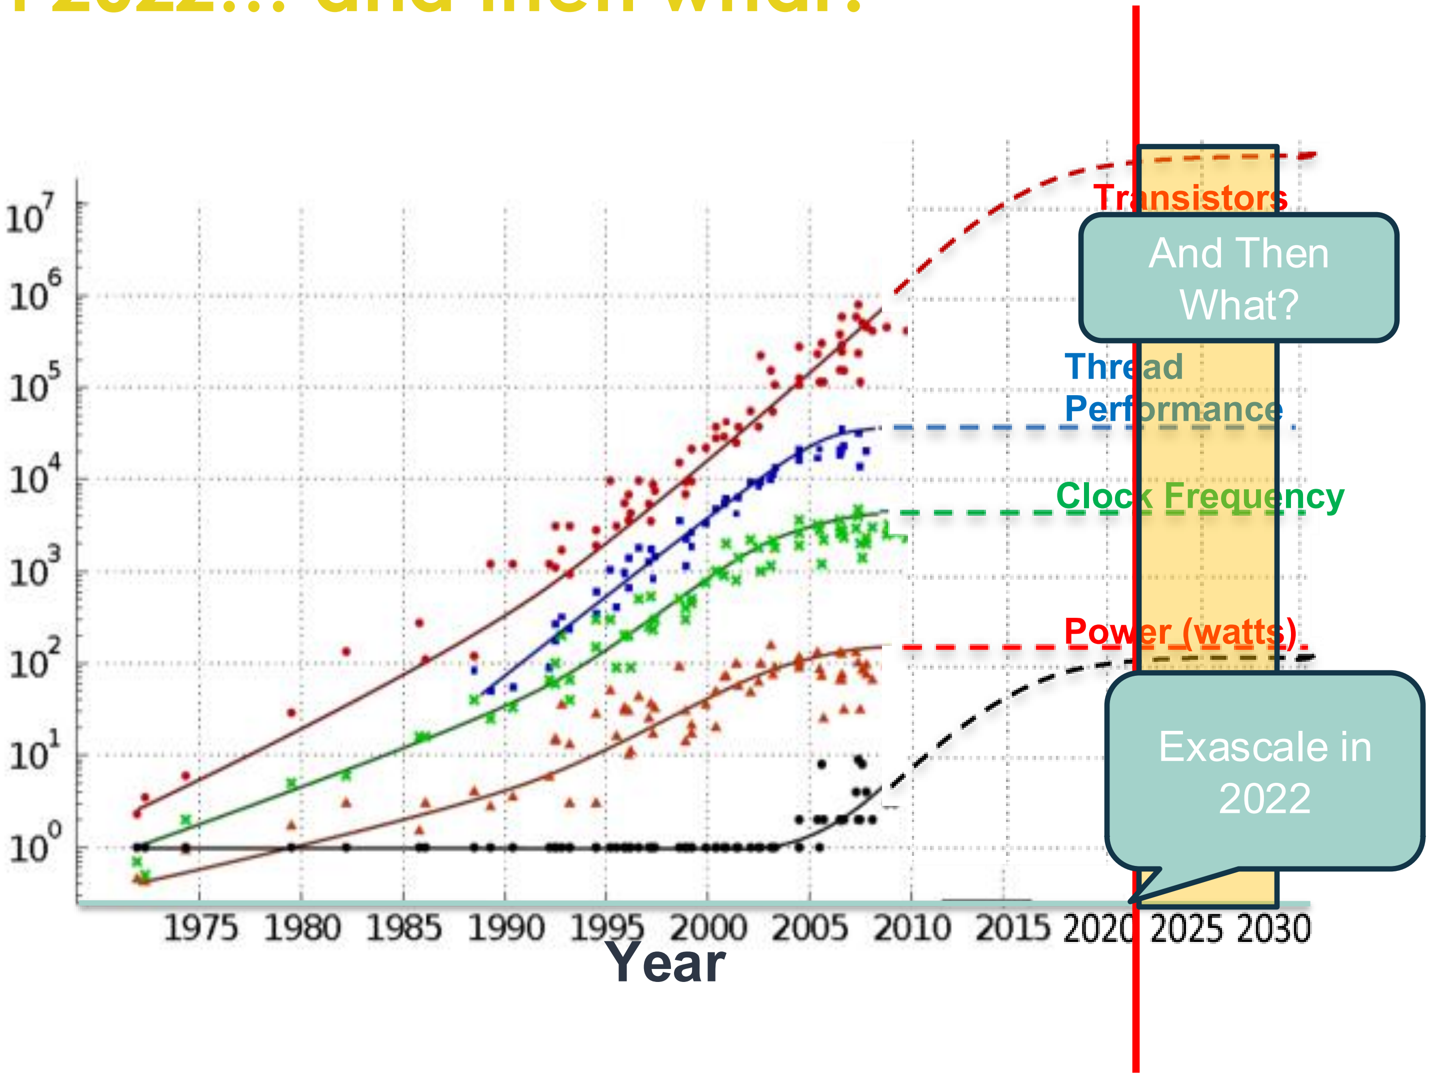

LINPACK: The *Hammering* Performance Benchmark

- ❖ LINPACK Benchmark - Longstanding benchmark started in 1979
- ❖ Lots of positive features; easy to understand and run; shows trends
- ❖ However, much has changed since 1979
- ❖ Arithmetic was expensive then and today it is over-provisioned and inexpensive
- ❖ LINPACK performance of computer systems is no longer strongly correlated to real application performance
- ❖ LINPACK benchmark based on dense matrix multiplication
- ❖ Designing a system for good LINPACK performance can lead to design choices that are wrong for today's applications

