The TOP500 Project

Looking Back Over 16 Years of Supercomputing Experience with Special Emphasis on the Industrial Segment

Hans Werner Meuer
hans@meuer.de
Prometeus GmbH & Universität Mannheim

Microsoft Industriekunden – Veranstaltung
Frankfurt /16. Oktober 2008 / Windows HPC Server 2008 Launch

31th List: The TOP10

	Manufacturer	Computer	Rmax [TF/s]	Installation Site	Country	Power [MW]	#Cores
1	IBM	Roadrunner BladeCenter QS22/LS21	1026	DOE/NNSA/LANL	USA	2.35	122,400
2	IBM	BlueGene/L eServer Blue Gene Solution	478.2	DOE/NNSA/LLNL	USA	2.33	212,992
3	IBM	Intrepid Blue Gene/P Solution	450.3	DOE/ANL	USA	1.26	163,840
4	Sun	Ranger SunBlade x6420	326	TACC	USA	2.00	62,976
5	Cray	Jaguar Cray XT4 QuadCore	205	DOE/ORNL	USA	1.58	30,976
6	IBM	JUGENE Blue Gene/P Solution	180	Forschungszentrum Juelich (FZJ)	Germany	0.50	65,536
7	SGI	Encanto SGI Altix ICE 8200	133.2	New Mexico Computing Applications Center	USA	0.86	14,336
8	HP	EKA Cluster Platform 3000 BL460c	132.8	Computational Research Laboratories, TATA SONS	India	1.60	14,384
9	IBM	Blue Gene/P Solution	112.5	IDRIS	France	0.32	40,960
10	SGI	SGI Altix ICE 8200EX	106.1	Total Exploration Production	France	0.44	10,240

Outline

- Mannheim Supercomputer Statistics & Top500 Project Start in 1993
- Competition between Manufacturers, Countries and Sites
- My Supercomputer Favorite in the Top500 Lists
- The 31st List as of June 2008
- Performance Development and Projection
- Bell's Law
- Supercomputing, quo vadis?
- Top500, quo vadis?

Mannheim Supercomputer Statistics 1986 – 1992

- Presented at ISC’86 – ISC’92/Mannheim Supercomputer Seminars
- Counting of Supercomputers in the World
- 530 systems in the year 1992

