



# TOP500 and Accidental Benchmarking

New Directions  
in Numerical  
Linear Algebra  
and High  
Performance  
Computing:  
Celebrating the  
70th Birthday of  
Jack Dongarra,  
July 8, 2021

Erich  
Strohmaier

# Confessions of an Accidental Benchmarker

- Appendix B of the Linpack Users' Guide
  - Designed to help users extrapolate execution Linpack software package
- First benchmark report from 1977;
  - Cray 1 to DEC PDP-10



# "APPENDIX B"

UNIT = 10\*\*6 TIME/( 1/3 100\*\*3 + 100\*\*2 )

*2/3 N^3 ops time*

Facility	TIME N=100 secs.	UNIT micro- secs.	Computer	Type	Compiler
NCAR	14.8 .049	0.14	CRAY-1	S	CFT, Assembly BLAS
LASL	4.64 .148	0.43	CDC 7600	S	FTN, Assembly BLAS
NCAR	3.54 .192	0.56	CRAY-1	S	CFT
LASL	3.27 .210	0.61	CDC 7600	S	FTN
Argonne	2.31 .297	0.86	IBM 370/195	D	H
NCAR	1.91 .359	1.05	CDC 7600	S	Local
Argonne	1.77 .388	1.33	IBM 3033	D	H
NASA Langley	1.40 .489	1.42	CDC Cyber 175	S	FTN
U. Ill. Urbana	1.86 .506	1.47	CDC Cyber 175	S	Ext. 4.6
LLL	1.24 .554	1.61	CDC 7600	S	CHAT, No optimize
SLAC	1.19 .579	1.69	IBM 370/168	D	H Ext., Fast mult.
Michigan	1.09 .631	1.84	Amdahl 470/V6	D	H
Toronto	.772 .890	2.59	IBM 370/165	D	H Ext., Fast mult.
Northwestern	.477 1.44	4.20	CDC 6600	S	FTN
Texas	.356 1.93*	5.63	CDC 6600	S	RUN
China Lake	.352 1.95*	5.69	Univac 1110	S	V
Yale	.265 2.59	7.53	DEC KL-20	S	F20
Bell Labs	.197 3.46	10.1	Honeywell 6080	S	Y
Wisconsin	.197 3.49	10.1	Univac 1110	S	V
Iowa State	.194 3.54	10.2	Itel AS/5 mod3	D	H
U. Ill. Chicago	.148 4.10	11.9	IBM 370/158	D	G1
Purdue	.124 5.69	16.6	CDC 6500	S	FUN
U. C. San Diego	.062 13.1	38.2	Burroughs 6700	S	H
Yale	.047 17.1*	49.9	DEC KA-10	S	F40


\* TIME(100) = (100/75)\*\*3 SGEFA(75) + (100/75)\*\*2 SGESL(75)

# TOP500 – WHY HPL?



- Adaptive definition of 'Supercomputer' for collecting market statistics
- Simple metric and procedure (few rules)
- Based on measured performance (system has to function)
- Floating point benchmark ('scientific computing' in early 90s)
- High performing (optimizable) to encourage adoption
- Broad system coverage
- HPL (High Performance Linpack) had widest coverage by a factor 2-3 x at least
  - In 1993 and still !

# PERFORMANCE DEVELOPMENT



# WHY DID LINPACK WORK SO WELL?



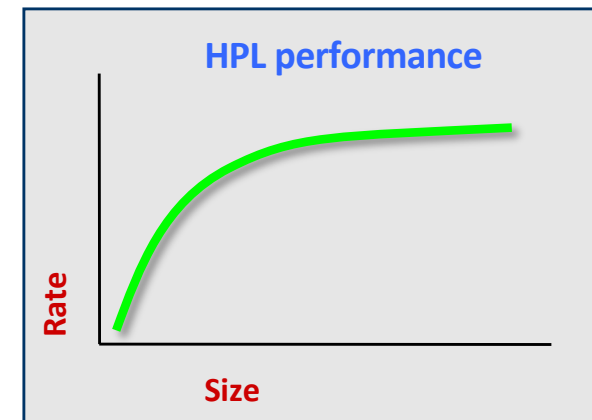
Many reasons, here are 2 essentials for the *TOP500*:

## 1) Easy and continuous scalable problem size

- Simplicity

## 2) Asymptotically best performance

- For both **system size** and **problem size**
- Brings out correct long-term trends



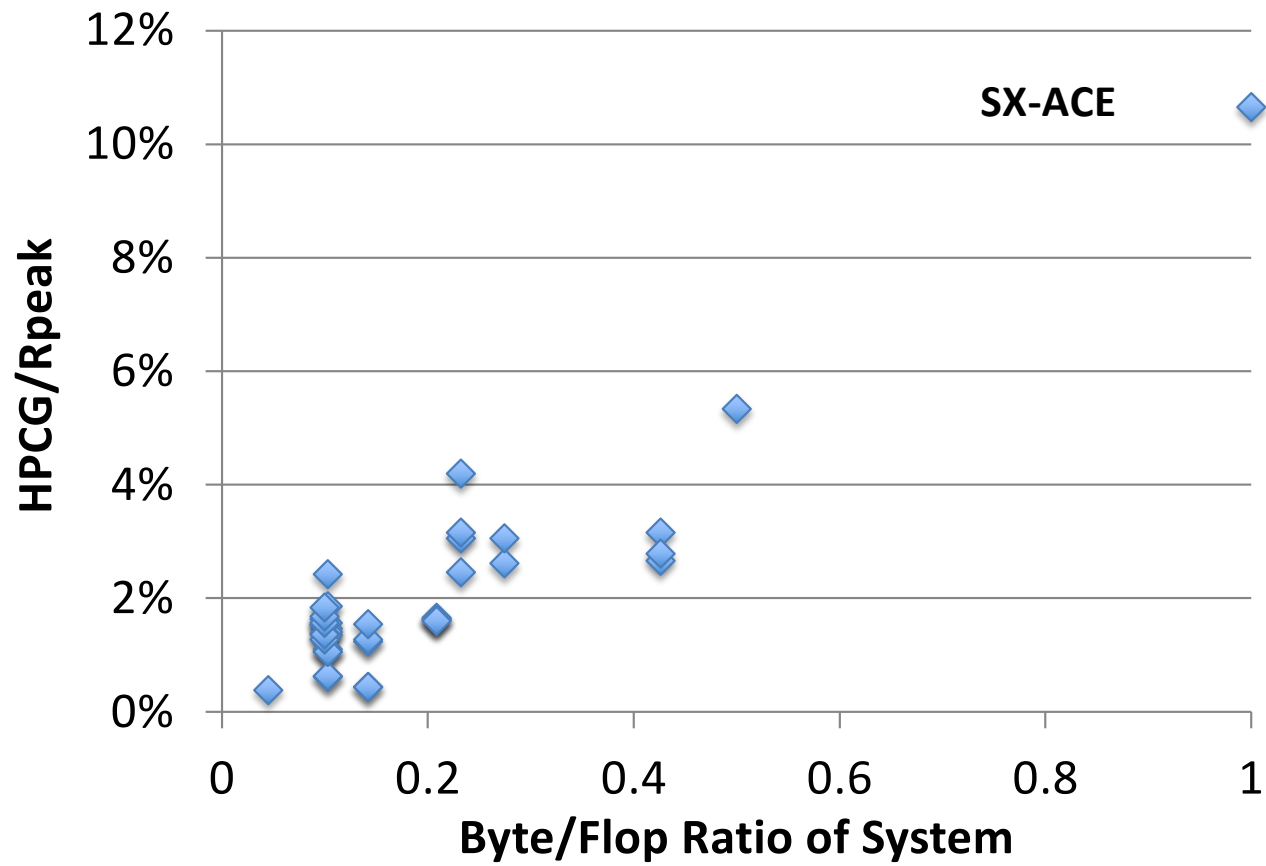
# CRITERIA FOR ADDITIONAL BENCHMARKS

---



- All the good things from HPL plus:
  - **Arithmetic Intensity: Flop/Bytes  $\sim O(1)$**
  - Main features all scale with  $O(n)$
- Does not correlate overly well with established BMs
- Changes relative rankings compared to TOP500