Performance Development – Jun’23

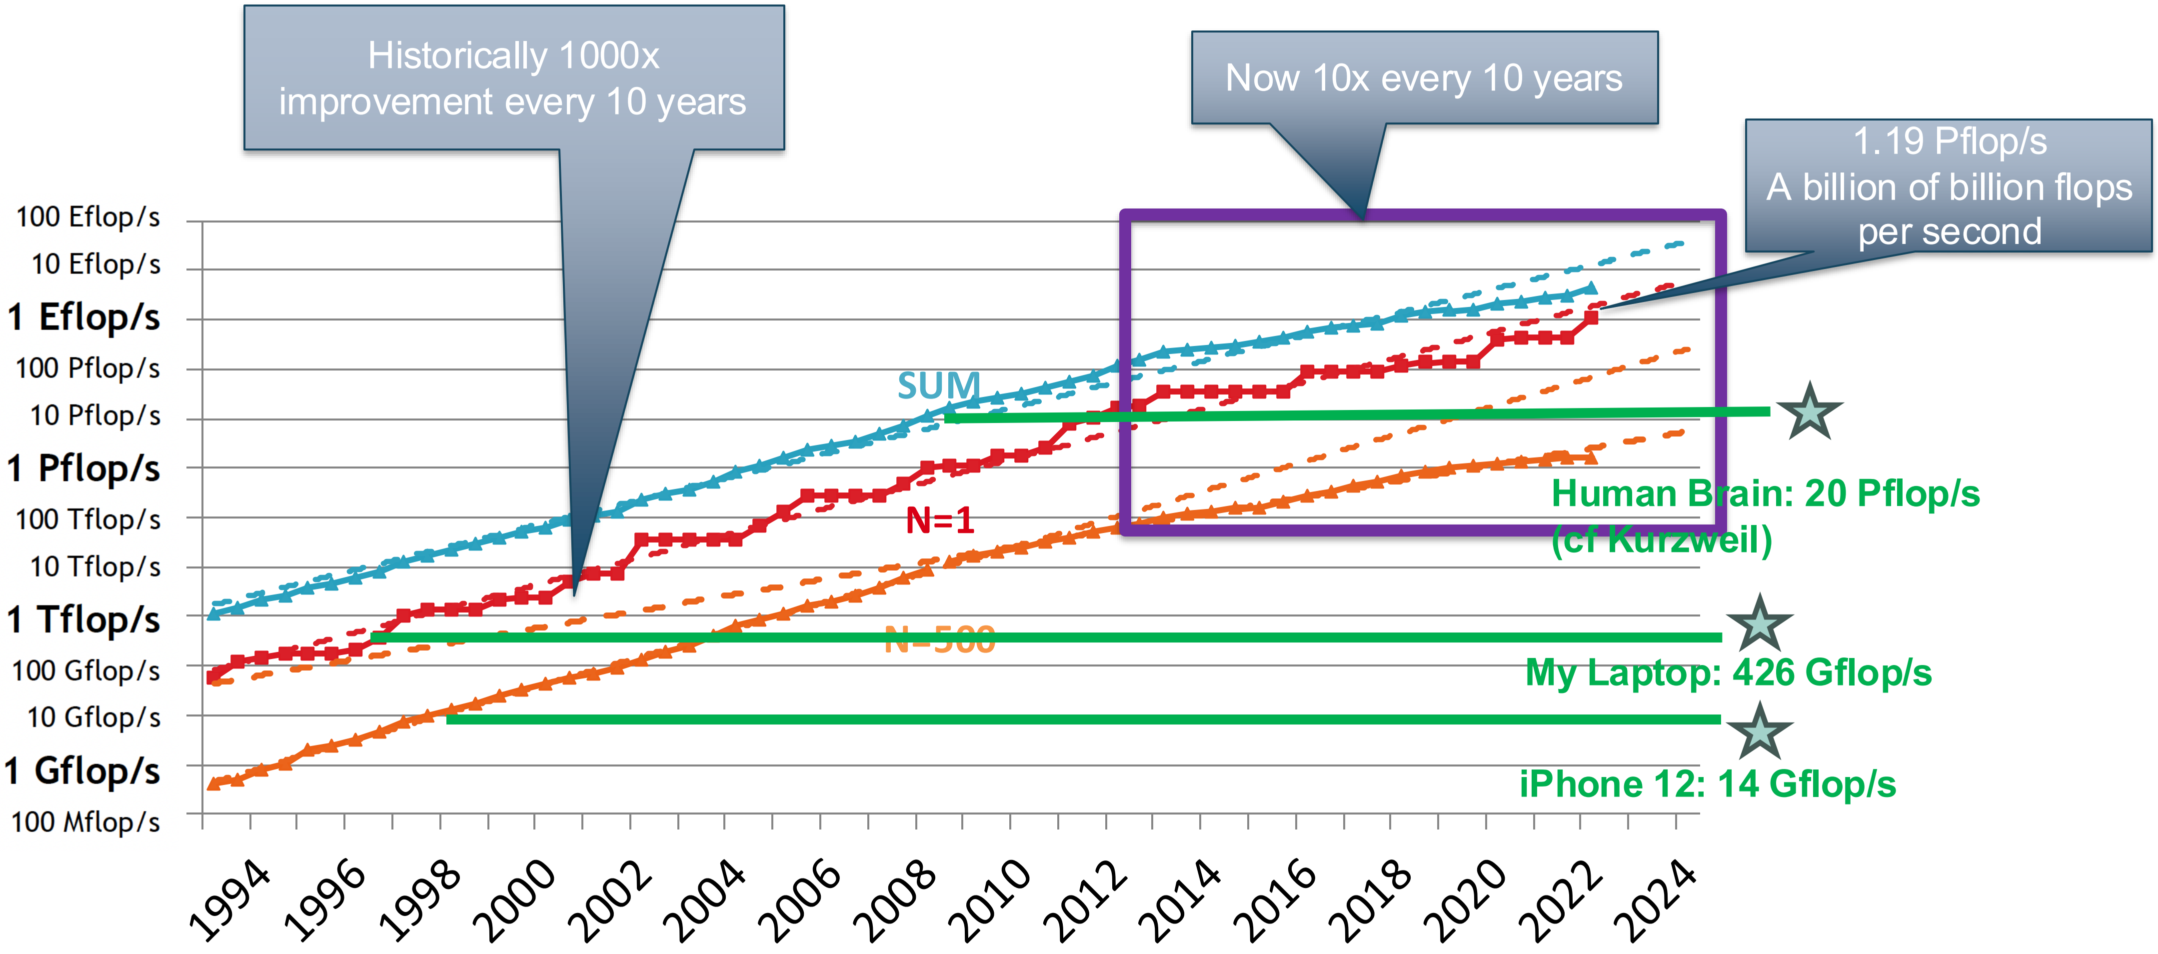

From *Desktop* to *ExaFlop*

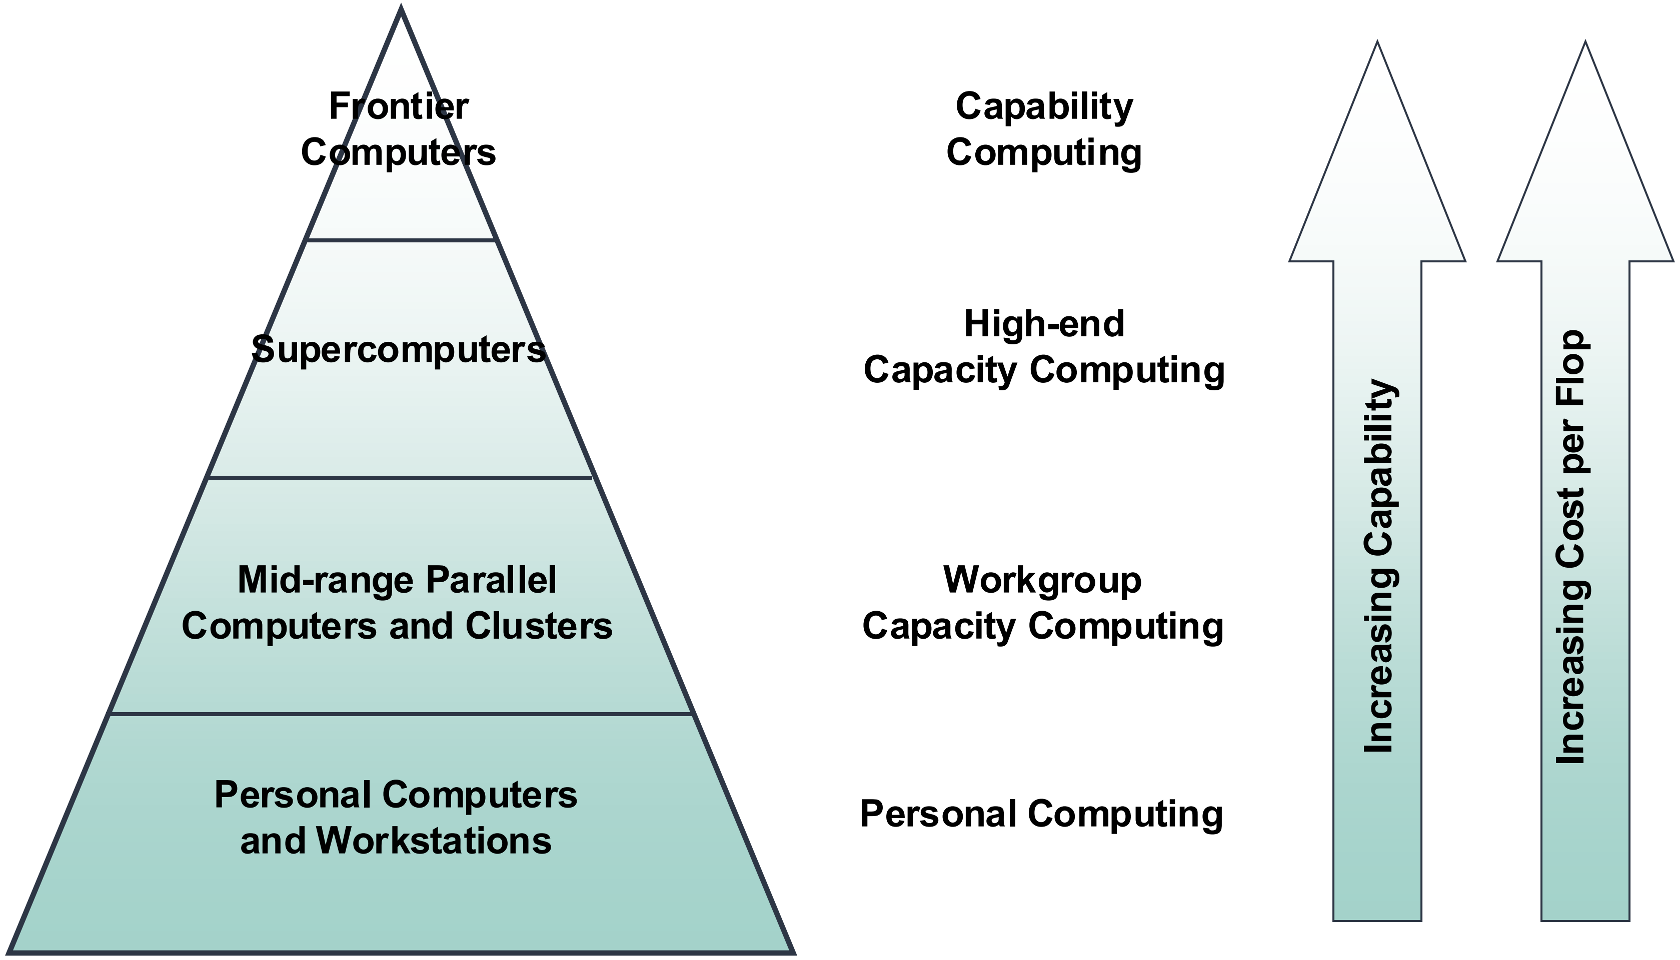

SIAM News by N. Trefethen
Harvard University and University of Oxford

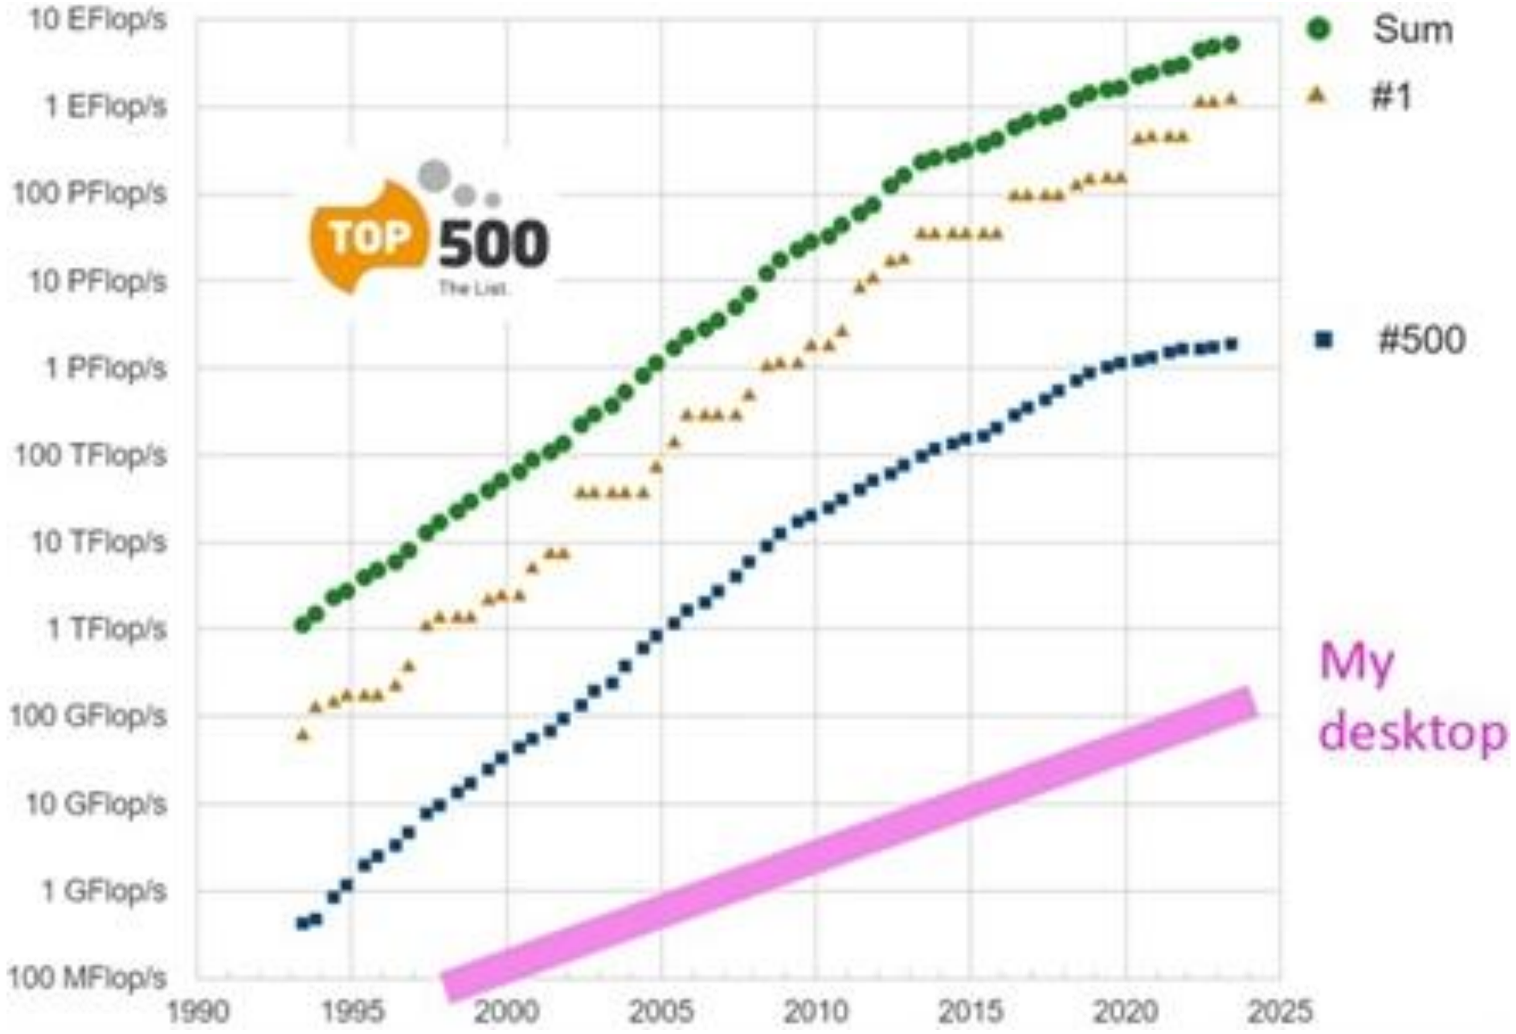

A Bifurcation in Moore's Law? https://sinews.siam.org/Details-Page/a-bifurcation-in-moores-law

SIAM News by SIAM Activity Group on Supercomputing by M. Bader, H. Ltaief, L. C. McInnes, and R. Yokota

How We Can All Play the High-performance Game: From Laptops to Supercomputers

https://sinews.siam.org/Details-Page/how-we-can-all-play-the-high-performance-game-from-laptops-to-supercomputers

Top10 fastest supercomputers for HPL (Nov'24)

#	Site	Manufacturer	TOP10 Computer of the TOP500	Country	Cores	Rmax [Pflops]	Power [MW]
1	Lawrence Livermore National Laboratory	HPE	El Capitan HPE Cray EX255a, AMD EPYC 24C 1.8GHz, Instinct MI300A, Slingshot-11	USA	11,039,616	1,742	29.6
2	Oak Ridge National Laboratory	HPE		USA	9,066,176	1,353	24.6
3	Argonne National Laboratory	Intel		USA	4,742,808	1,012	38.7
4	Microsoft Azure	Microsoft		USA	1,123,200	561.2	
5	Eni S.p.A. Center for Computational Science	HPE		Italy	3,143,520	477.9	8.5
6	RIKEN Center for Computational Science	Fujitsu	Fugaku Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D	Japan	7,630,848	442.0	29.9
7	Swiss National Supercomputing Centre (CSCS)	HPE	Alps HPE Cray EX254n, NVIDIA Grace 72C 3.1GHz, GH200, Slingshot-11	Switzerland	2,121,600	434.9	7.1
8	EuroHPC / CSC	HPE	LUMI HPE Cray EX235a, AMD EPYC 64C 2.0GHz, Instinct MI250X, Slingshot-11	Finland	2,752,704	379.7	7.1
9	EuroHPC / CINECA	EVIDEN	Leonardo Atos BullSequana XH2000, Xeon 32C 2.6GHz, NVIDIA A100, HDR Infiniband	Italy	1,824,768	241.2	7.5
10	Lawrence Livermore National Laboratory	HPE	Tuolumne HPE Cray EX255a, AMD EPYC 24C 1.8GHz, Instinct MI300A, Slingshot-11	USA	1,161,216	208.1	3.4

Top10 fastest supercomputers for HPL-MxP (Nov'24)

Rank	Site	Computer	Cores	HPL Rmax (Eflop/s)	TOP500 Rank	HPL-MxP (Eflop/s)	Speedup
1	DOE/SC/ANL **USA**	**Aurora**, HPE Cray EX, Intel Max 9470 52C, 2.4 GHz, Intel GPU MAX, Slingshot-11	8,159,232	1.012	3	11.6	11.5
2	DOE/SC/ORNL **USA**	**Frontier**, HPE Cray EX235a, AMD Zen-3 (Milan) 64C 2GHz, AMD MI250X, Slingshot-11	8,560,640	1.353	2	11.4	8.4
3	EuroHPC/CSC **Finland**	**LUMI**, HPE Cray EX235a, AMD Zen-3 (Milan) 64C 2GHz, AMD MI250X, Slingshot-11	2,752,704	0.380	8	2.35	6.2
4	RIKEN Center for Comput. Science, **Japan**	**Fugaku**, Fujitsu A64FX 48C 2.2GHz, Tofu D	7,630,848	0.442	6	2.00	4.5
5	EuroHPC/CINECA **Italy**	**Leonardo**, BullSequana XH2000, Xeon 8358 32C 2.6GHz, NVIDIA A100, QR NVIDIA HDR100 IB	1,824,768	0.241	9	1.80	7.5
6	CII, Institute of Science **Japan**	**TSUBAME 4**, HPE Cray XD665, AMD EPYC 9654 96C 2.4GHz, NVIDIA H100, Mellanox NDR200	172,800	0.025	47	0.64	25
7	NVIDIA **USA**	**Selene**, DGX SuperPOD, AMD EPYC 7742 64C 2.25 GHz, Mellanox HDR, NVIDIA A100	555,520	0.063	23	0.63	9.9
8	DOE/SC/LBNL/NERSC **USA**	**Perlmutter**, HPE Cray EX235n, AMD EPYC 7763 64C 2.45 GHz, Slingshot-10, NVIDIA A100	761,856	0.079	19	0.59	7.5
9	FZJ - Jülich **Germany**	**JUWELS Booster Module**, Bull Sequana XH2000 , AMD EPYC 24C 2.8GHz, HDR IB, NVIDIA A100	449,280	0.044	33	0.47	10
10	GENCI-CINES **France**	**Adastra**, HPE Cray EX235a, AMD EPYC 64C 2GHz, AMD 250X, Slingshot-11	319,072	0.046	30	0.30	6.5

Architecture efficiency tracked by the Green 500

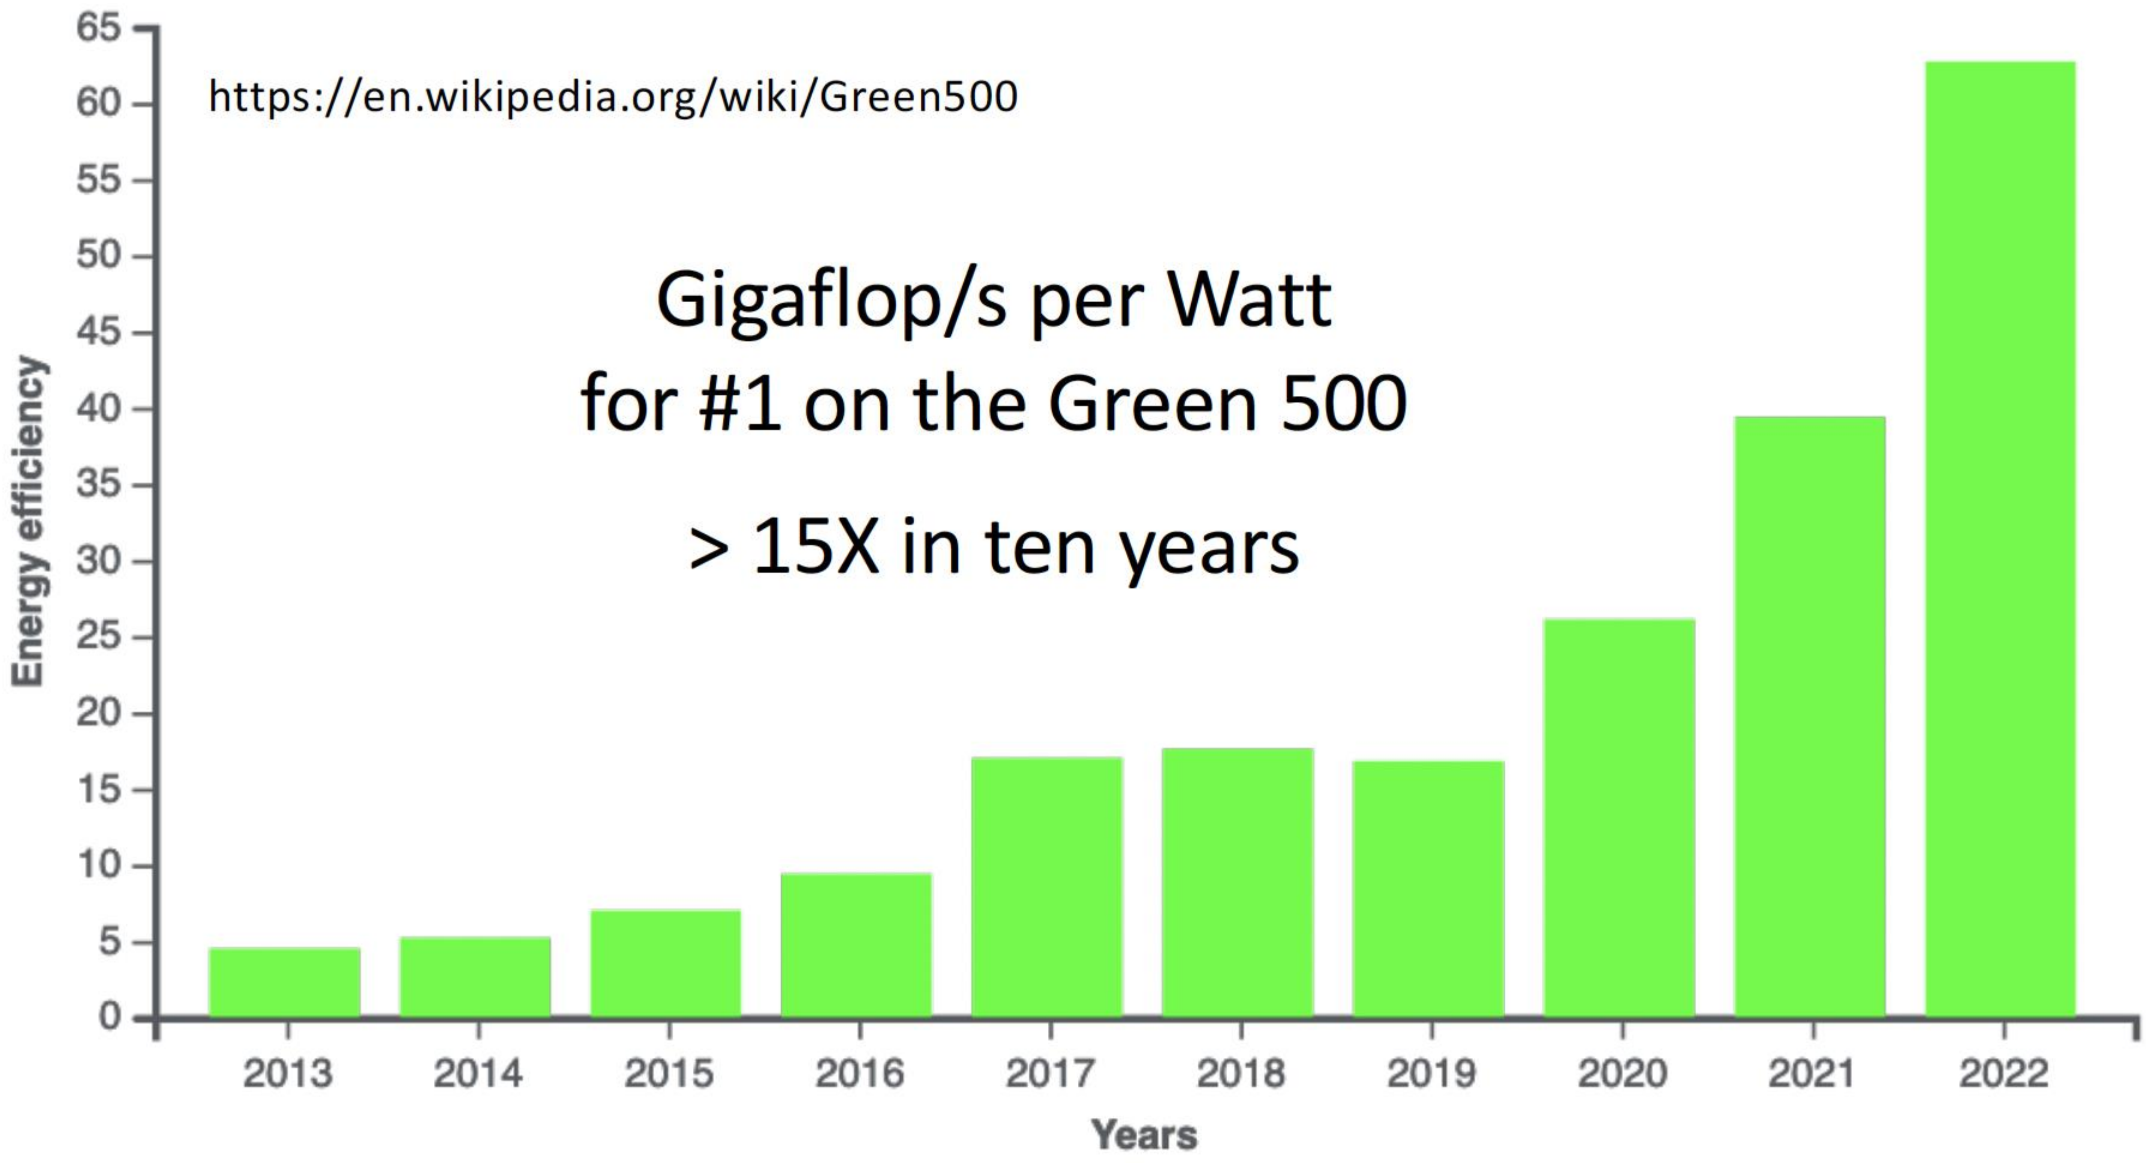

Most energy efficient architectures – Nov’24

Computer	Host	Interconnect	Device	Rmax/ Power
JEDI, BullSequana XH3000	Grace Hopper Superchip 72C 3GHz	Quad-Rail NVIDIA InfiniBand NDR200	NVIDIA Grace	***72.73**
ROMEO-2025, BullSequana XH3000	Grace Hopper Superchip 72C 3GHz	Quad-Rail NVIDIA InfiniBand NDR200	NVIDIA Grace	***70.91**
Adastra 2, HPE Cray EX255a	AMD EPYC 24C 1.8GHz	Slingshot-11	AMD MI300A	***69.10**
Isambard-AI phase 1, HPE Cray EX254n	NVIDIA Grace 72C 3.1GHz	Slingshot-11	NVIDIA Grace	***68.83**
Capella, Lenovo ThinkSystem SD650 V3	AMD EPYC 9334 32C 2.7GHz	InfiniBand NDR200	NVIDIA H100	***68.05**
JETI - JUPITER Exascale Transition Instrument, BullSequana XH3000	Grace Hopper Superchip 72C 3GHz	Quad-Rail NVIDIA InfiniBand NDR200	NVIDIA Grace	***67.96**
Helios GPU, HPE Cray EX254n	NVIDIA Grace 72C 3.1GHz	Slingshot-11	NVIDIA Grace	***66.95**
Henri, ThinkSystem SR670 V2	Intel Xeon Platinum 8362 32C 2.8GHz	Infiniband HDR	NVIDIA H100	**65.40**
HoreKa-Teal, ThinkSystem SD665-N V3	AMD EPYC 9354 32C 3.25GHz	Infiniband NDR200	NVIDIA Grace	**62.96**
rzAdams, HPE Cray EX255a	AMD EPYC 24C 1.8GHz	Slingshot-11	AMD MI300A	**62.80**

[Gflops/Watt]