Manufacturer Share

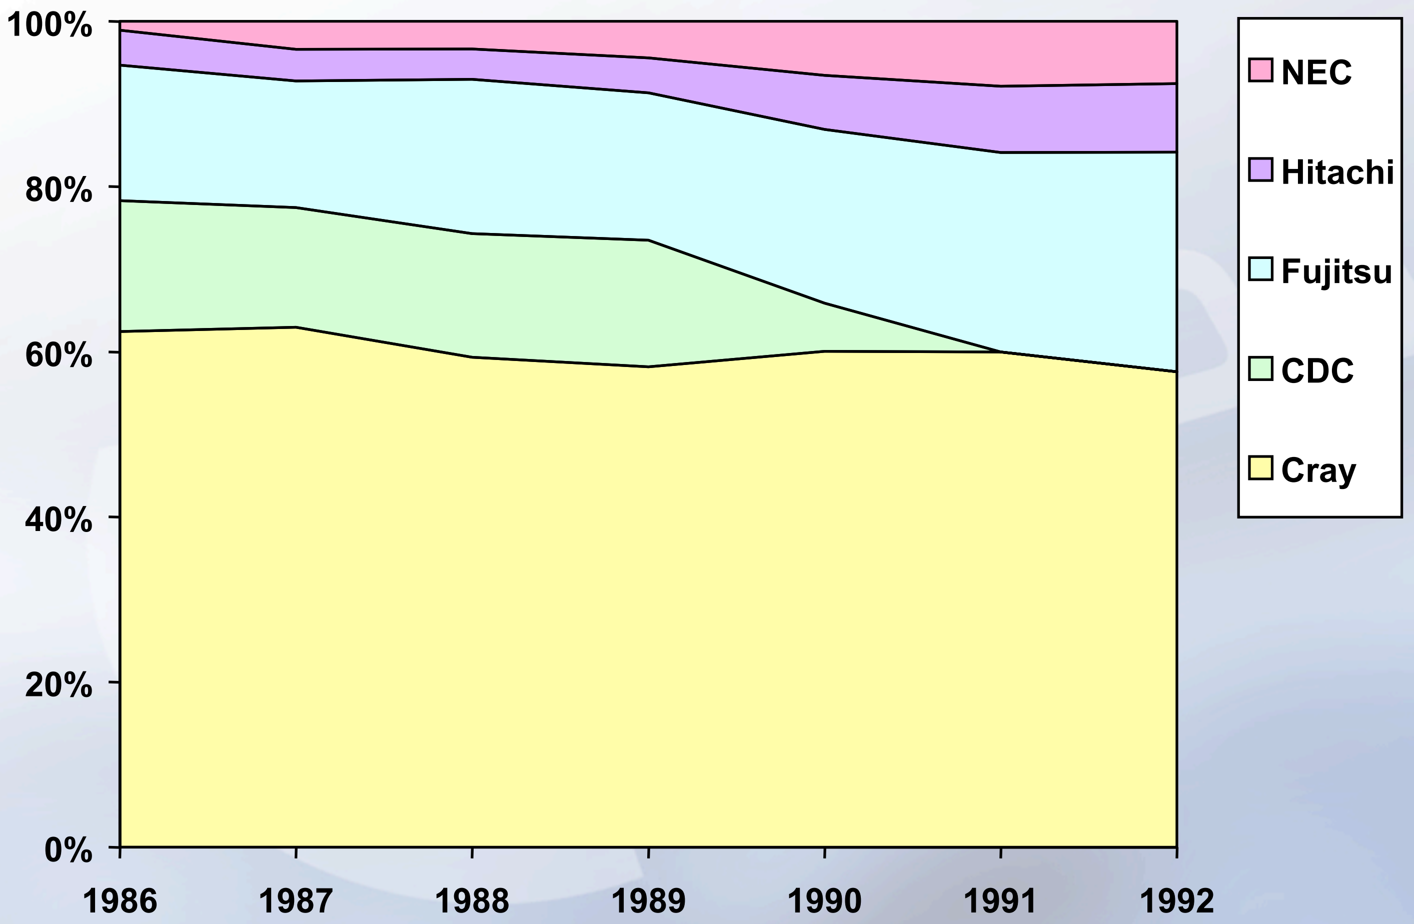

Countries Share

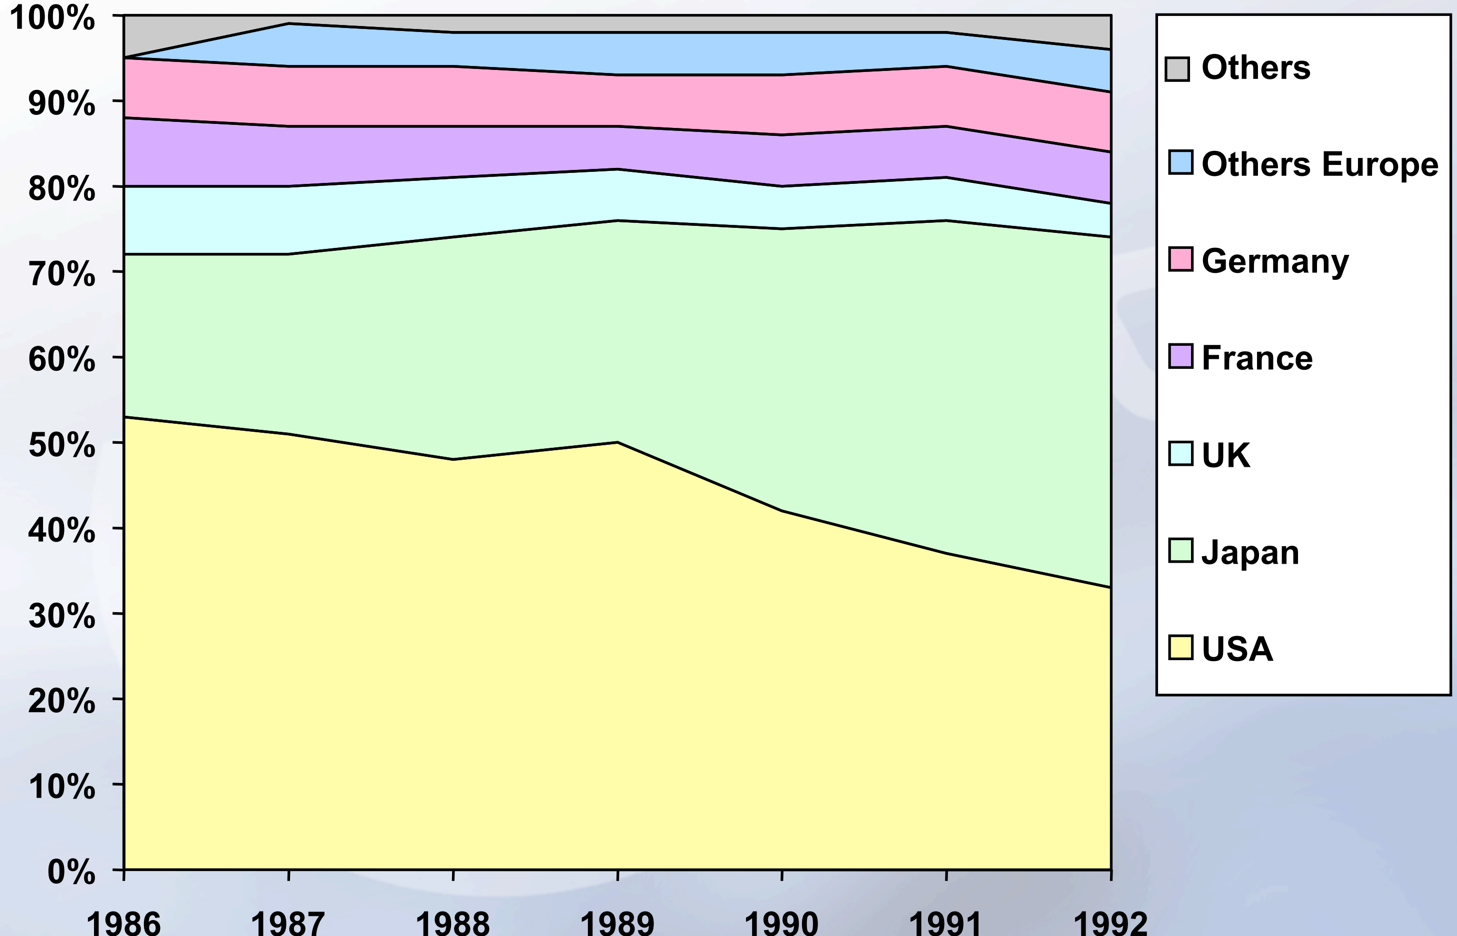

Deficits of the Mannheim Supercomputer Statistics

- Not very reliable
- Vector Supercomputers only – „numerical“/ mainly „scientific and engineering“ applications
- Definition of „Supercomputer“ necessary

Top500 Procedure

- Listing the 500 most powerful computers in the world
- Yardstick: Rmax of Linpack

  Solve Ax=b, dense problem, matrix is random
- Update twice a year:
  - ISC'xy in June in Germany
  - SC'xy in November in the U.S.
- All information available from the TOP500 webserver at: www.top500.org

Top500 Authors

Project was started in spring 1993 by:

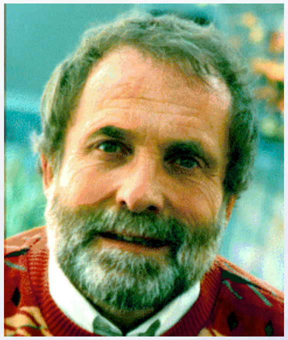

Hans W. Meuer

Erich Strohmaier

Authors since:

06/1993
Jack Dongarra

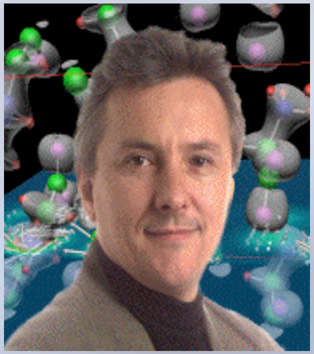

11/2000
Horst Simon

Dongarra's TOP500 List of World's Fastest Supercomputers Released at Mannheim Conference

„CRPC researcher Jack Dongarra of the University of Tennessee and Oak Ridge National Laboratory is one of three renowned computer scientists who assemble the legendary TOP500 List of the world's fastest supercomputers. Released twice a year since 1993, the list features sites with the most powerful computer systems, determined with information from a questionnaire sent to high-performance computer (HPC) experts, computational scientists, manufacturers, and the Internet community at large. Dongarra, Hans Meuer of the University of Mannheim, and Erich Strohmaier of the University of Tennessee released their June 1999 TOP500 List at the 14th Mannheim Supercomputing Conference and Seminar, held June 10-12 in Mannheim, Germany."

Source: http://www.crpc.rice.edu/WhatsNew/top500.html
The Center for Research on Parallel Computation at Rice University

Outline

- Mannheim Supercomputer Statistics
  & Top500 Project Start in 1993
- Competition between Manufacturers,
  Countries and Sites
- My Supercomputer Favorite
  in the Top500 Lists
- The 31st List as of June 2008
- Performance Development and Projection
- Bell's Law
- Supercomputing, quo vadis?
- Top500, quo vadis?

Top500 Status

- 1st Top500 List in June 1993 at ISC'93 in Mannheim
  .
  .
  .
  .
- 30th Top500 List on November 13, 2007 at SC07 in Reno
- **31st Top500 List on June 18, 2008 at ISC'08 in Dresden**
- 32nd Top500 List on November 18, 2008 at SC08 in Austin
- 33rd Top500 List on June 24, 2009 at ISC'09 in Hamburg
- Acknowledged by HPC-users, manufacturers and media

Competition between Manufacturers, Countries and Sites

1st List as of 06/1993

Countries	Count	Share %
USA	225	45.00 %
Japan	111	22.20 %
Germany	59	11.80 %
France	26	5.20 %
United Kingdom	25	5.00 %
Australia	9	1.80 %
Italy	6	1.20 %
Netherlands	6	1.20 %
Switzerland	4	0.80 %
Canada	3	0.60 %
Denmark	3	0.60 %
Korea	3	0.60 %
Others	20	4.00 %
Totals	**500**	**100 %**

31st List as of 06/2008

Countries	Count	Share %
USA	257	51.40%
Japan	22	4.40 %
Germany	46	9.20 %
France	34	6.80 %
United Kingdom	53	10.60 %
Australia	1	0.20 %
Italy	5	1.00 %
Netherlands	5	1.00 %
Switzerland	6	1.20 %
Canada	2	0.40 %
Denmark	—	—
Korea	2	0.40 %
China	12	2.40 %
India	6	1.20 %
Others	49	9.80 %
Totals	**500**	**100 %**