# HPCG - BYTES/FLOPS







# HPL-AI Benchmark Top 5 List, June 2021

Rank	Site	Computer	Cores	HPL Rmax (Eflop/s)	TOP500 Rank	HPL-AI (Eflop/s)	Speedup
1	RIKEN Center for Computational Science Japan	<b>Fugaku</b> , Fujitsu A64FX, Tofu D	7,630,848	0.442	1	2.0	4.5x
2	DOE/SC/ORNL USA	<b>Summit</b> , AC922 IBM POWER9, IB Dual-rail FDR, NVIDIA Volta V100	2,414,592	0.149	2	1.15	7.7x
3	NVIDIA USA	<b>Selene</b> , DGX SuperPOD, AMD EPYC 7742 64C 2.25 GHz, Mellanox HDR, NVIDIA A100	555,520	0.063	6	0.63	9.9x
4	DOE/SC/LBNL/NERSC USA	<b>Perlmutter</b> , HPE Cray EX235n, AMD EPYC 7763 64C 2.45 GHz, Slingshot-10, NVIDIA A100	761,856	0.065	5	0.59	9.1x
5	Forschungszentrum Juelich (FZJ) Germany	<b>JUWELS Booster Module</b> , Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, Mellanox HDR InfiniBand, NVIDIA Ampere A100, Atos	449,280	0.044	8	0.47	10x