#1 Systems on the Top500 Over the Past 30 Years

Top500 List (# of times)	Computer	HPL r_{max} (Tflop/s)	Procs/Cores	Matrix Size	Hours To BM	MW
6/93 (1)	TMC CM-5/1024 (DOE LANL)	.060	1,024	52,224	0.4	
11/93 (1)	Fujitsu Numerical Wind Tunnel (Nat. Aerospace Lab of Japan)	.124	140	31,920	0.1	1.
6/94 (1)	Intel XP/S140 (DOE SNL)	.143	3,680	55,700	0.2	
11/94-11/95 (3)	Fujitsu Numerical Wind Tunnel (Nat. Aerospace Lab of Japan)	.170	140	42,000	0.1	1.
6/96 (1)	Hitachi SR2201/1024 (Univ. of Tokyo)	.220	1,024	138,240	2.2	
11/96 (1)	Hitachi CP-PACS/2048 (Univ of Tsukuba)	.368	2,048	103,680	0.6	
6/97-6/00 (7)	Intel ASCI Red (DOE SNL)	2.38	9,632	362,880	3.7	.85
11/00-11/01 (3)	IBM ASCI White, SP Power3 375 MHz (DOE LLNL)	7.23	8,192	518,096	3.6	
6/02-6/04 (5)	NEC Earth-Simulator (JAMSTEC)	35.9	5,120	1,000,000	5.2	6.4
11/04-11/07 (7)	IBM BlueGene/L (DOE LLNL)	478.	212,992	1,000,000	0.4	1.4
6/08-6/09 (3)	IBM Roadrunner –PowerXCell 8i 3.2 Ghz (DOE LANL)	1,105.	129,600	2,329,599	2.1	2.3
11/09-6/10 (2)	Cray Jaguar - XT5-HE 2.6 GHz (DOE ORNL)	1,759.	224,162	5,474,272	17	6.9
11/10 (1)	NUDT Tianhe-1A, X5670 2.93Ghz NVIDIA (NSC Tianjin)	2,566.	186,368	3,600,000	3.4	4.0
6/11-11/11 (2)	Fujitsu K computer, SPARC64 VIIIfx (RIKEN)	10,510.	705,024	11,870,208	29	9.9
6/12 (1)	IBM Sequoia BlueGene/Q (DOE LLNL)	16,324.	1,572,864	12,681,215	23	7.9
11/12 (1)	Cray XK7 Titan AMD + NVIDIA Kepler (DOE ORNL)	17,590.	560,640	4,423,680	0.9	8.2
6/13-11/15 (6)	NUDT Tianhe-2 Intel IvyBridge + Xeon Phi (NSCC Guangzhou)	33,862.	3,120,000	9,960,000	5.4	17.8
6/16-11/17 (4)	Sunway TaihuLight System (NSCC Wuxi)	93,014.	10,549,600	12,288,000	3.7	15.4
6/18-11/19 (4)	IBM Summit Power9 + Nvidia Volta (DOE ORNL)	148,600	2,414,592	16,473,600	3.3	10.1
6/20-11/22 (4)	Fujitsu Fugaku ARM A64FX (RIKEN)	442,010	7,630,828	21,288,960	4.4	29.9
6/22 - ? (4)	HPE Frontier AMD + AMD (DOE ORNL)	1,102,000	7,733,248	24,440,832	2.5	21.1
11/24 - ?	HPE El Capitan AMD + AMD (DOE LLNL)	1,742,000	11,039,616.	?	?	29.6