Competition between Manufacturers, Countries and Sites

Countries/Systems

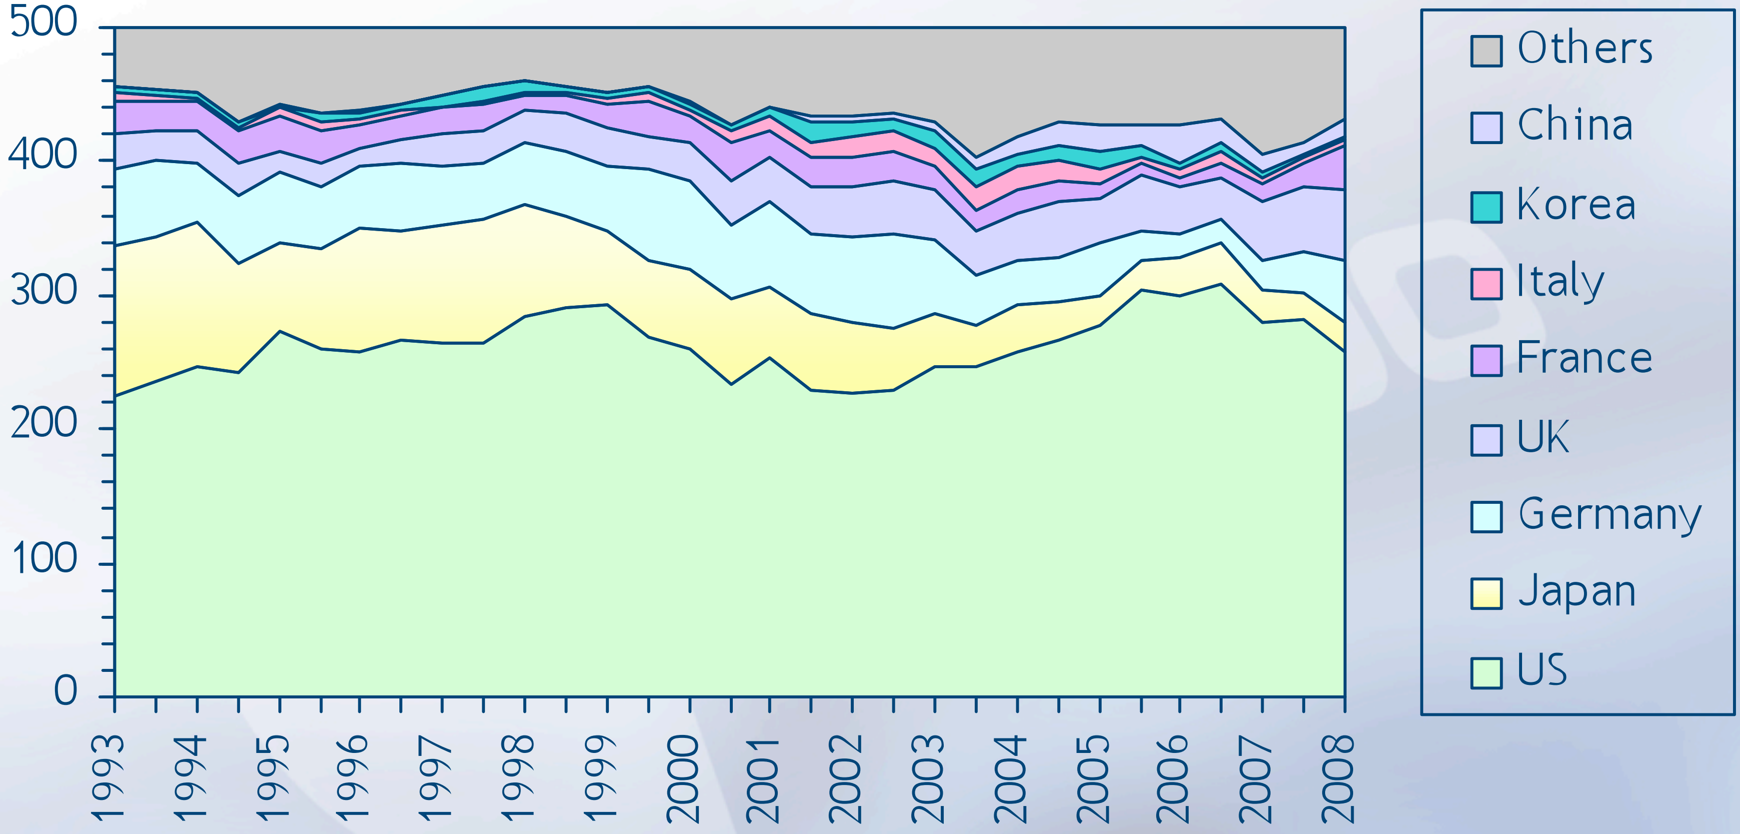

TOP500 SUPERCOMPUTER SITES

Competition between Manufacturers, Countries and Sites

1st List as of 06/1993

Manufacturers	Count	Share %
Cray Research	205	41.00 %
Fujitsu	69	13.80 %
Thinking Machines	54	10.80 %
Intel	44	8.80 %
Convex	36	7.20 %
NEC	32	6.40 %
Kendall Square Research	21	4.20 %
MasPar	18	3.60 %
Meiko	9	1.80 %
Hitachi	6	1.20 %
Parsytec	3	0.60 %
nCube	3	0.60 %
Totals	**500**	**100 %**

31st List as of 06/2008

Manufacturers	Count	Share %
Cray Inc.	16	3.20 %
Fujitsu	4	0.80 %
Thinking Machines	—	—
Intel	1	0.20 %
Hewlett-Packard	183	36.60%
NEC	1	0.20 %
Kendall Square Research	—	—
MasPar	—	—
Meiko	—	—
Hitachi/Fujitsu	1	0.20 %
Parsytec	—	—
nCube	—	—
IBM	210	42.00 %
SGI	22	4.40 %
Dell	27	5.40 %
Others	35	7.00 %
Totals	**500**	**100 %**