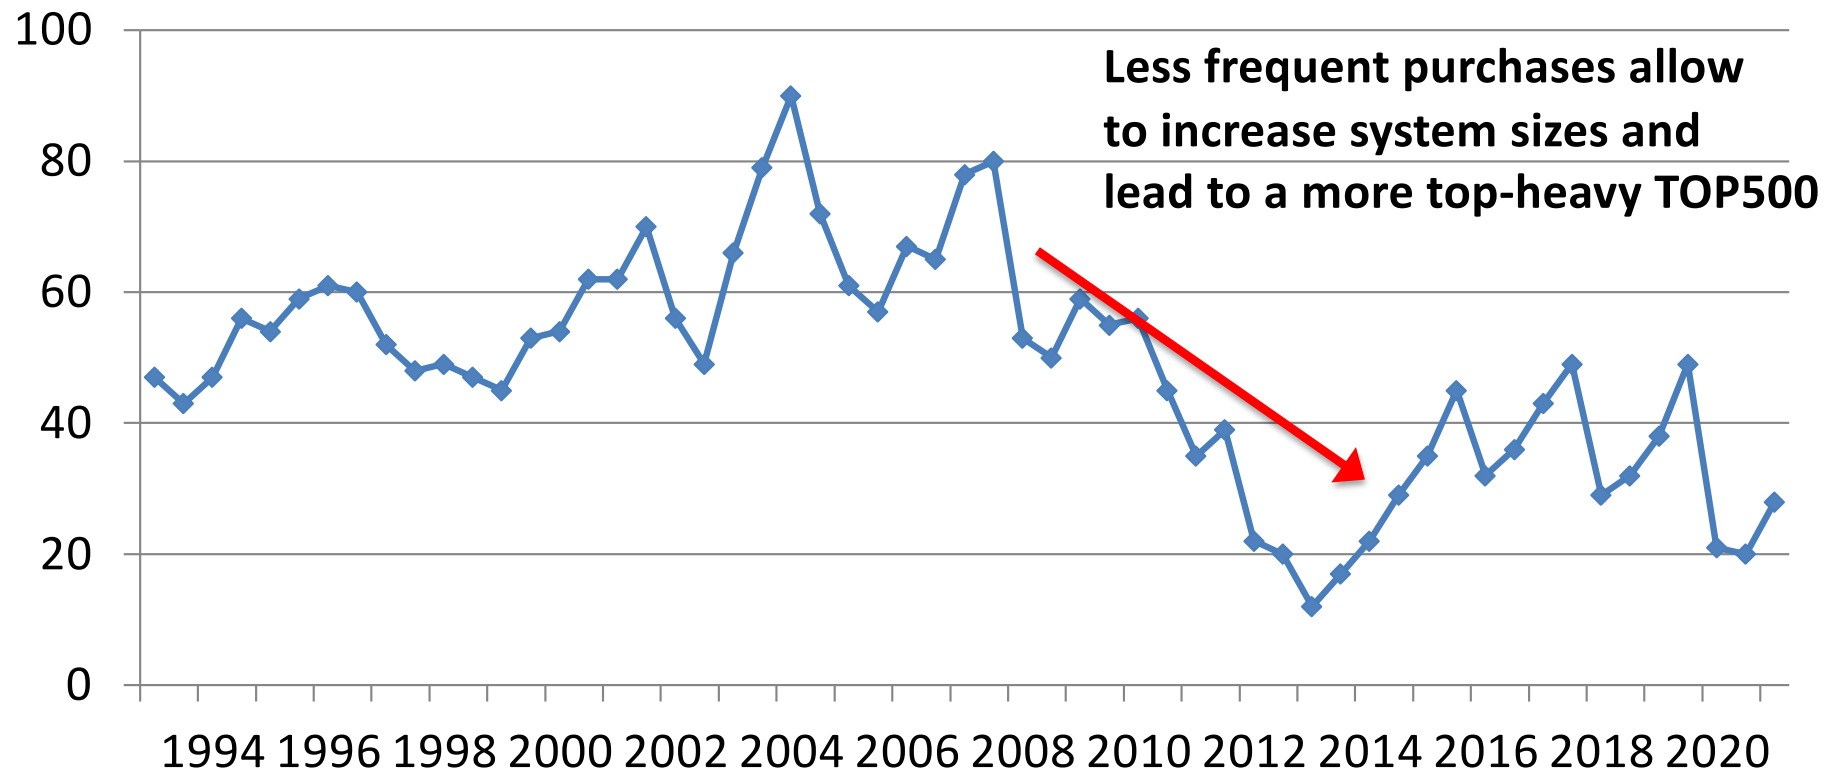
# PERFORMANCE DEVELOPMENT




# AVERAGE SYSTEM AGE



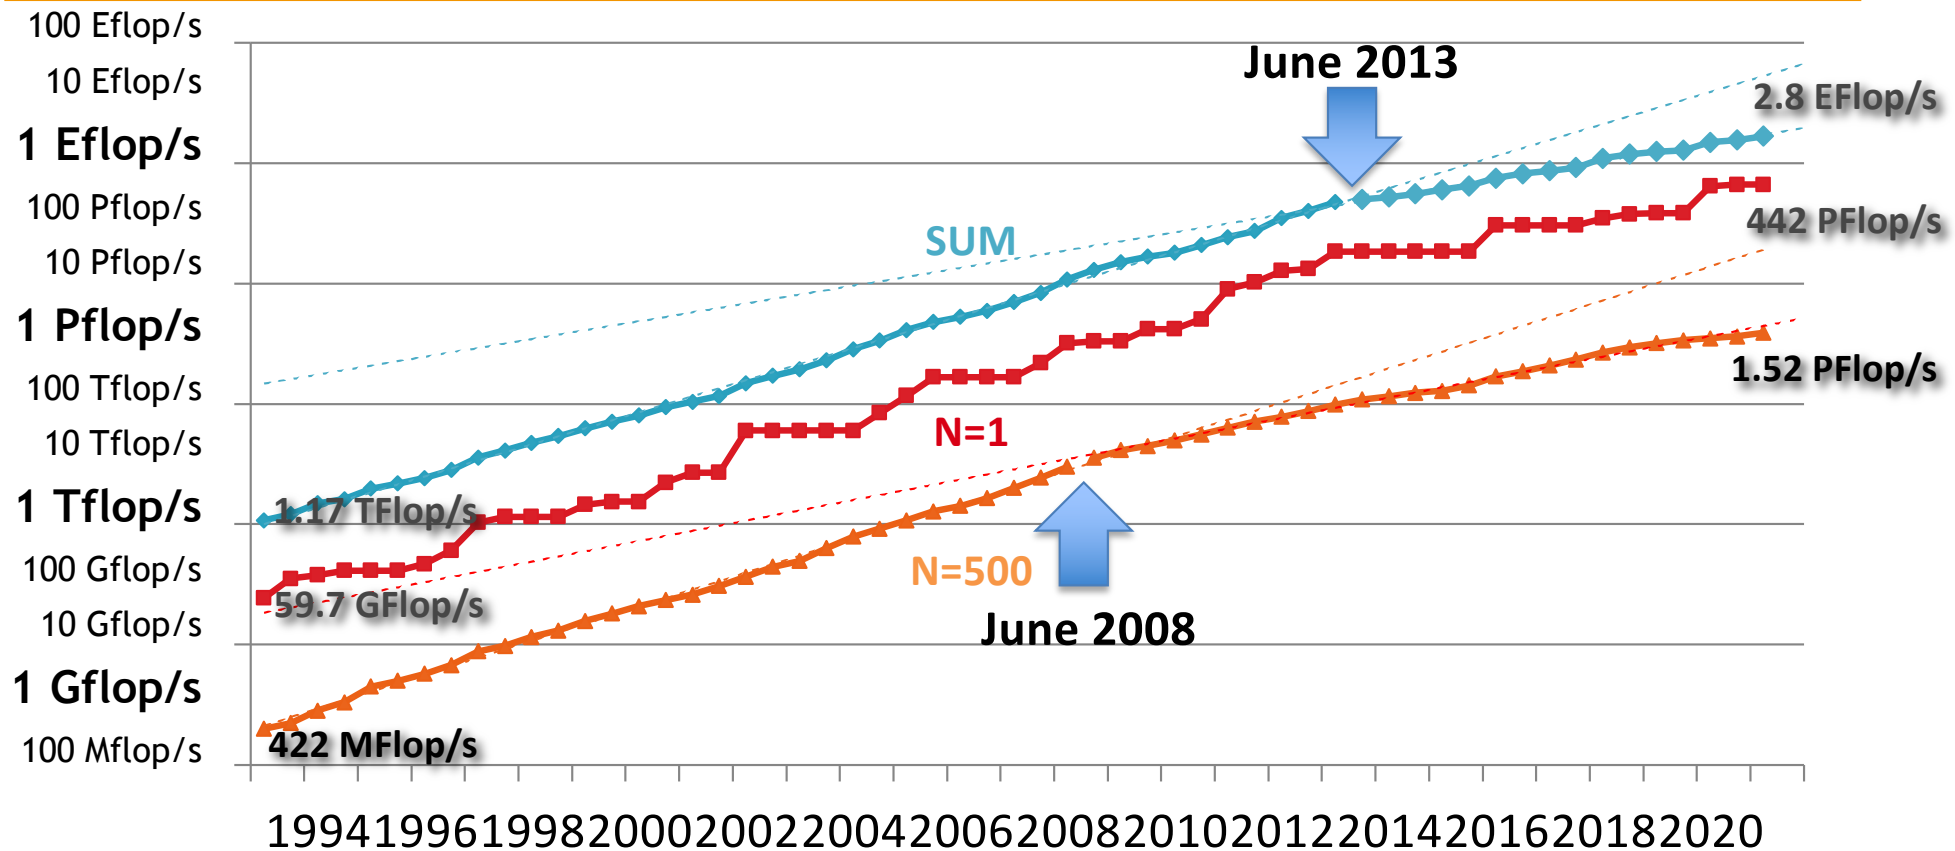
# RANK AT WHICH HALF OF TOTAL PERFORMANCE IS ACCUMULATED



# ANNUAL PERFORMANCE INCREASE OF THE TOP500



# PERFORMANCE DEVELOPMENT




## RESEARCH / COMMERCIAL MARKETS

---



- Markets for scientific computing and for commercial data processing are very different.
- Extract proper sub-samples for these markets from the full TOP500 list
  - TOP100 Research and Academic installations
  - TOP100 Industry (and Vendor) installations
    - Could try to separate out Industry installations but difficult to do
  - Ignore "Government, Classified, Others" for now
  - 100 works reasonably well, more is difficult

# ENTRY LEVEL SYSTEM SIZE





## "SUMMARY"

---

- HPL might have been an afterthought, but it does not look like an accident!
- Linear Algebra seems to be everything we need to measure HPC performance ...
- HPL served us well to analyze performance trends in the TOP500 - and continues to do so!