11
7
3

Top500 Country/System Share

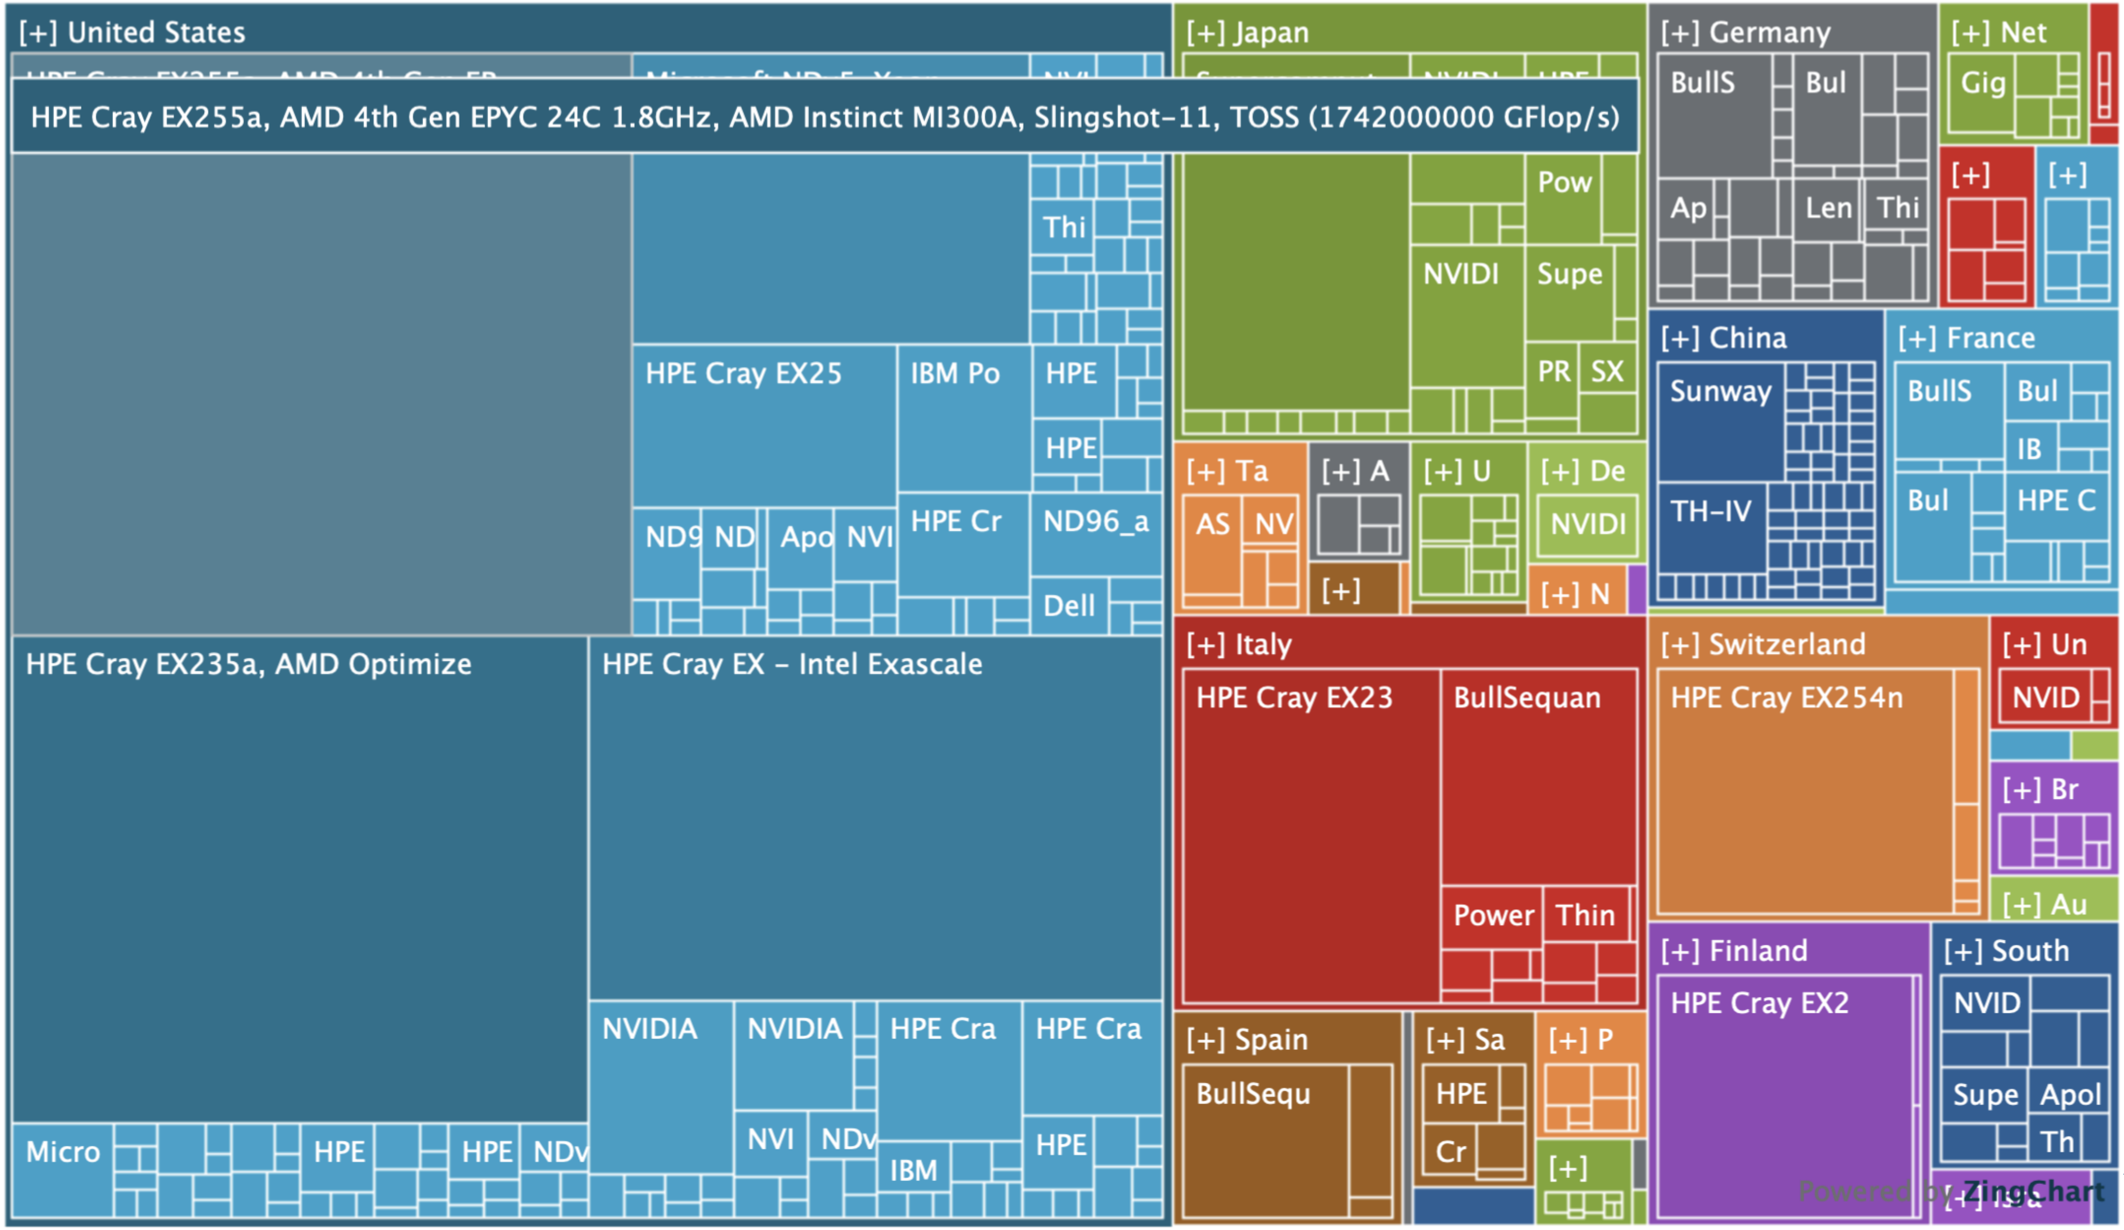

Top500 Performance Breakdown – Jun'23

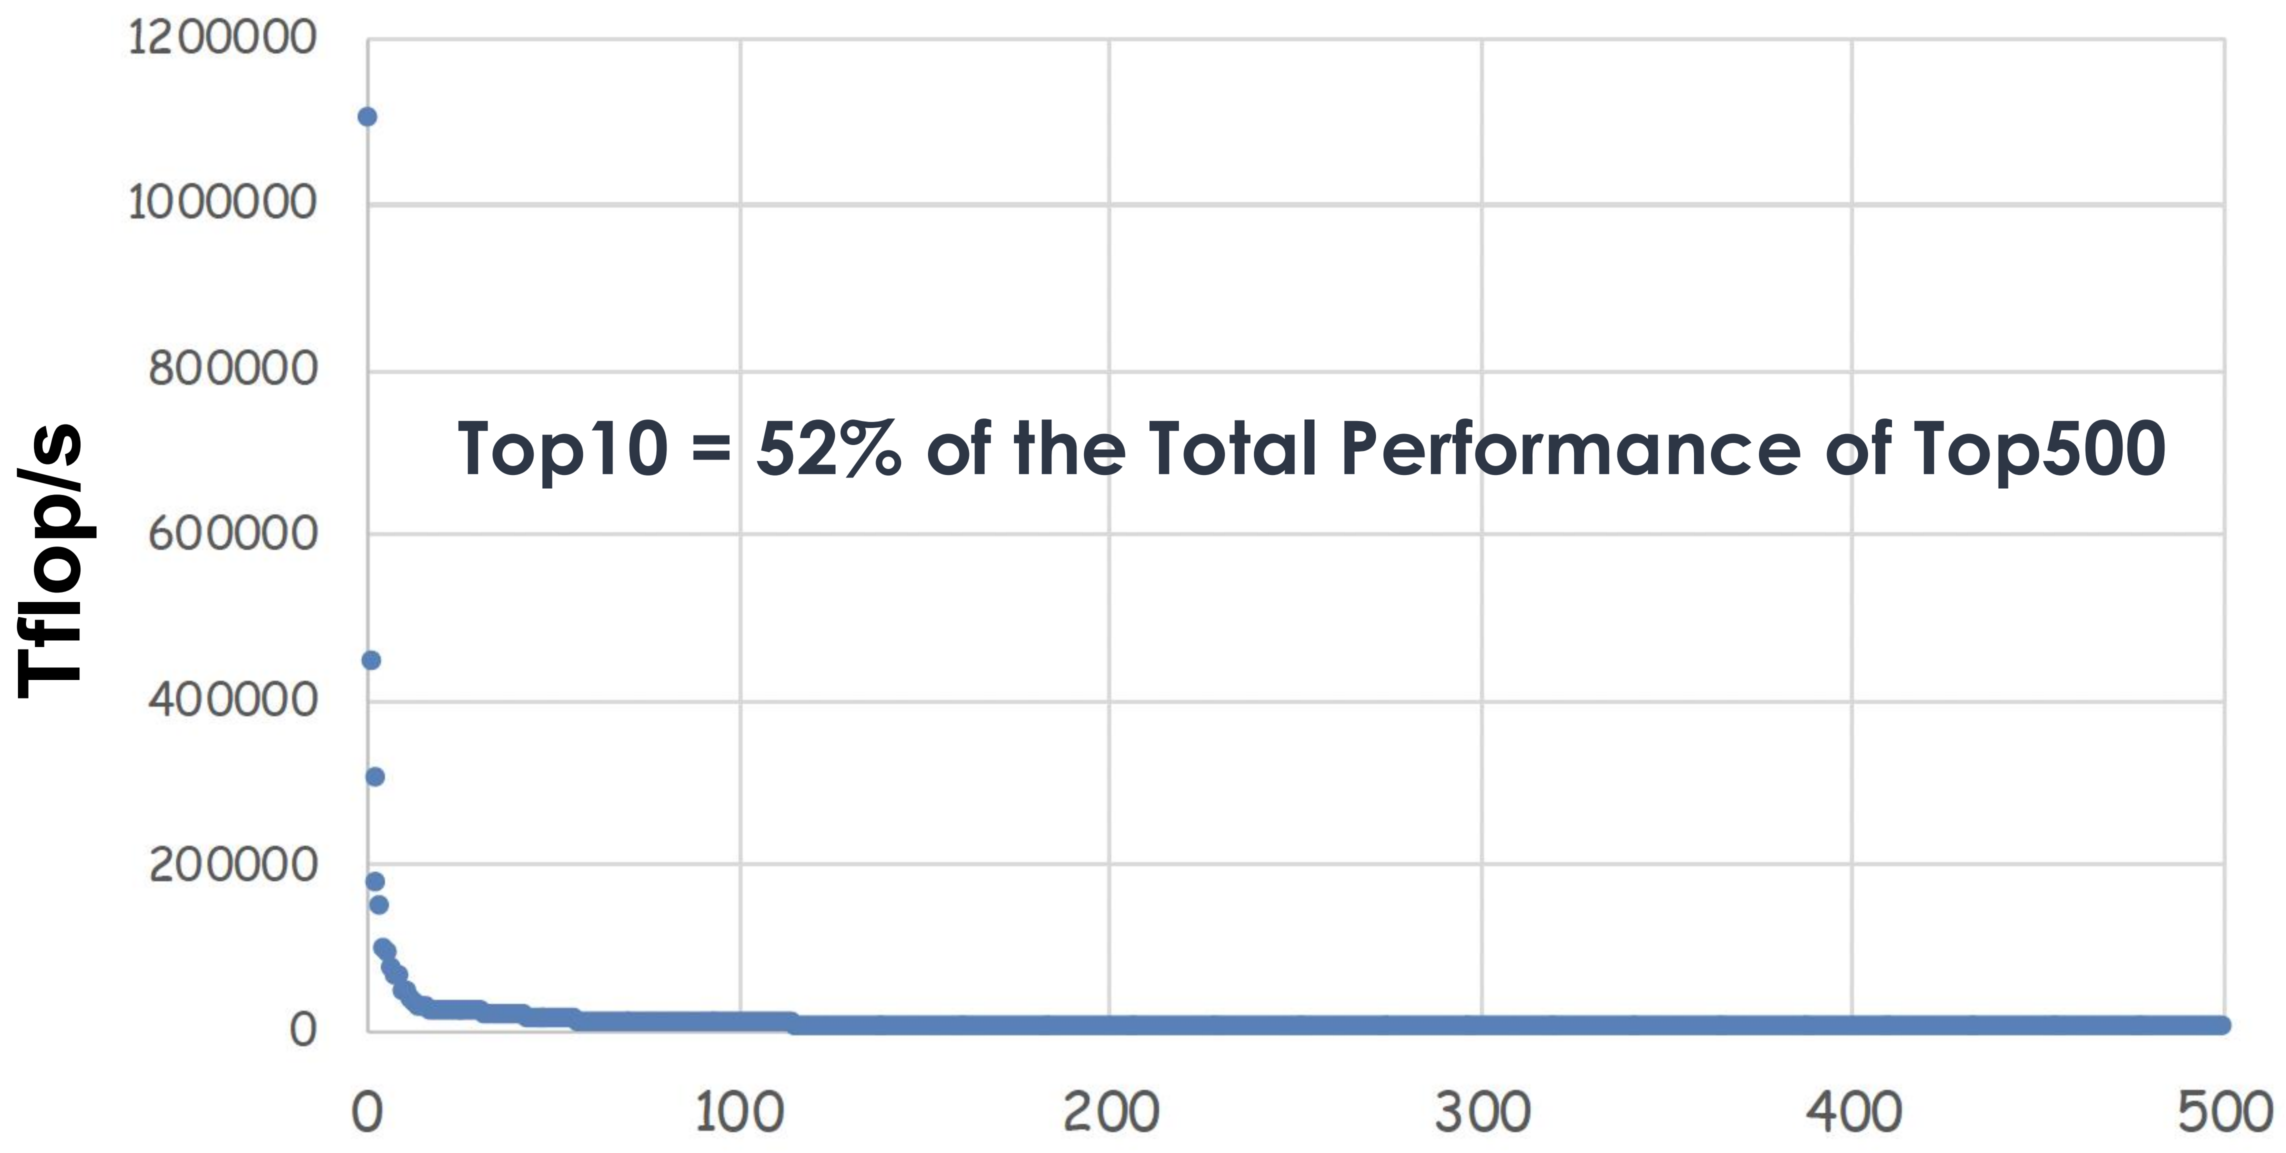

Top500 System Replacement Rate – Jun'23

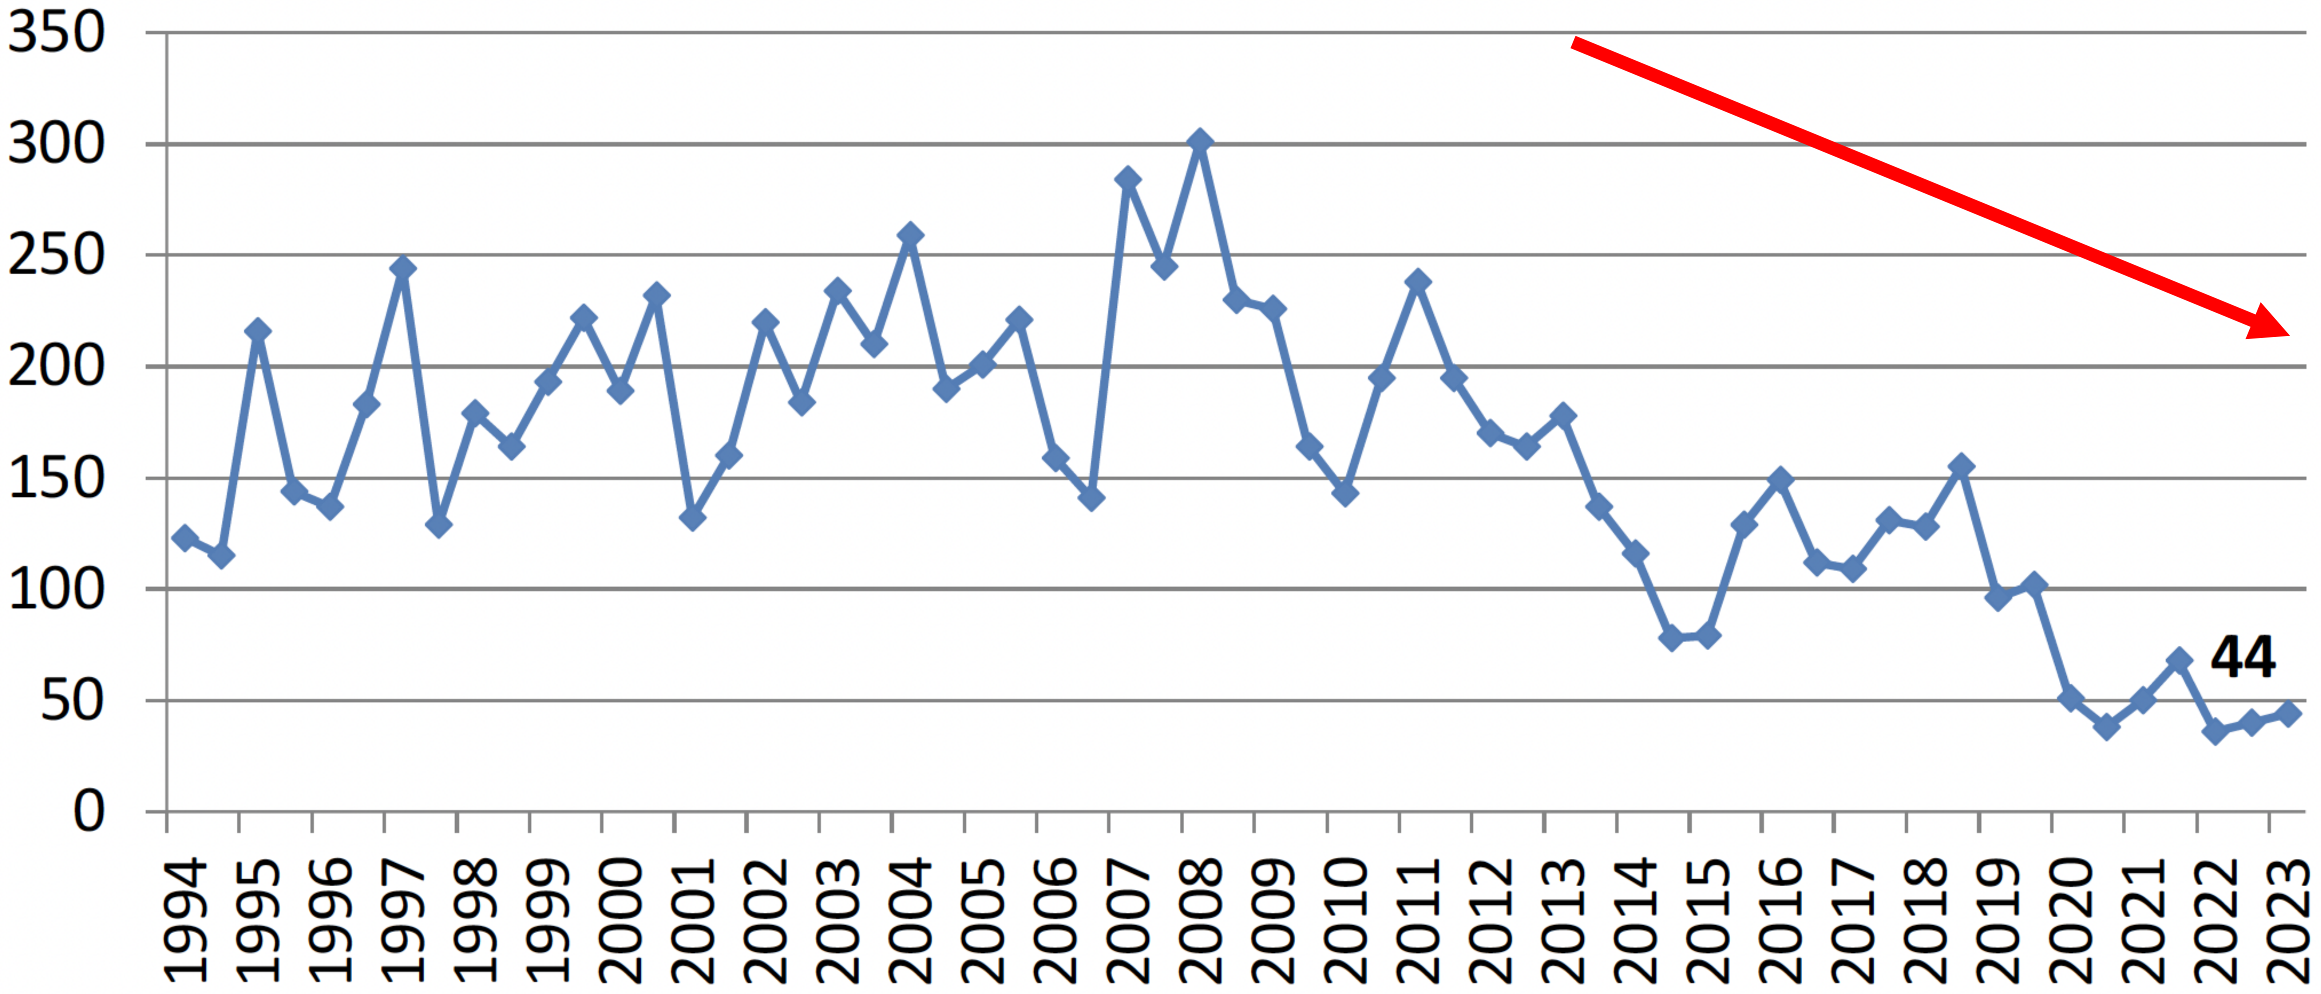

Frequencies of Cores – Jun'23

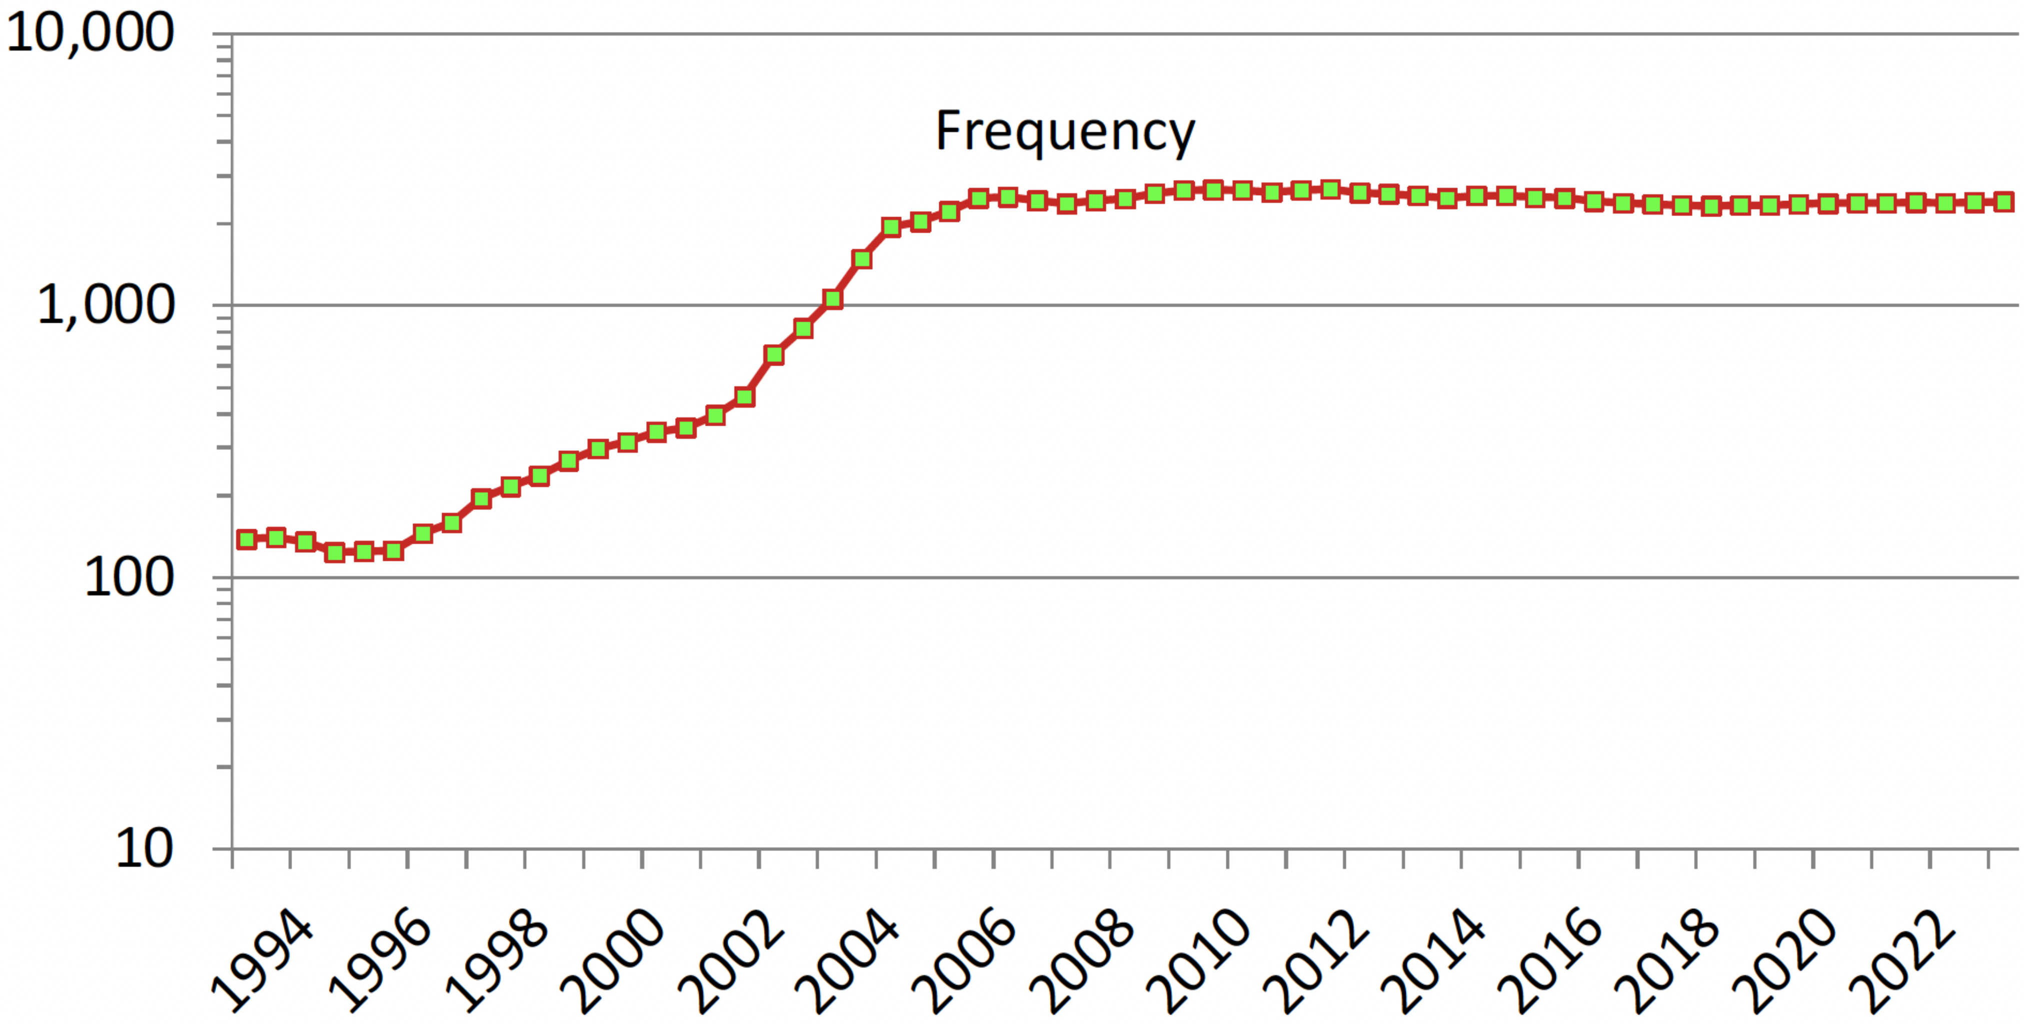

R_MAX PER CORES – Jun'23

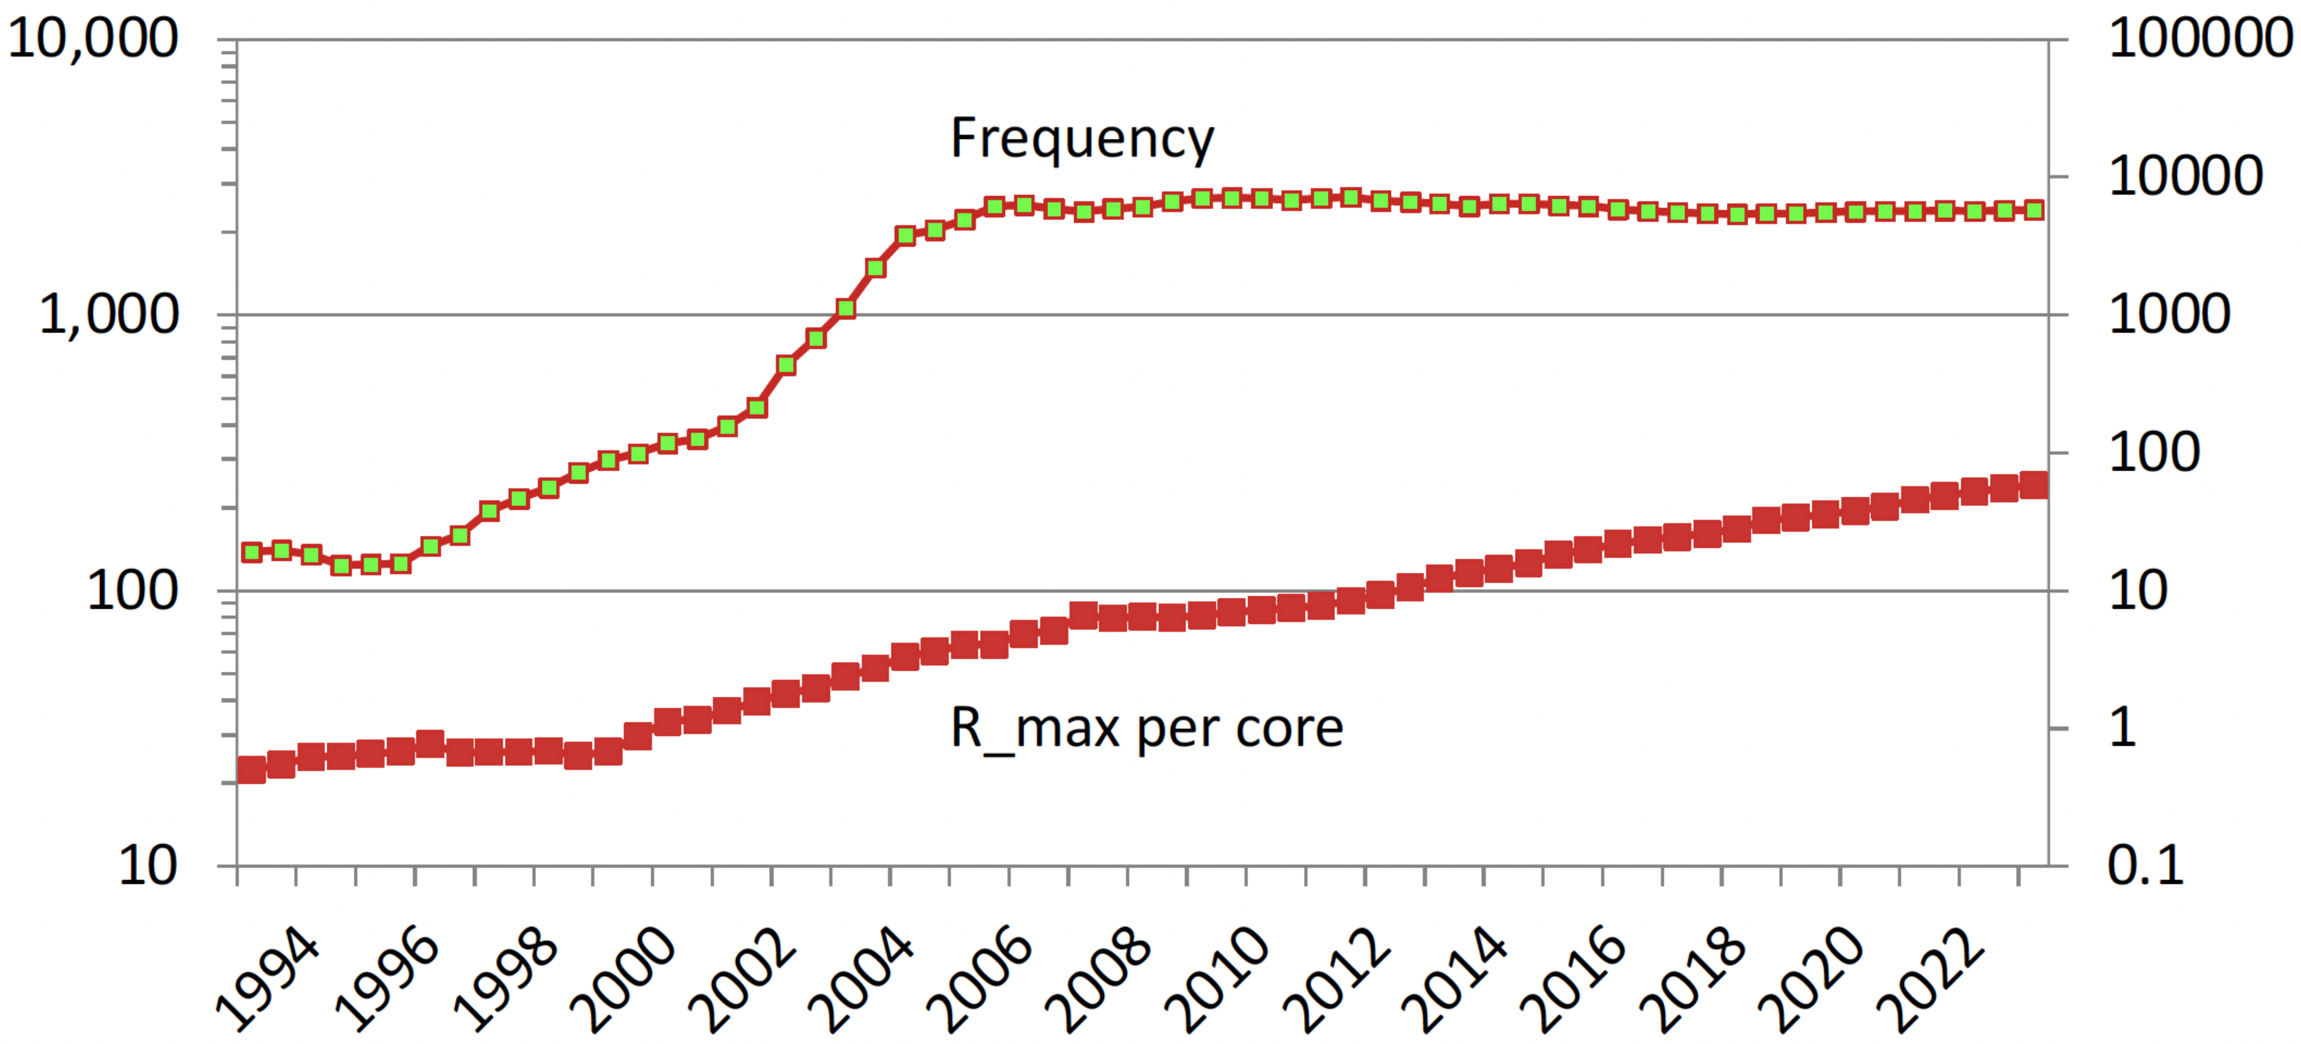

R_MAX PER CPU SOCKET – Jun'23

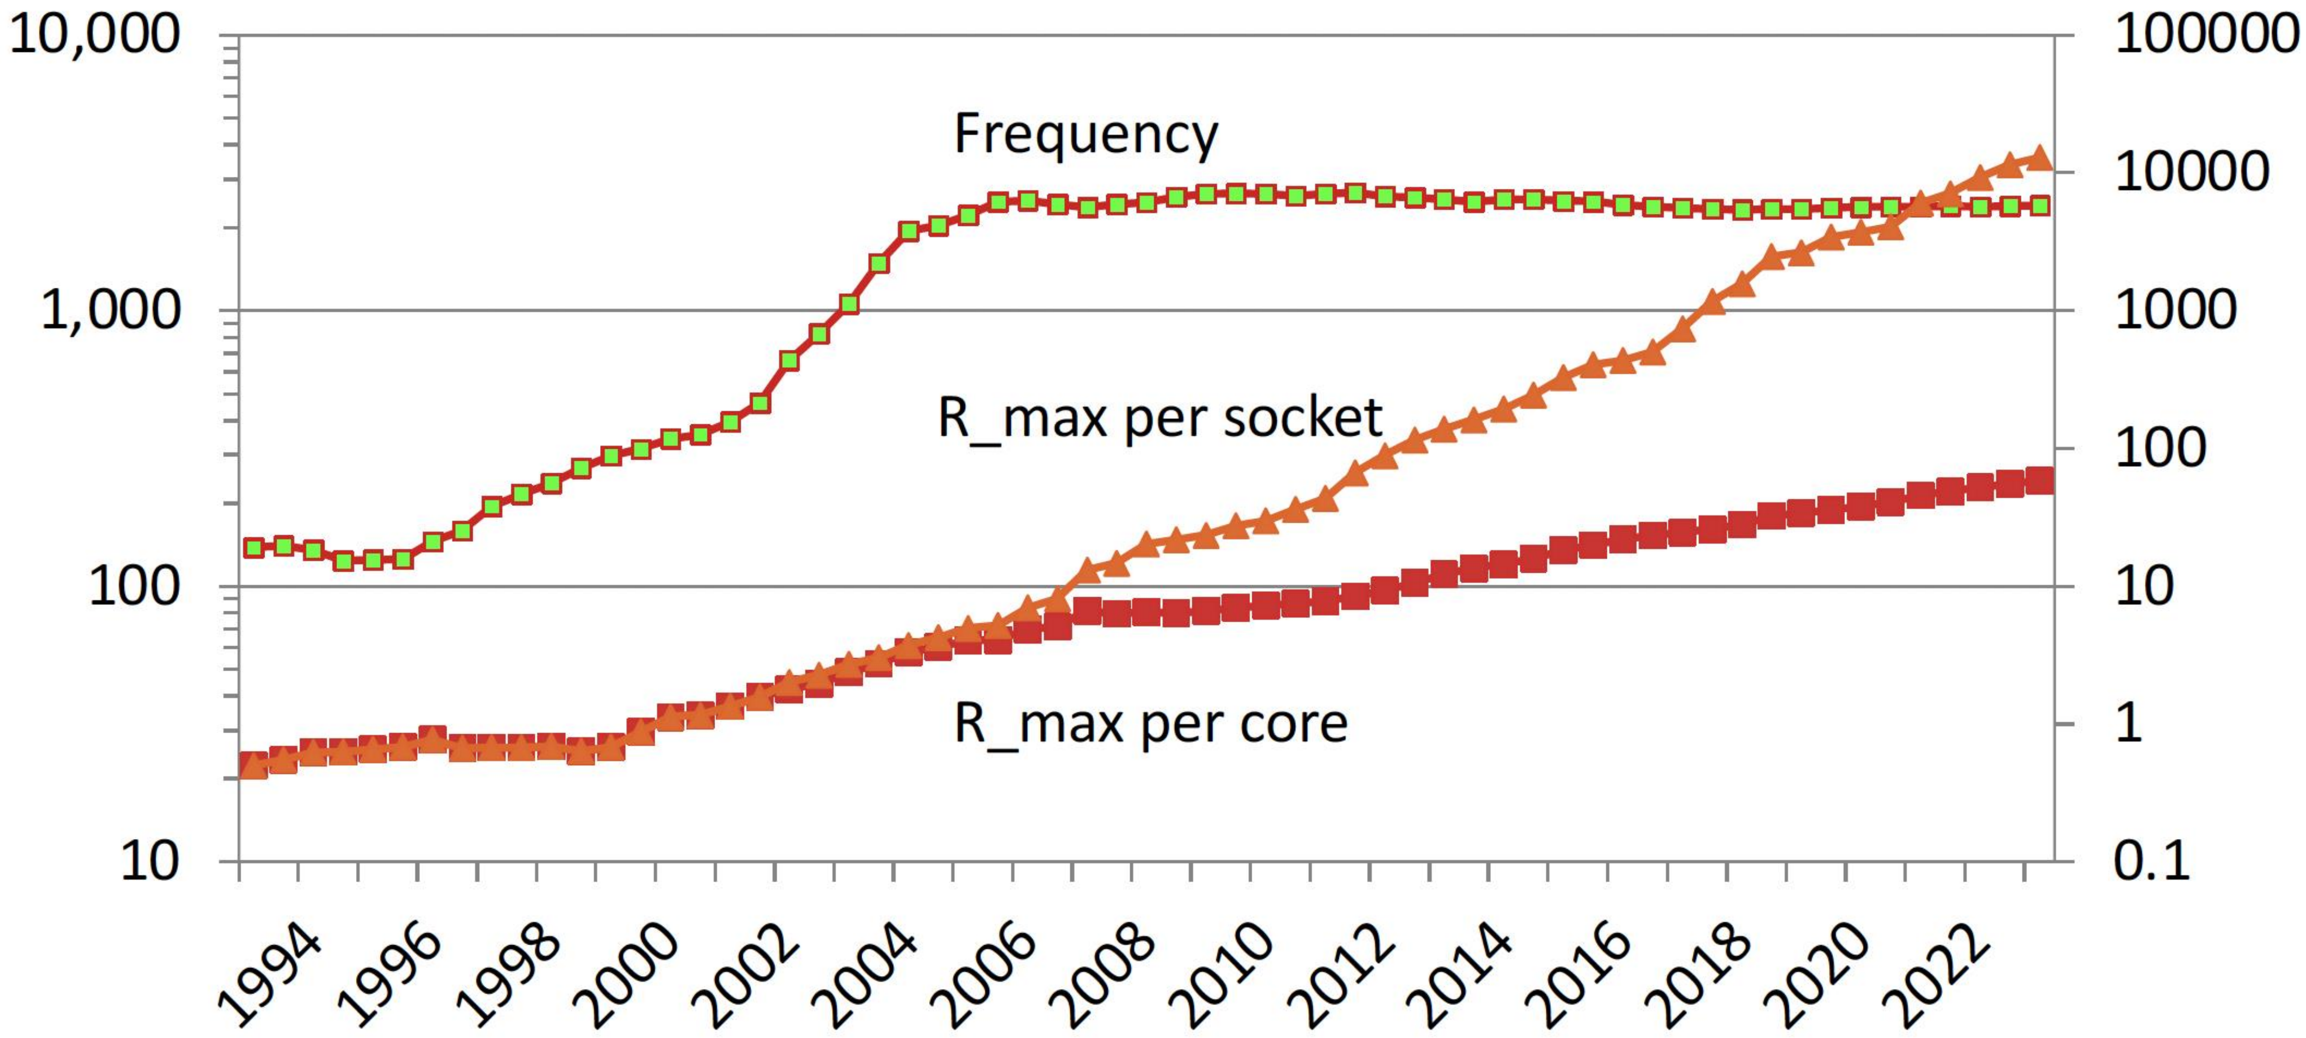

Top500 Vendor/System Share – Nov'24

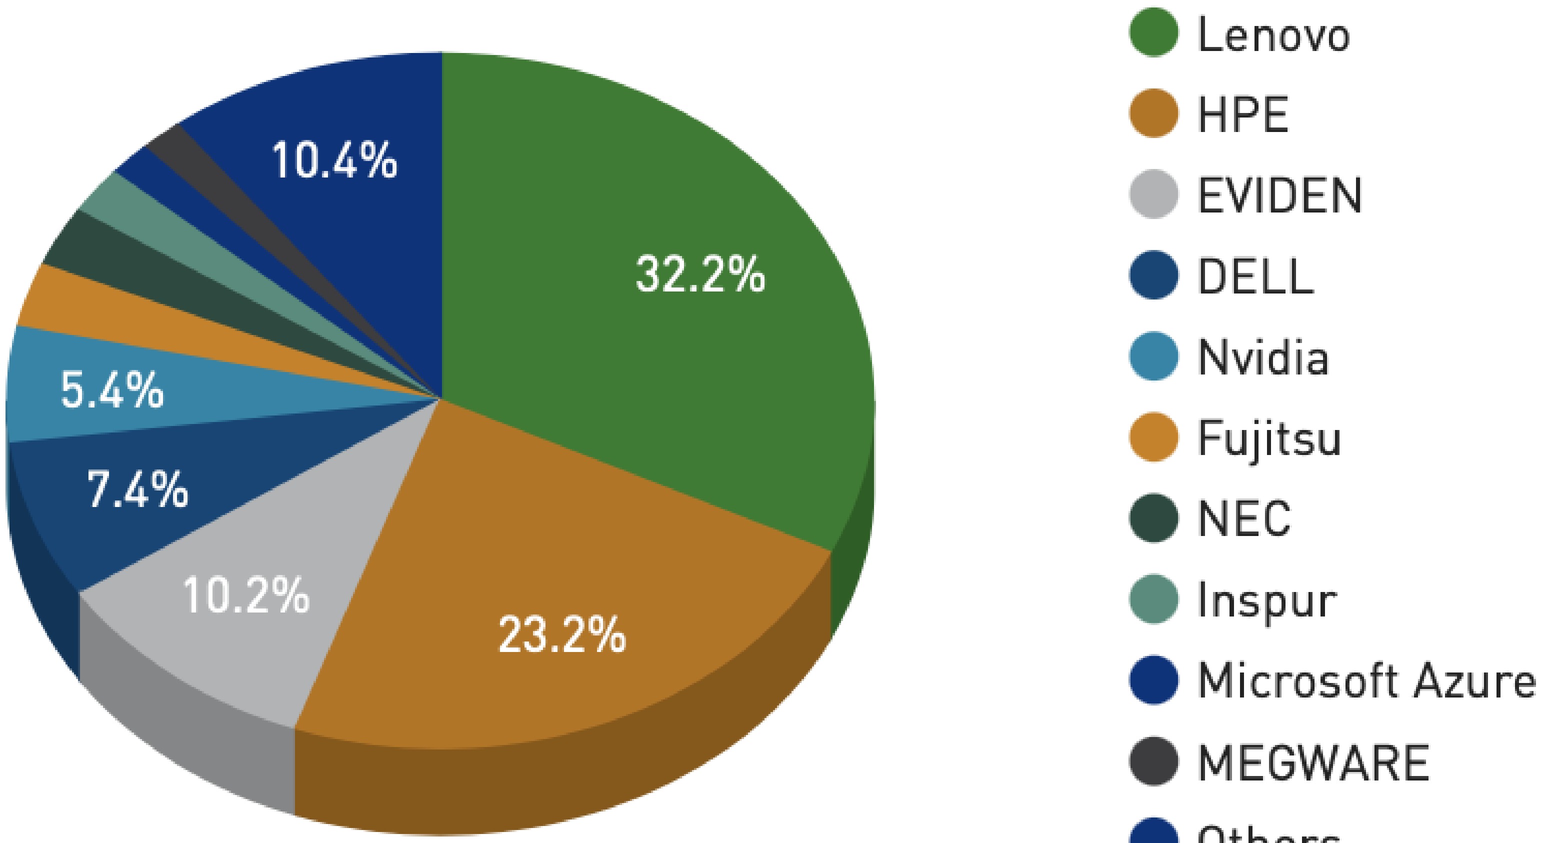

Top500 Vendor Performance Share – Nov’24

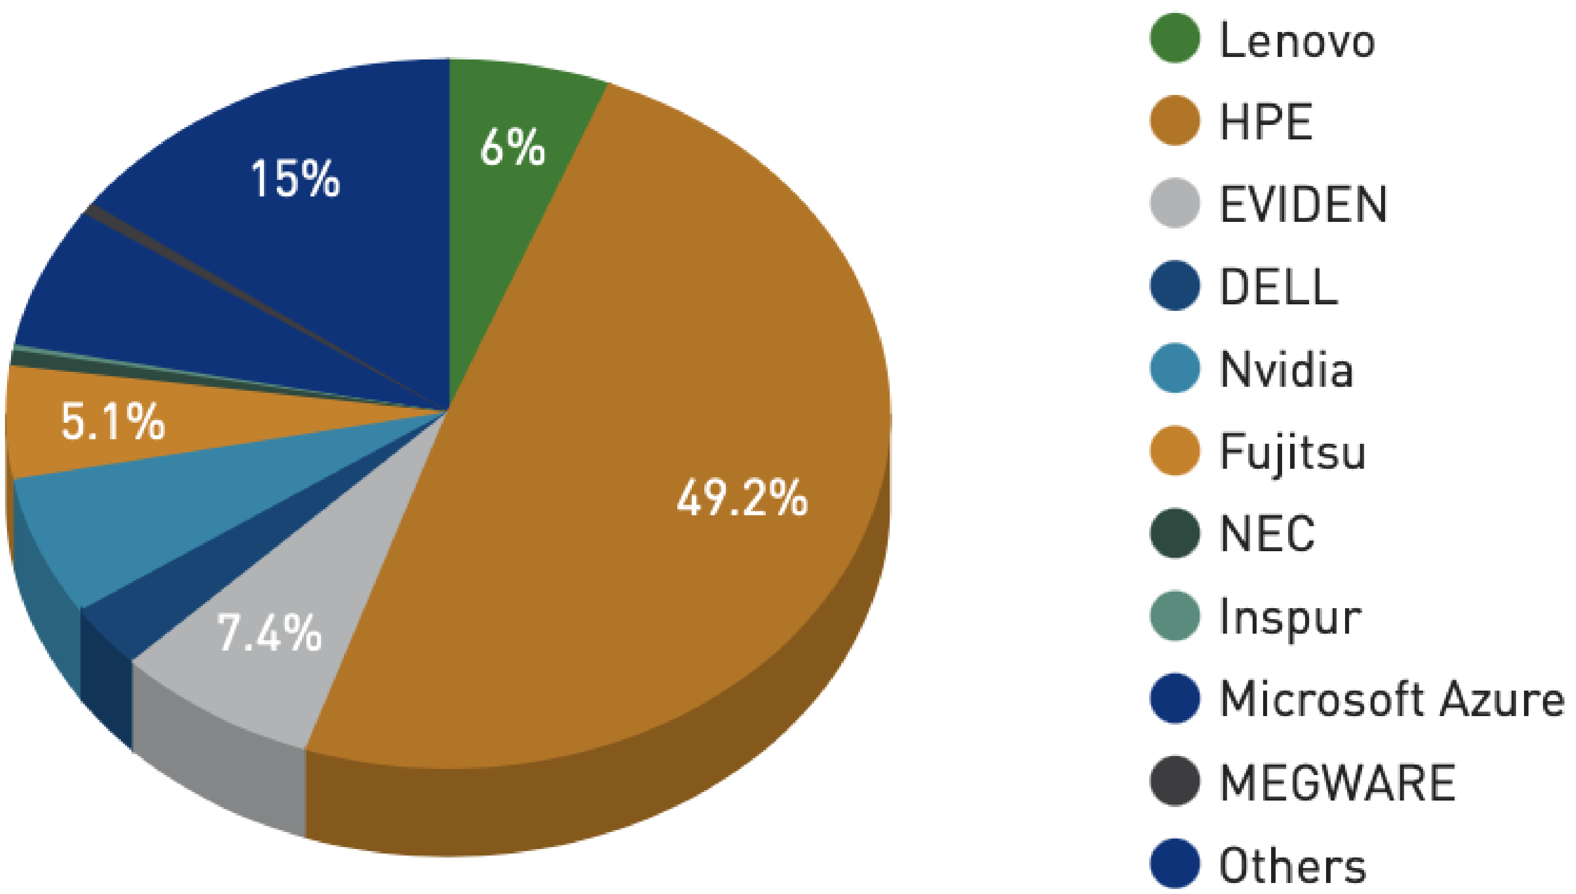

Top500 Accelerator/Co-Processor System Share – Nov’24

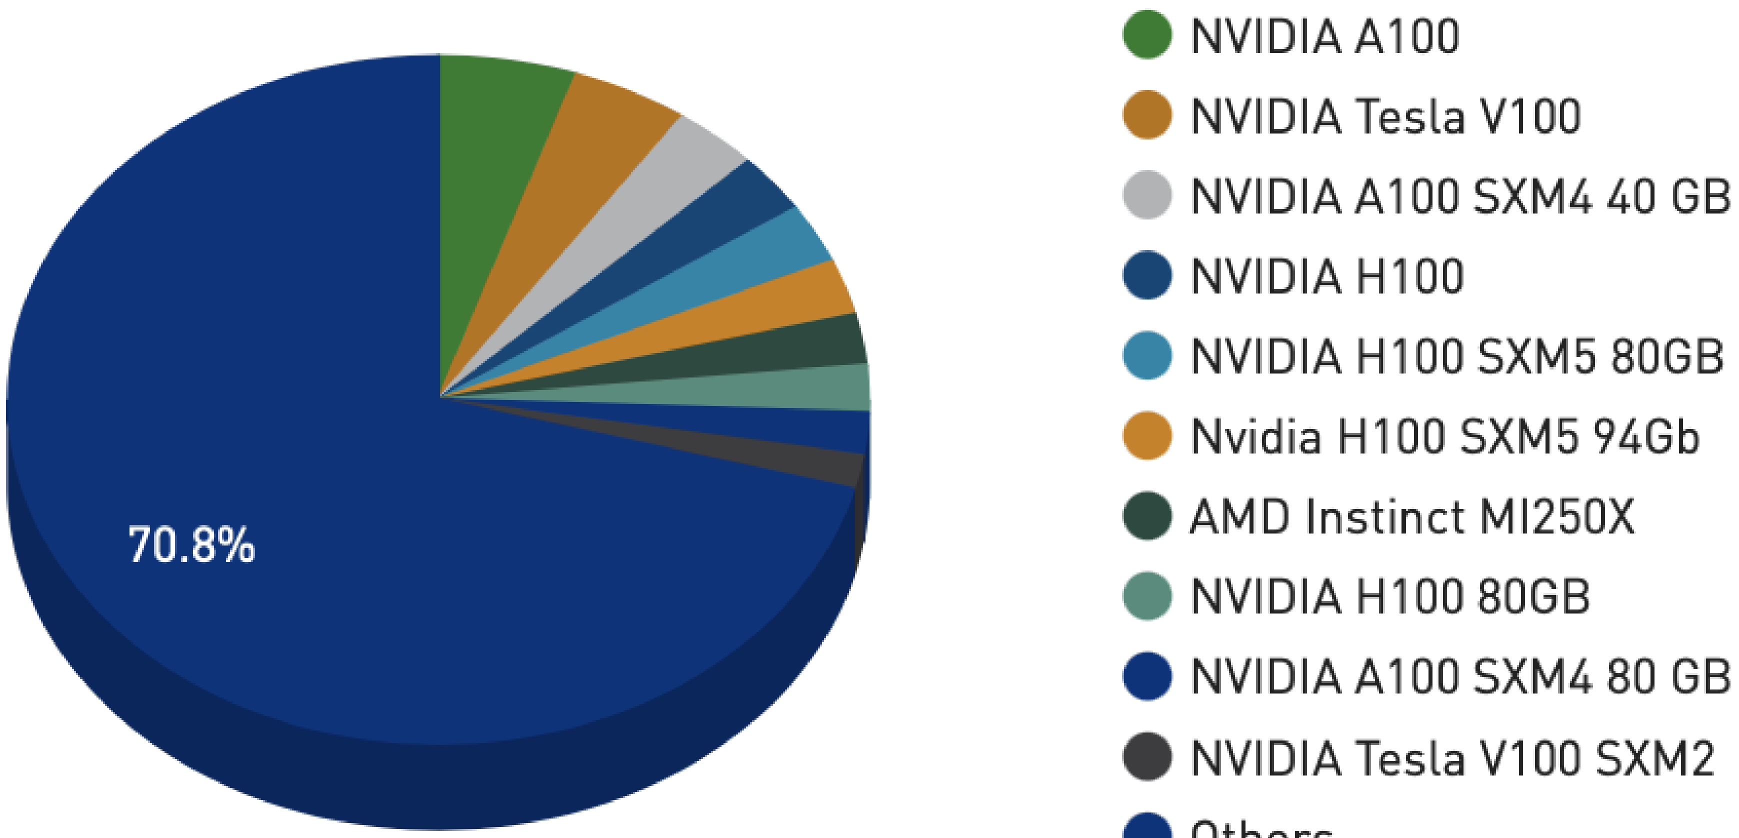

Top500 Accelerator/Co-Processor Performance Share – Nov'24

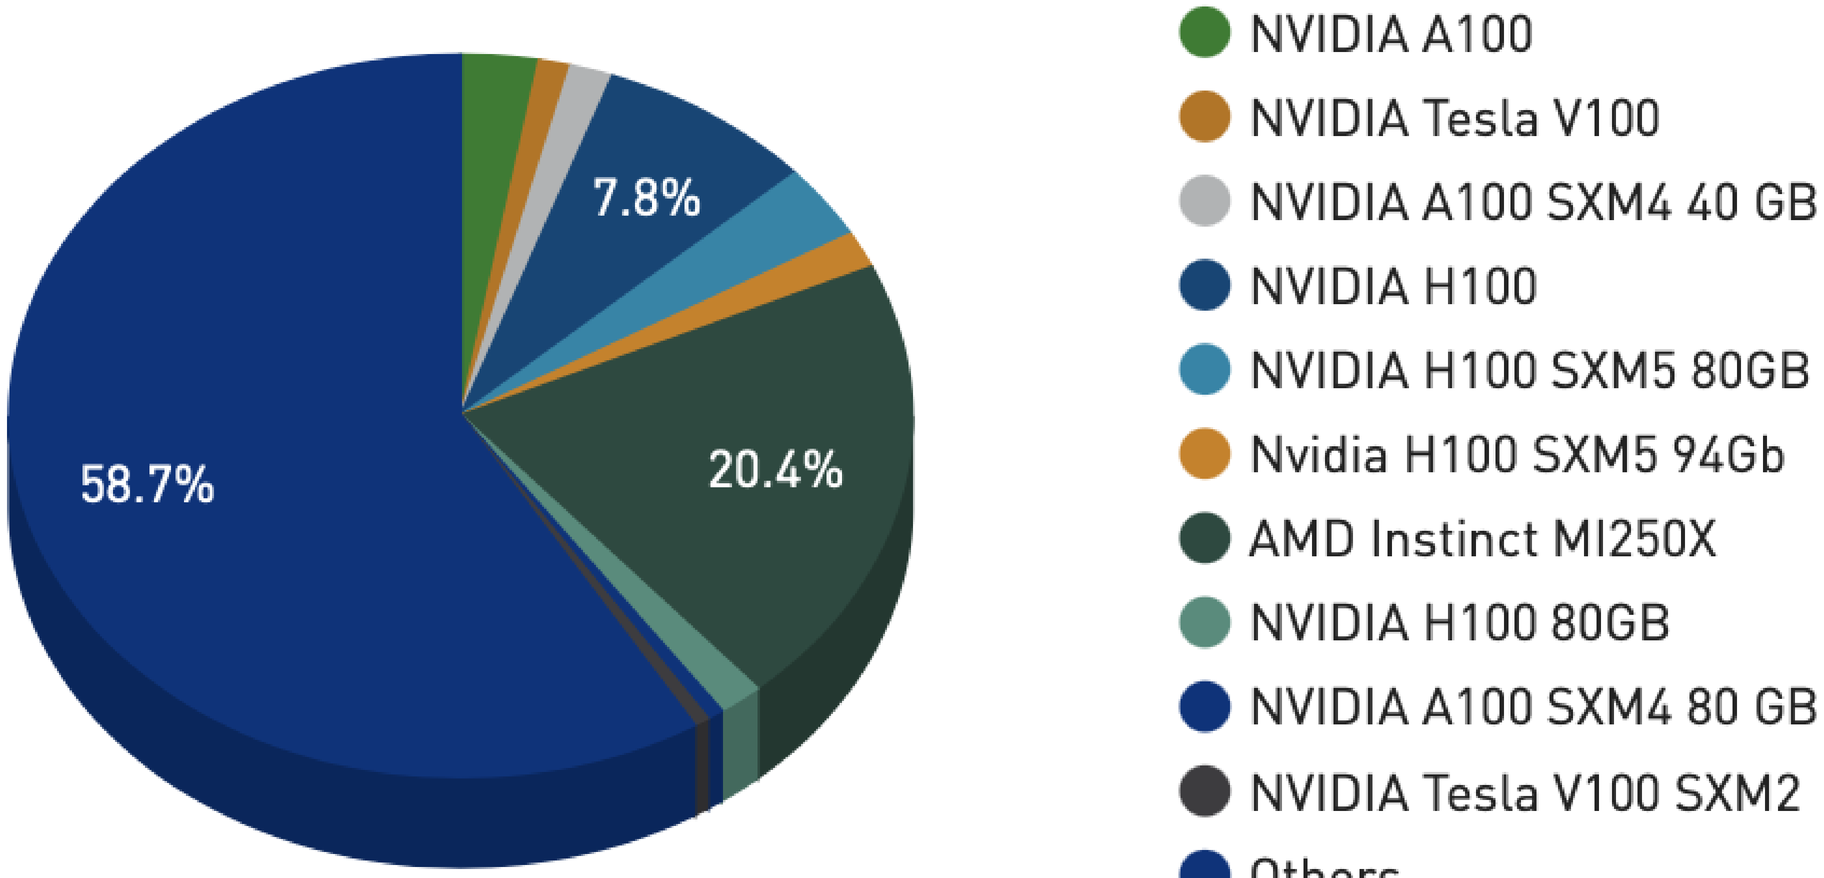

Projected Performance Development

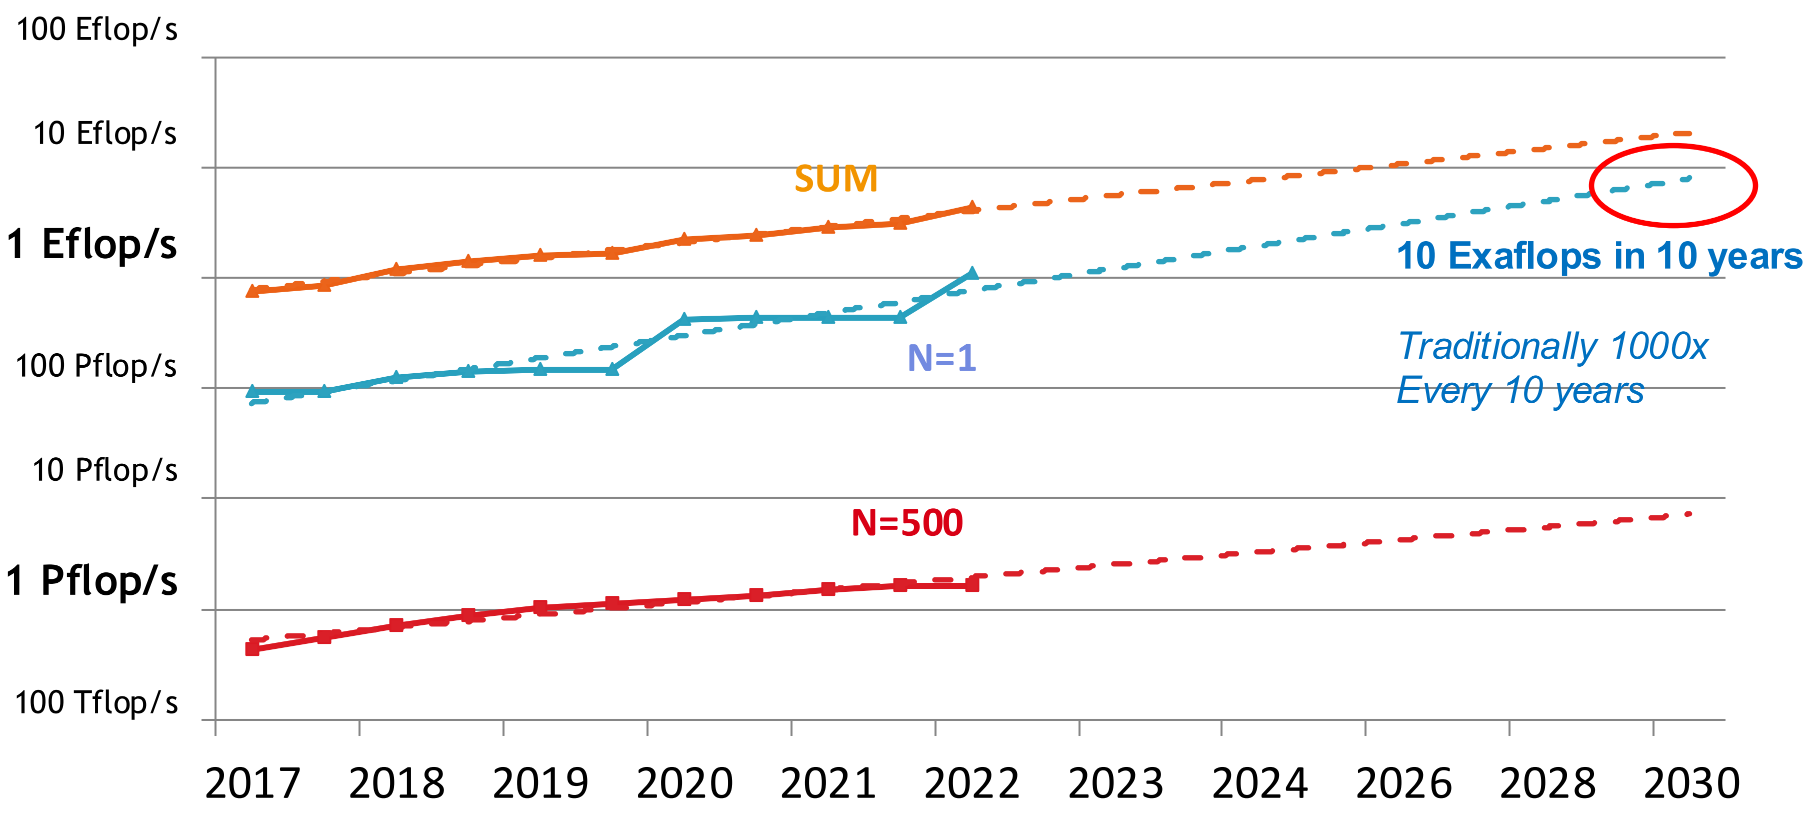

The rise of commercial off-the-shelf microprocessor technologies