Manufacturer/Systems

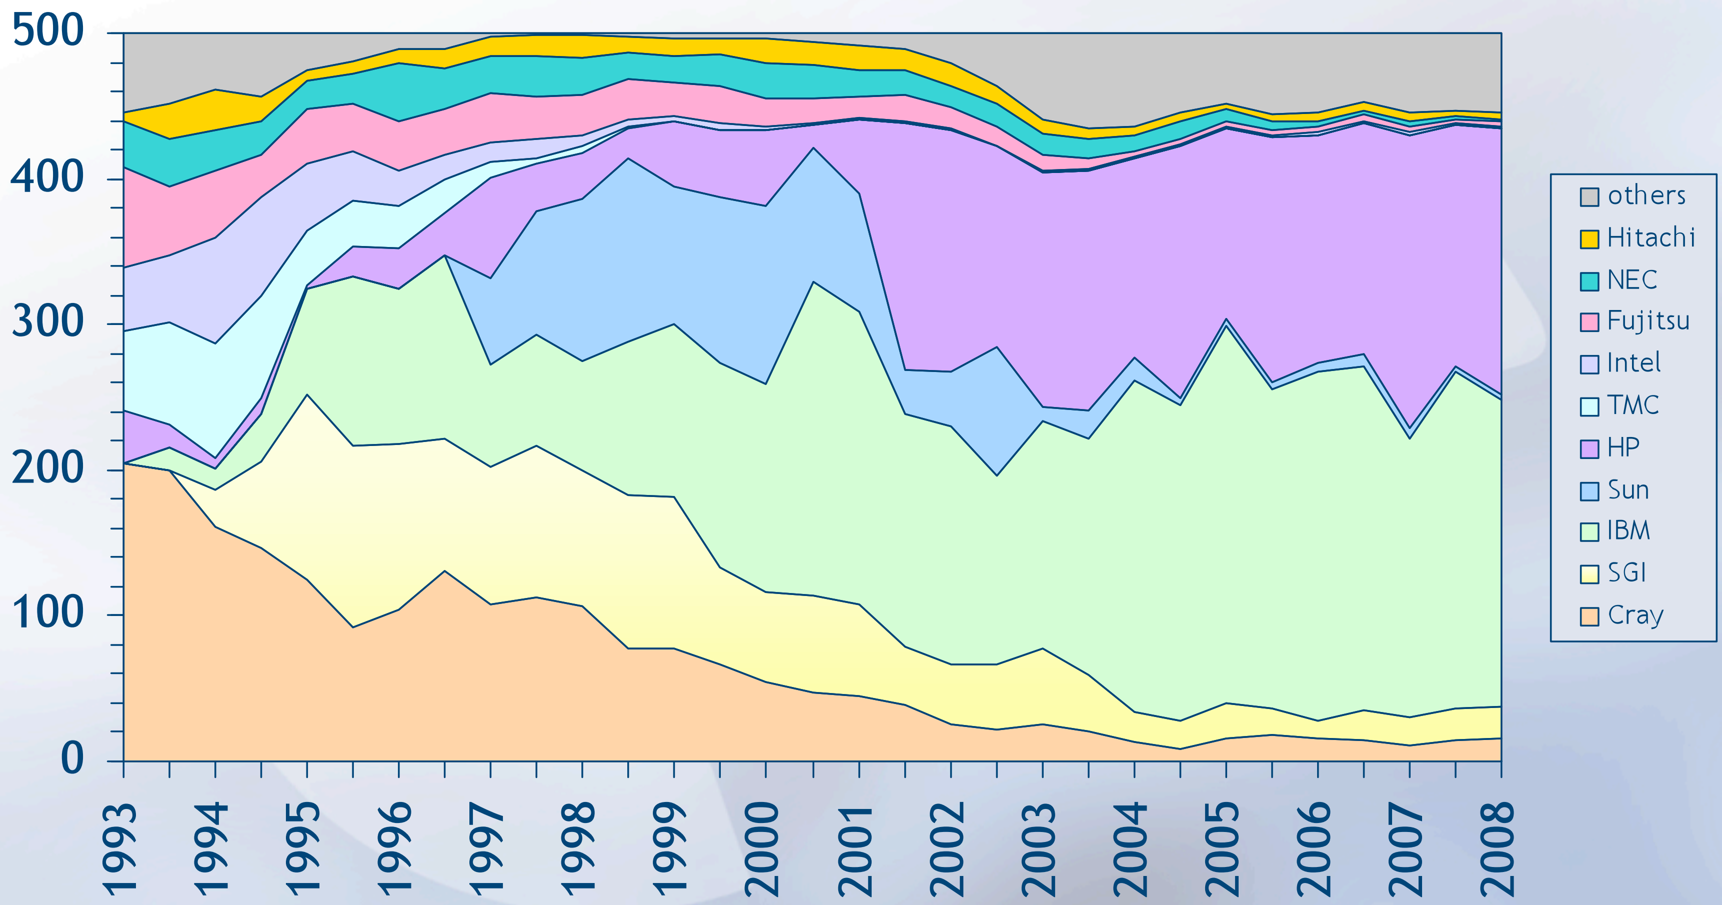

Top Sites through 31 Lists

Rank	Site	Country	% over time
1	Lawrence Livermore National Laboratory	United States	5.4
2	Sandia National Laboratories	United States	3.63
3	Los Alamos National Laboratory	United States	3.59
4	Government	United States	3.23
5	The Earth Simulator Center	Japan	1.93
6	National Aerospace Laboratory of Japan	Japan	1.65
7	Oak Ridge National Laboratory	United States	1.43
8	NCSA	United States	1.3
9	NASA/Ames Research Center/NAS	United States	1.24
10	NERSC/LBNL	United States	1.18
11	University of Tokyo	Japan	1.18
12	Pittsburgh Supercomputing Center	United States	1.11
13	Semiconductor Company (C)	United States	1.1
14	Naval Oceanographic Office (NAVOCEANO)	United States	1.06
15	ECMWF	United Kingdom	1.02
16	ERDC MSRC	United States	0.91
17	Forschungszentrum Juelich (FZJ)	Germany	0.87
18	IBM Thomas J. Watson Research Center	United States	0.86
19	Japan Atomic Energy Research Institute	Japan	0.8
20	Minnesota Supercomputer Center	United States	0.72

Outline

- Mannheim Supercomputer Statistics & Top500 Project Start in 1993
- Competition between Manufacturers, Countries and Sites
- ➔ My Supercomputer Favorite in the Top500 Lists
- The 31st List as of June 2008
- Performance Development and Projection
- Bell's Law
- Supercomputing, quo vadis?
- Top500, quo vadis?

My Supercomputer Favorite in the Top500 Lists

Top10 in the 06/1997 list

	Manufacturer	Computer	Rmax [GF/s]	Installation Site	Country	Year	#Proc
1	Intel	ASCI Red	1068	Sandia National Laboratories	USA	1996	7264
2	Hitachi	CP-PACS/2048	368.2	Center for Computational Science	Japan	1996	2048
3	Fujitsu	Numerical Wind Tunnel	229	National Aerospace Laboratory	Japan	1996	167
4	Hitachi	SR2201/1024	220.4	University of Tokyo	Japan	1996	1024
5	Cray	T3E	176	Forschungszentrum Jülich	Germany	1996	512
6	Cray	T3E	176	Government	USA	1997	512
7	Cray	T3E	176	Max-Planck-Gesell. MPI/IPP	Germany	1997	512
8	Cray	T3E	176	NASA/Goddard Space Flight Center	USA	1996	512
9	Cray	T3E	176	Pittsburgh Supercomp. Center	USA	1996	512
10	Cray	T3E	176	University Stuttgart	Germany	1996	512