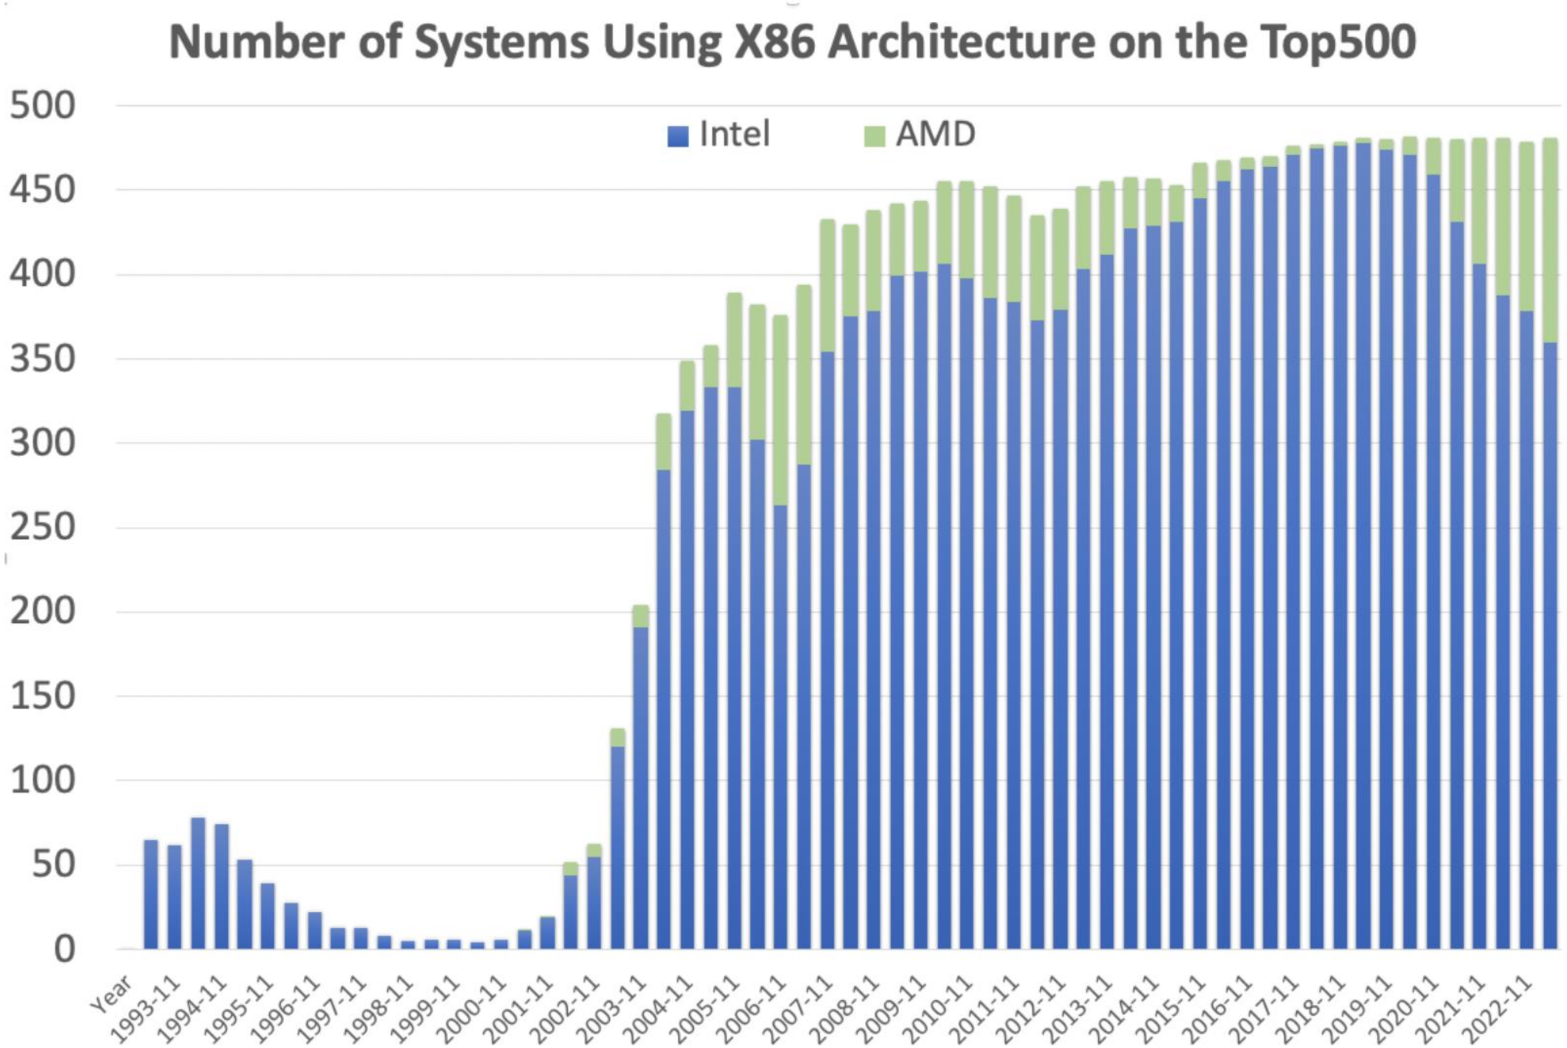

The economic context in 2023 has changed dramatically

Learning from Attack of the Killer Micros: more about the economic model than technology alone
John Markoff, May 6, 1991

… Hyperscalers are dominant in the market!

Dan Reed, 2022
https://arxiv.org/pdf/2203.02544.pdf

$4.23 trillion!!!

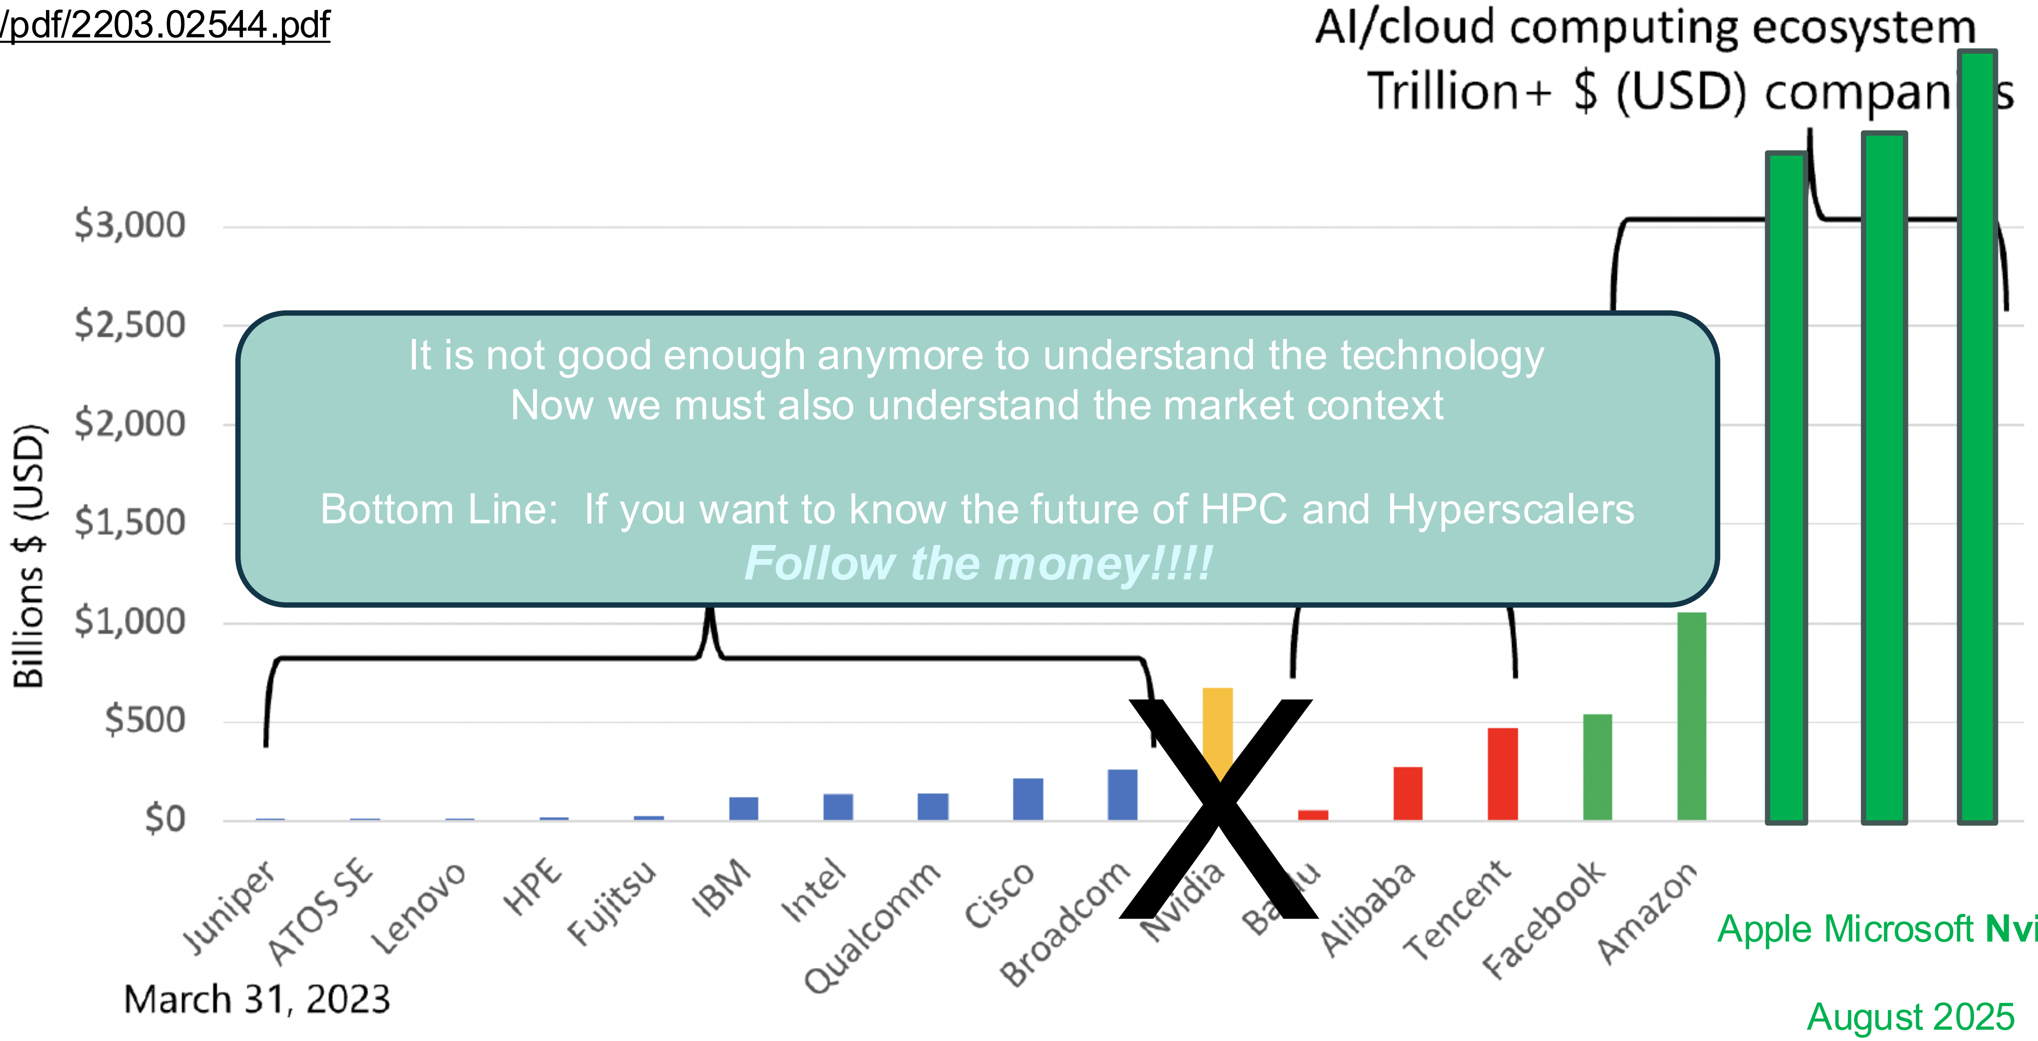

It is not good enough anymore to understand the technology
Now we must also understand the market context

Bottom Line: If you want to know the future of HPC and Hyperscalers
Follow the money!!!!

Apple Microsoft **Nvidia**

August 2025

HPCG: The Benchmark for Scientific Applications
hpcg-benchmark.org w/ J. Dongarra, P. Luszczek and M. Heroux

- High Performance Conjugate Gradients (HPCG).
- Solves *Ax=b*, *A* large, sparse, *b* known, *x* computed.
- An optimized implementation of PCG contains essential computational and communication patterns that are prevalent in a variety of methods for discretization and numerical solution of PDEs

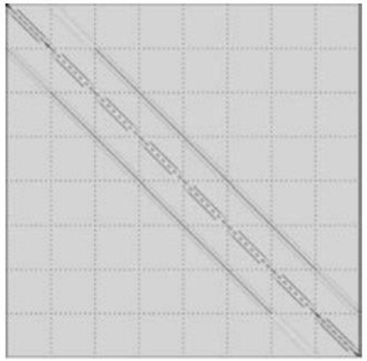

- Patterns:
  - Dense and sparse computations.
  - Dense and sparse collectives.
  - Multi-scale execution of kernels via MG (truncated) V cycle.
  - Data-driven parallelism (unstructured sparse triangular solves).
- Strong verification (via spectral properties of PCG).

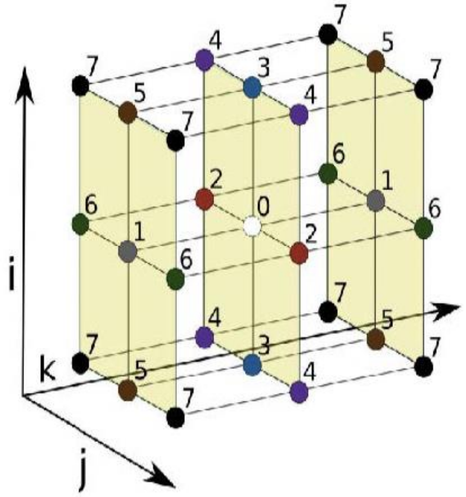

27-point stencil operator

HPCG Top 10, June 2023

#	T	Site	Manufacturer	Computer	Country	HPCG [Pflop/s]	HPL [Pflop/s]	HPCG/ Peak	HPCG/ HPL
1	2	RIKEN-CCS	Fujitsu	Fugaku Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D	Japan	16.005	442.0	3.0%	3.6%
2	1	Oak Ridge National Laboratory	HPE	Frontier HPE Cray EX235a, AMD EPYC 64C 2.0GHz, Instinct MI250X, Slingsh.-10	USA	14.054	1,194.0	0.8%	1.2%
3	3	EuroHPC / CSC	HPE	LUMI HPE Cray EX235a, AMD EPYC 64C 2.0GHz, Instinct MI250X, Slingsh.-10	Finland	3.408	309.1	0.8%	1.1%