My Supercomputer Favorite in the Top500 Lists

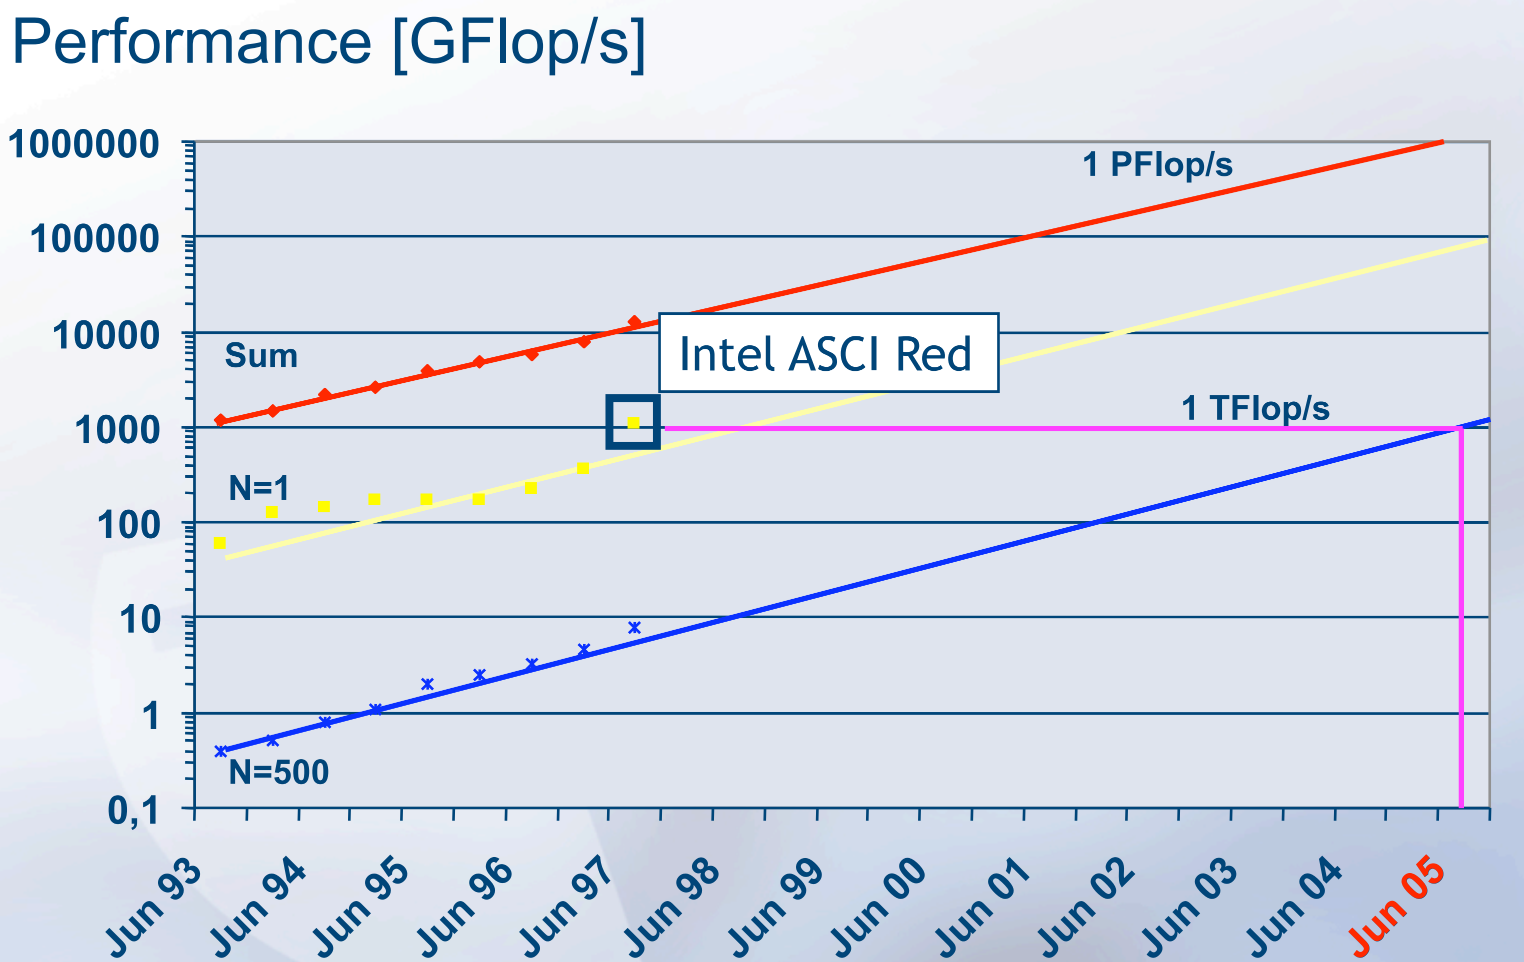

My Supercomputer Favorite in the Top500 Lists

Number of Teraflop/s Systems in the Top500

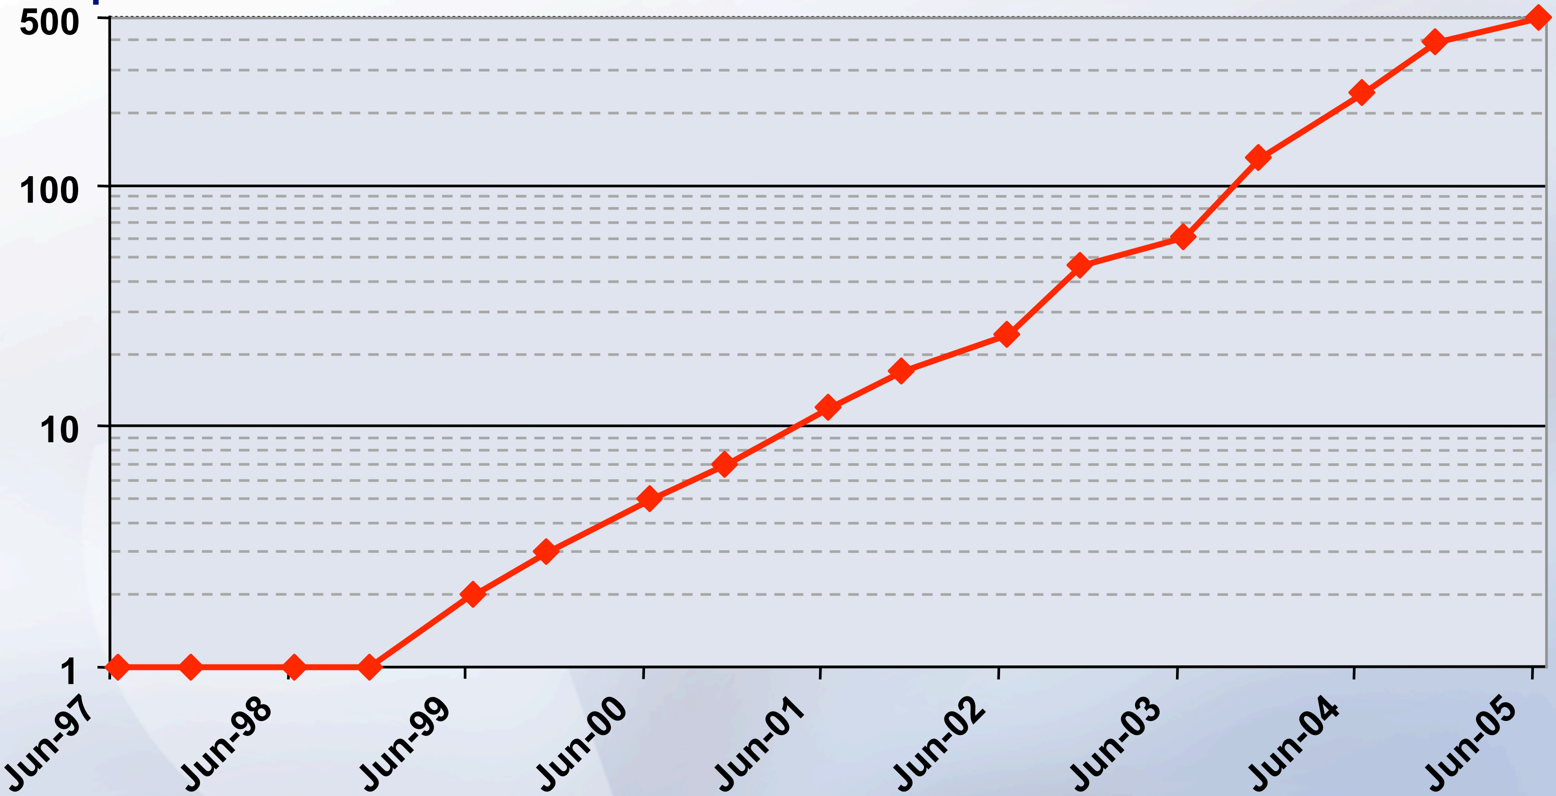

My Supercomputer Favorite in the Top500 Lists

Outline

- Mannheim Supercomputer Statistics & Top500 Project Start in 1993
- Competition between Manufacturers, Countries and Sites
- My Supercomputer Favorite in the Top500 Lists
- ➔ The 31st List as of June 2008
- Performance Development and Projection
- Bell's Law
- Supercomputing, quo vadis?
- Top500, quo vadis?