4	4	EuroHPC / CINECA	Atos	Leonardo Atos BullSequana XH2000, Xeon 32C 2.6GHz, NVIDIA A100, HDR Infiniband	Italy	3.114	238.7	1.0%	1.3%
5	5	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.07 GHz, Volta GV100, EDR	USA	2.926	148.6	1.5%	2.0%
6	8	NERSC - Lawrence Berkeley National Laboratory	HPE	Perlmutter HPE Cray EX235n, AMD EPYC 64C 2.45GHz, NVIDIA A100, Slingshot-10	USA	1.905	70.9	2.0%	2.7%
7	6	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1 GHz, Volta GV100, EDR	USA	1.796	94.6	1.4%	1.9%
8	9	NVIDIA Corporation	NVIDIA	Selene DGX A100 SuperPOD, AMD 64C 2.25GHz, NVIDIA A100, Mellanox HDR	USA	1.623	63.5	2.1%	2.6%
9	13	Forschungszentrum Jülich (FZJ)	Atos	JUWELS Booster Module BullSequana XH2000, AMD EPYC 24C 2.8GHz, NVIDIA A100, Mell. HDR	Germany	1.275	44.1	1.8%	2.9%
10	23	Saudi Aramco	HPE	Dammam-7 Cray CS-Storm, Xeon 20C 2.5GHz, NVIDIA T. V100, IB HDR 100	Saudi Arabia	0.881	22.4	1.6%	3.9%

Number of Systems with HPCG value – Jun'23

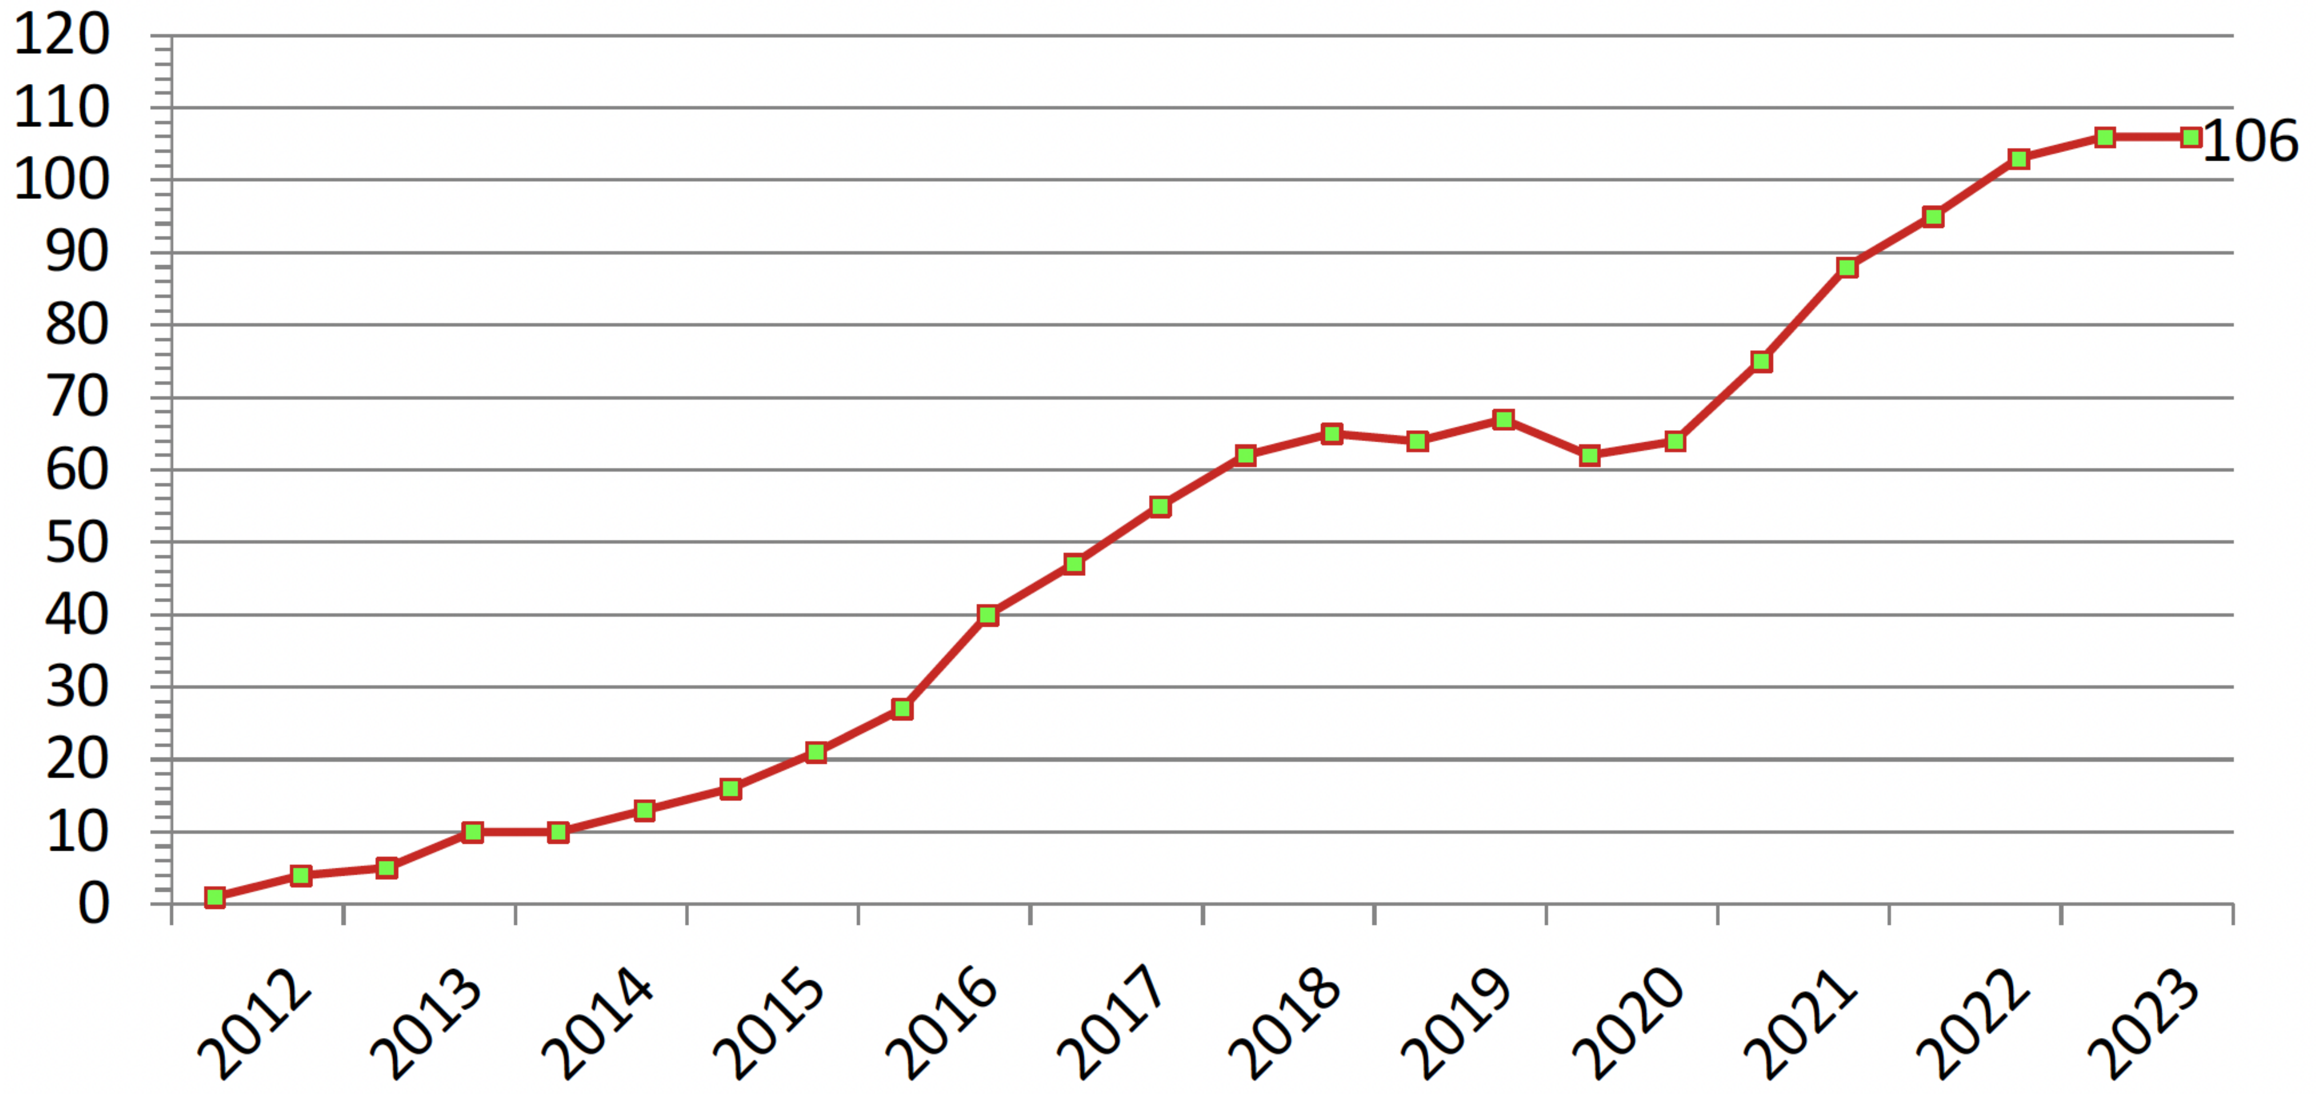

Top10: R_MAX Vs HPCG – Jun’23

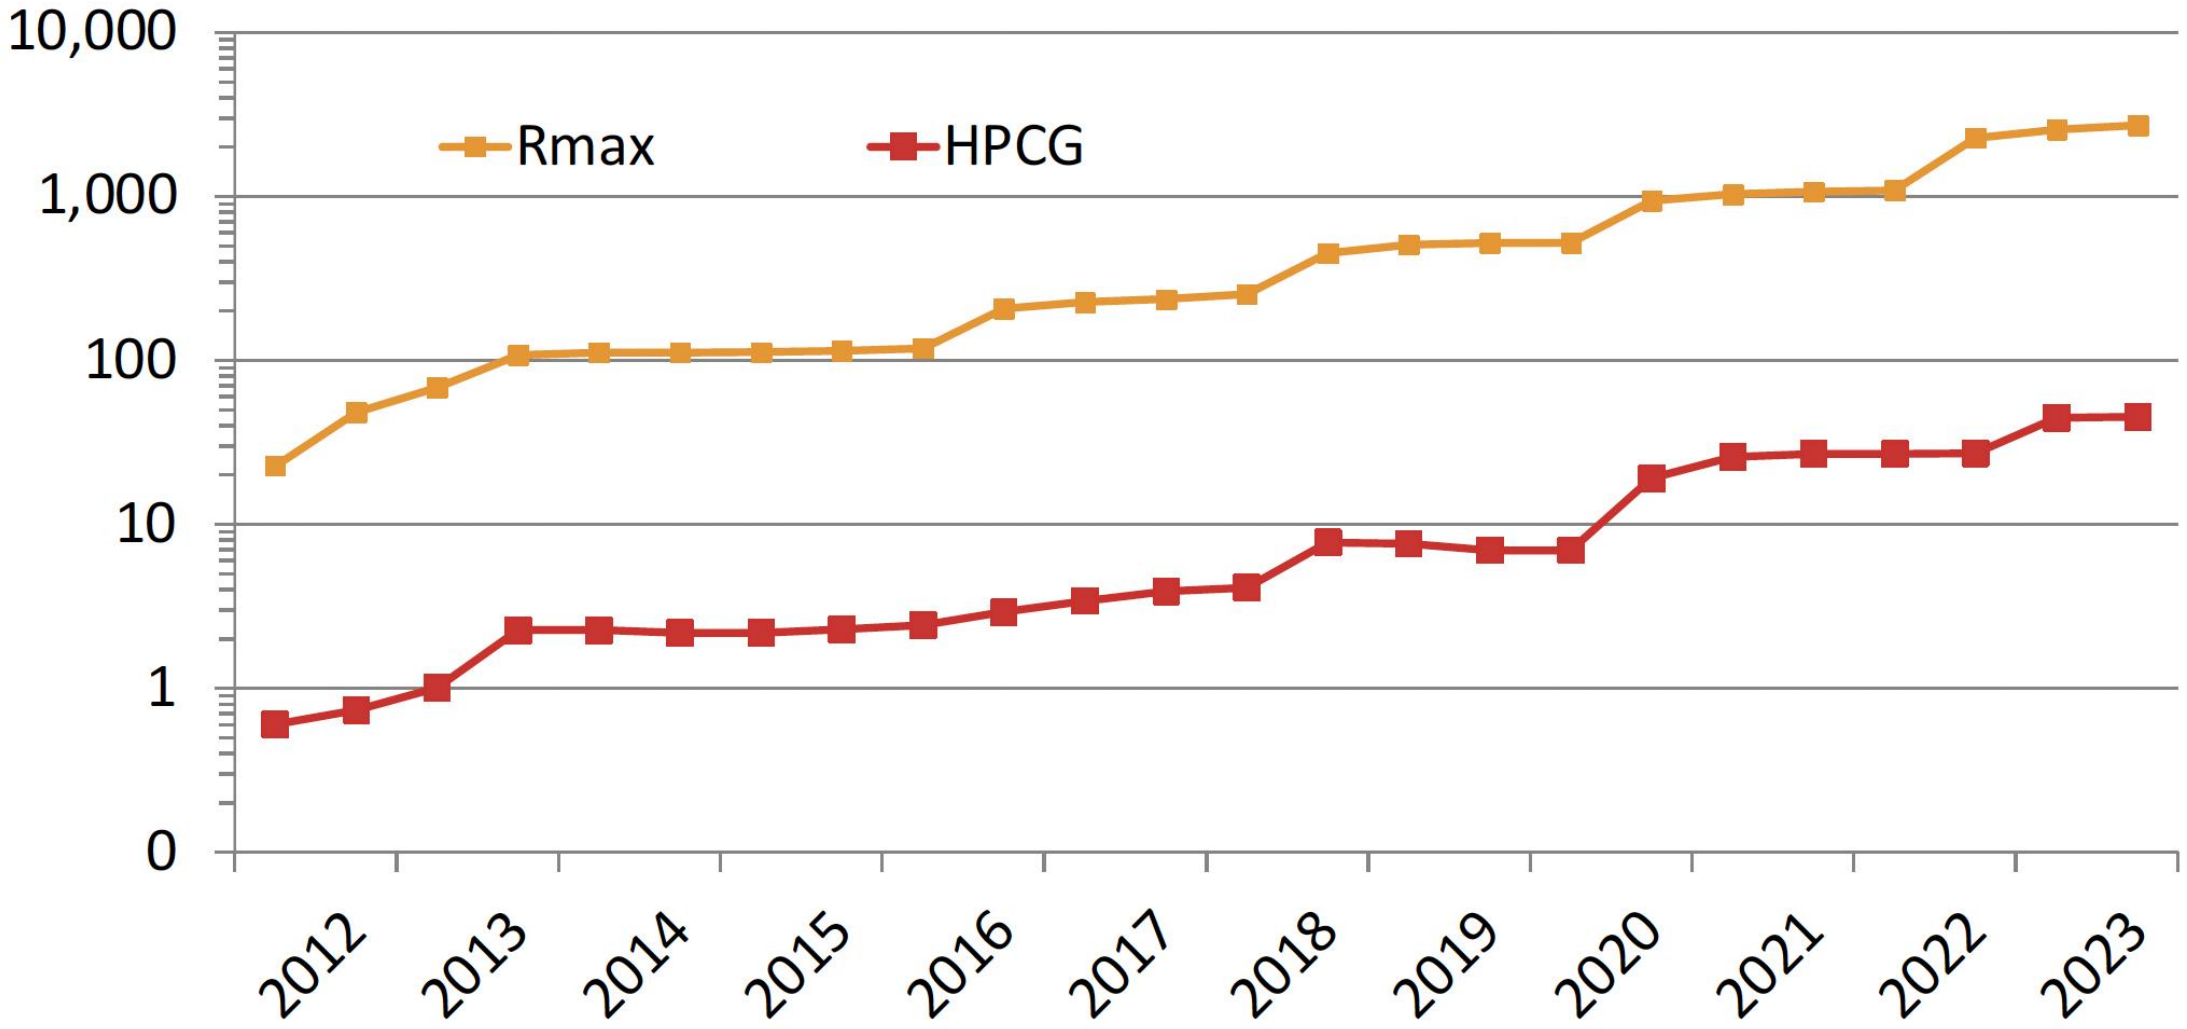

Introducing a new metric to reduce the electricity bill ☺

Improving the "science per Joule" (or per unit time) involves:

architecture

algorithm/software

application

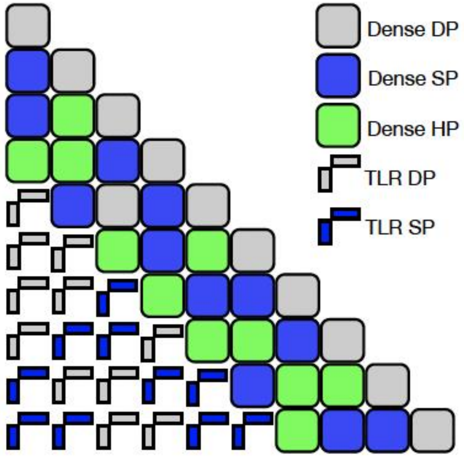

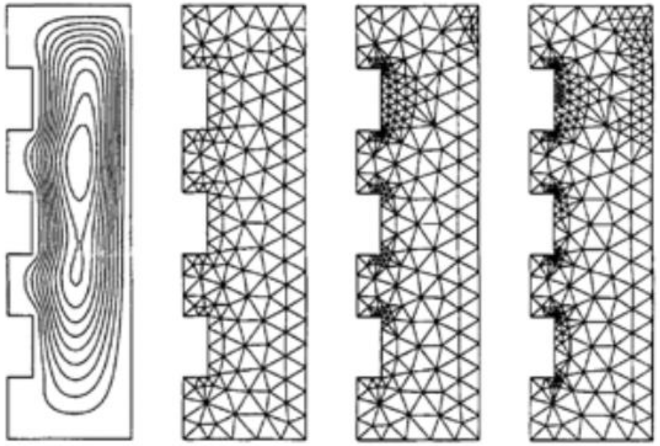

In a fortunate world, these are orthogonal: the *desired app* can employ the *best algorithm* on the most *efficient hardware.*

Hardware Landscape

AMD Epyc Genoa

High cache capacity
High memory bandwidth
x86 programming env
Memory-bound workloads

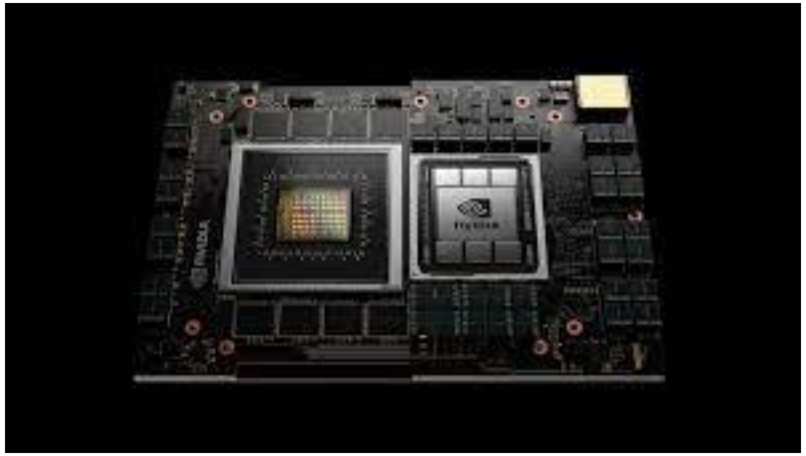

NVIDIA Grace Hopper

High speed CPU-GPU interconnect
Memory coherency
Support for mixed precisions
Compute-bound workloads

Hardware Landscape

AMD Epyc Genoa

High cache capacity
High memory bandwidth
x86 programming env
Memory-bound workloads

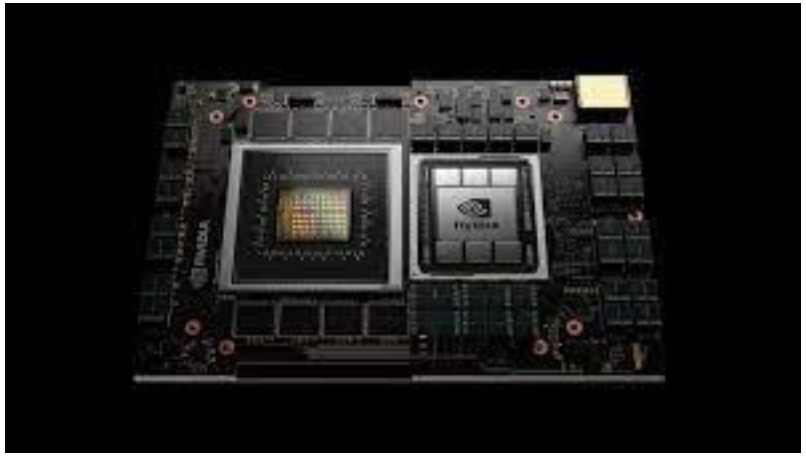

NVIDIA Grace Hopper

High speed CPU-GPU interconnect
Memory coherency
Support for mixed precisions
Compute-bound workloads

Cerebras CS-2
Graphcore IPU

AI-focused chip
Flat memory hierarchy
High SRAM bandwidth
Inference

Hardware Landscape

AMD Epyc Genoa

High cache capacity
High memory bandwidth
x86 programming env
Memory-bound workloads

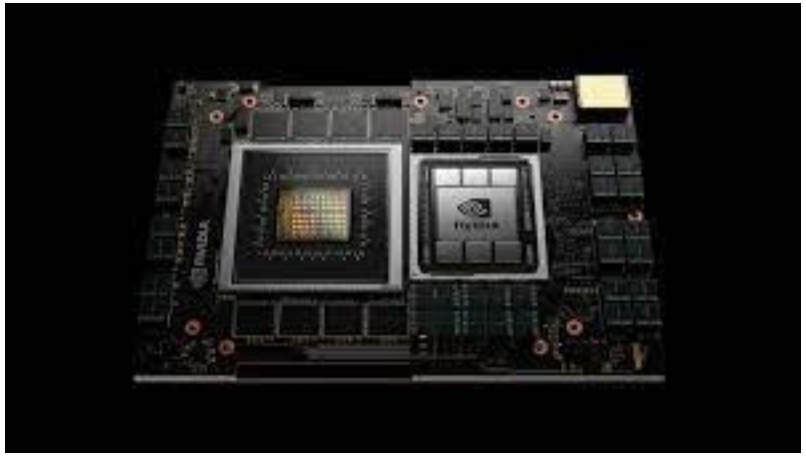

NVIDIA Grace Hopper

High speed CPU-GPU interconnect
Memory coherency
Support for mixed precisions
Compute-bound workloads

Cerebras CS-2
Graphcore IPU

AI-focused chip
Flat memory hierarchy
High SRAM bandwidth
Inference

How do we reconcile this hostile environment with our applications?