The 31st List as of June 2008

Processor Architecture/Systems

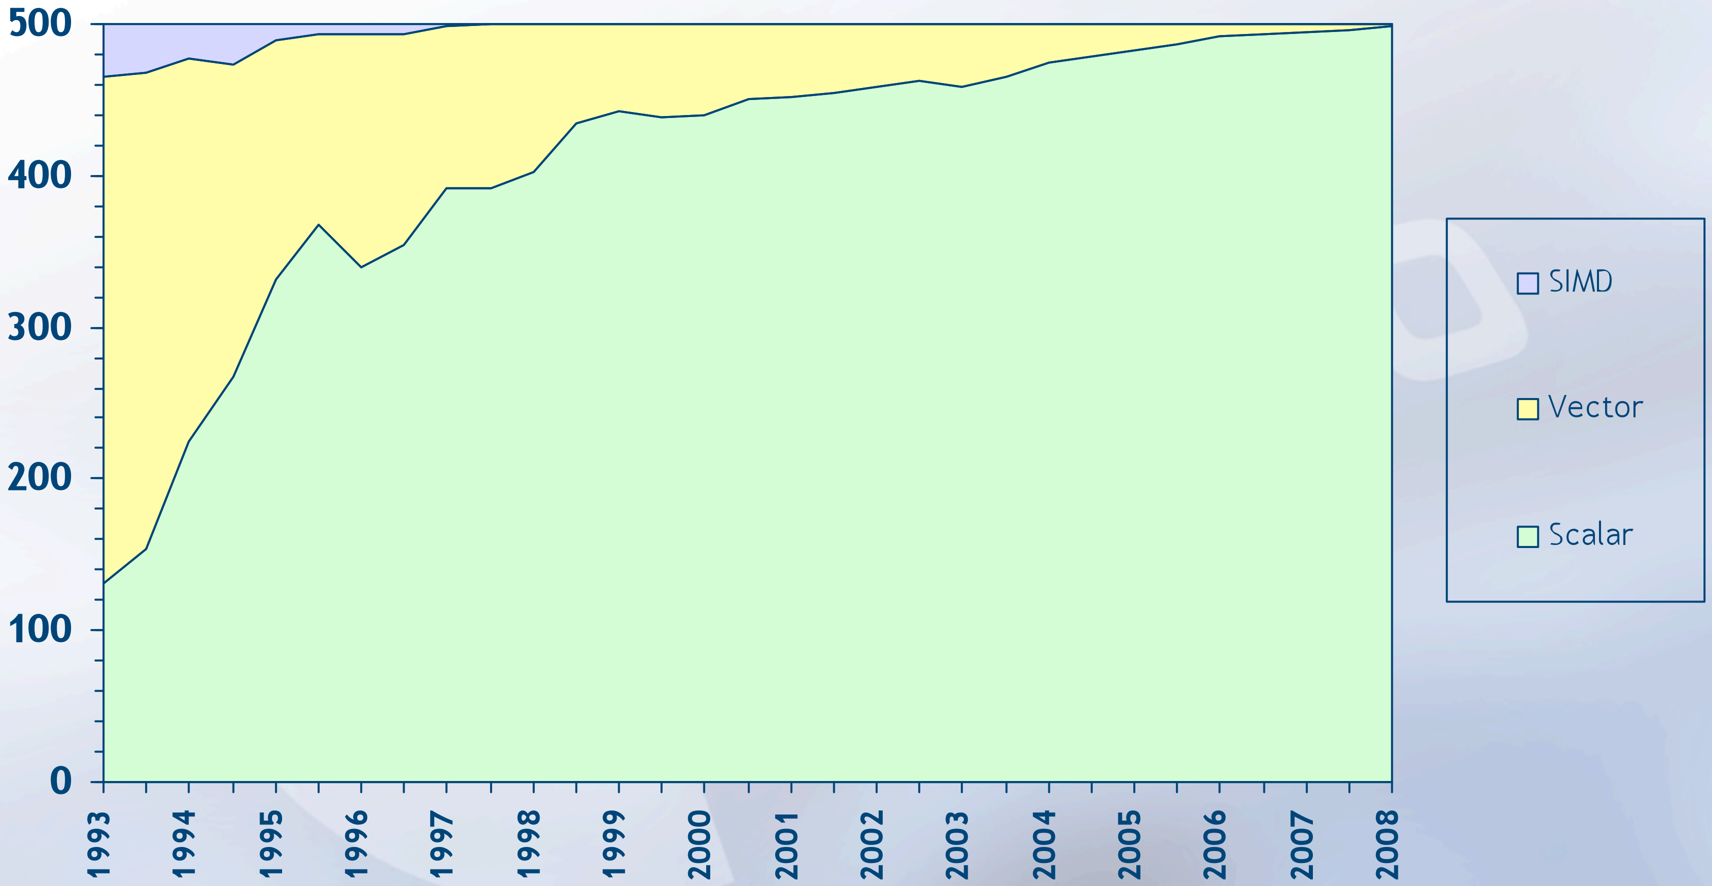

The 31st List as of June 2008

Operating Systems/Systems

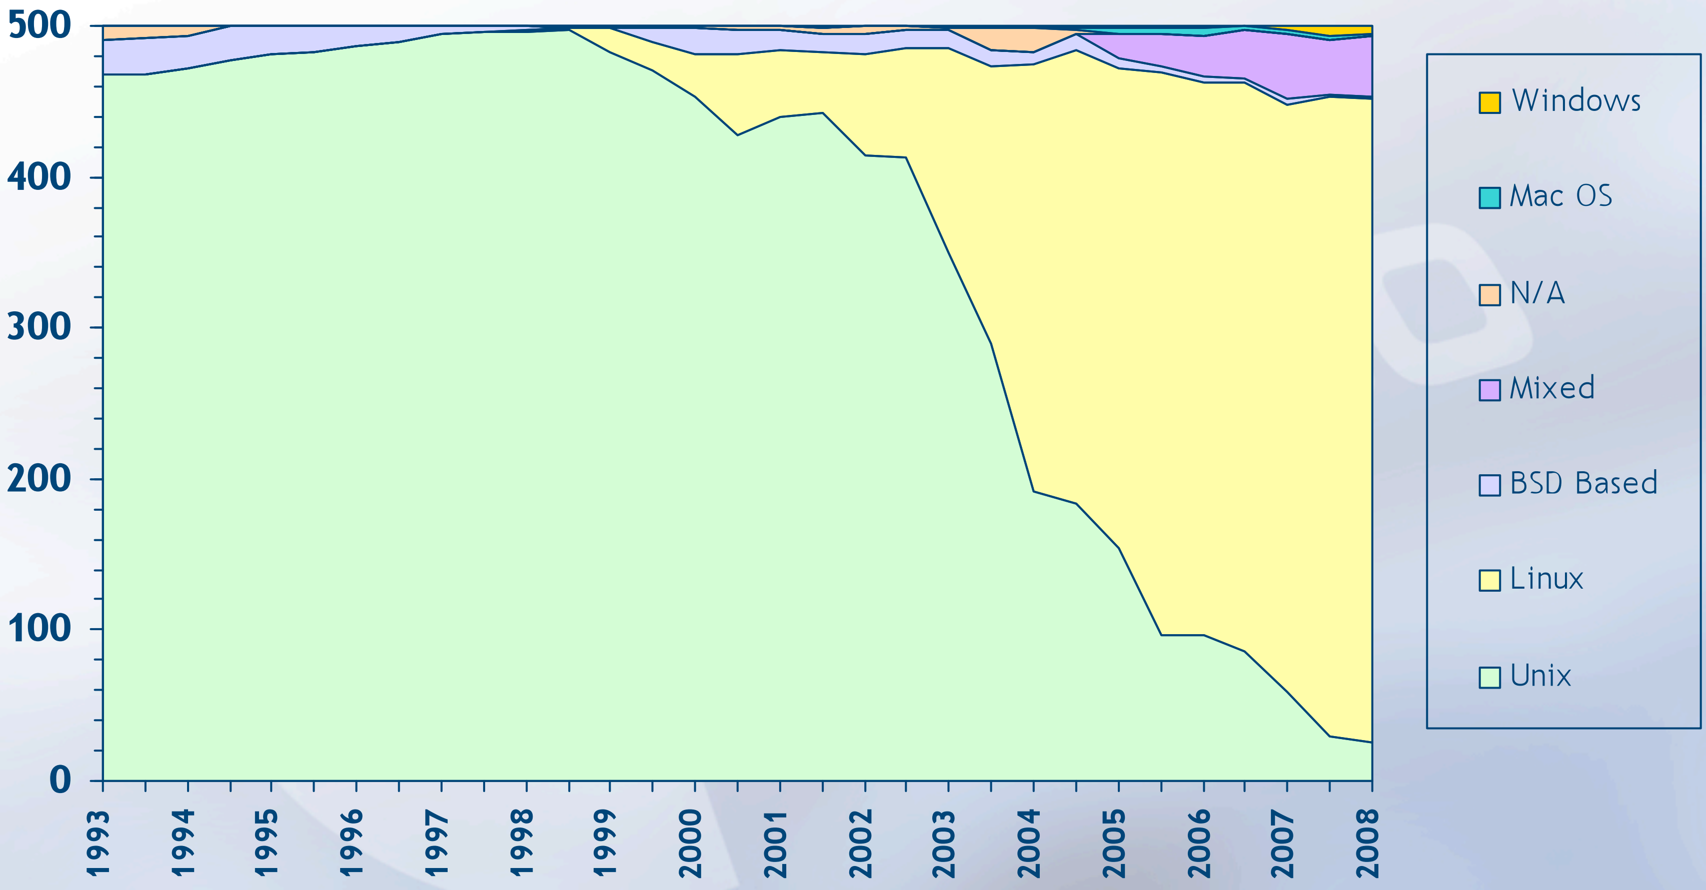

The 31st List as of June 2008

Windows in the 06/2008 list

Rank	Site	Computer/Year Vendor	Cores	R_{max}	R_{peak}	Power
23	NCSA United States	Abe - PowerEdge 1955, 2.33 GHz, Infiniband, Windows Server 2008/Red Hat Enterprise Linux 4 / 2007 Dell	9600	68.48	89.59	
39	HPC2N - Umea University Sweden	Akka - BladeCenter HS21 Cluster, Xeon QC HT 2.5 GHz, IB, Windows HPC 2008/CentOS / 2008 IBM	5376	46.04	53.76	173.21
100	Universitaet Aachen/RWTH Germany	PRIMERGY RX200 S4 Cluster, Quad Core 3Ghz, Infiniband, Windows HPC 2008 / 2008 UNICORNER/Fujitsu-Siemens	2048	18.81	24.58	
203	Financial Institution (Q) France	BladeCenter LS21, Opteron Dual Core 2.2 GHz, GigEthernet, Windows / 2008 IBM	6400	14.18	28.16	416.45
302	Microsoft Windows HPC Group United States	Rainier - PowerEdge 1955, 1.86 GHz, Cisco Infiniband, Windows OS / 2007 Dell	2048	11.75	15.24	

The 31st List as of June 2008

Application Area/Systems (June 2008)

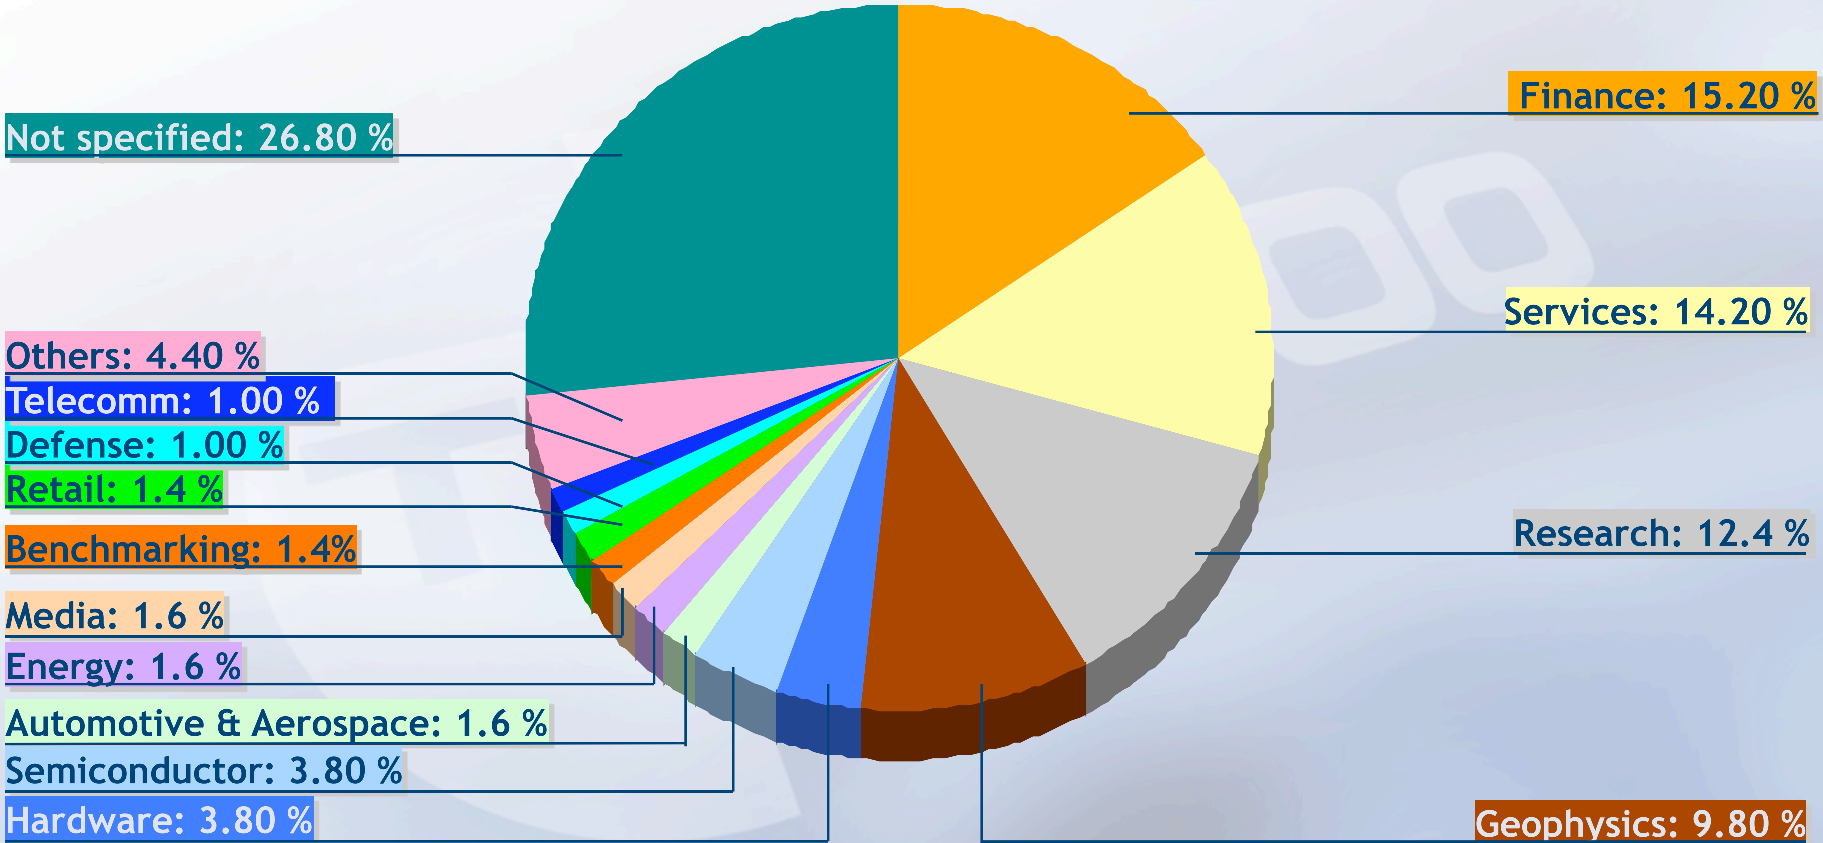

The 31st List as of June 2008

Processor Generations/Systems (June 2008)

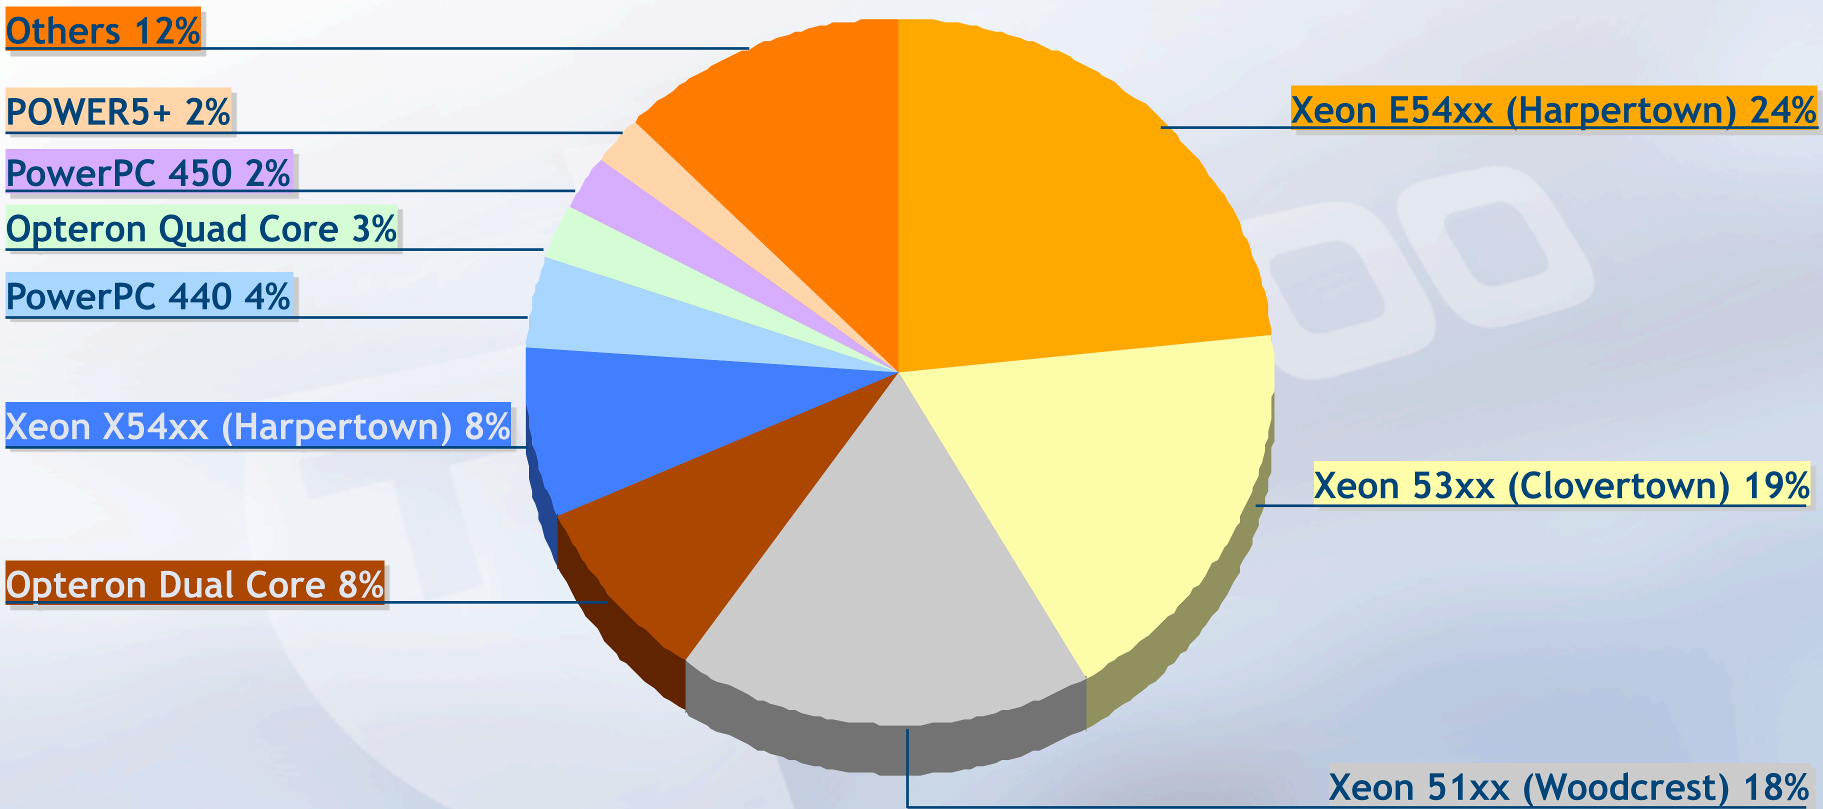

The 31st List as of June 2008

Multi-Core Level/Systems (June 2008)

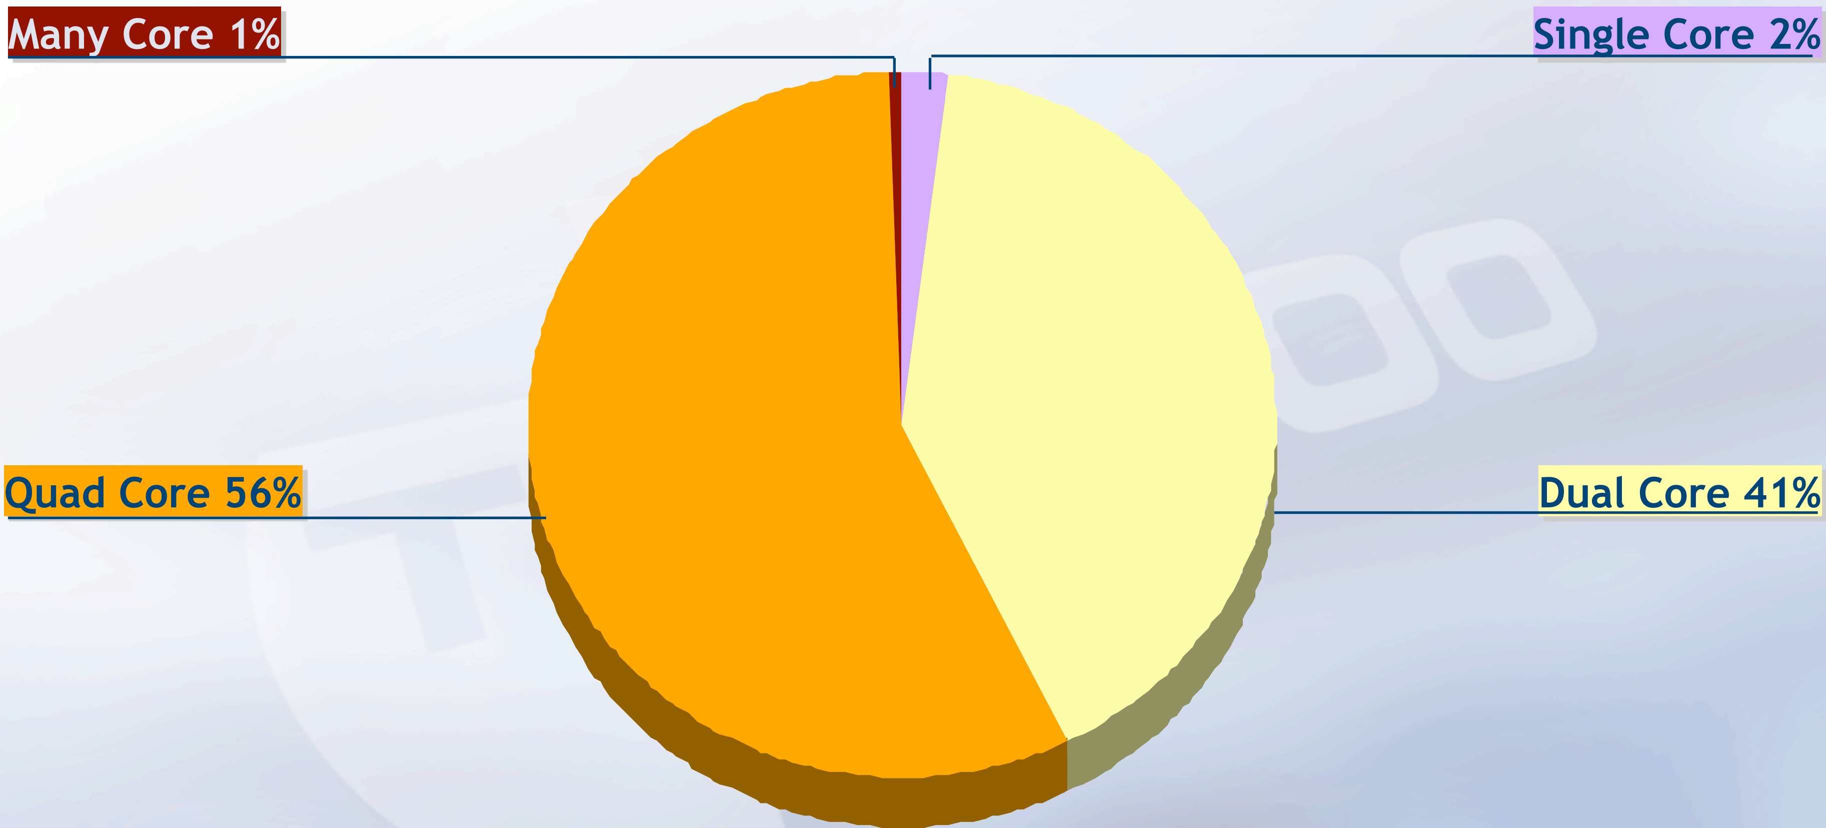

The 31st List as of June 2008

Interconnect Family/Systems

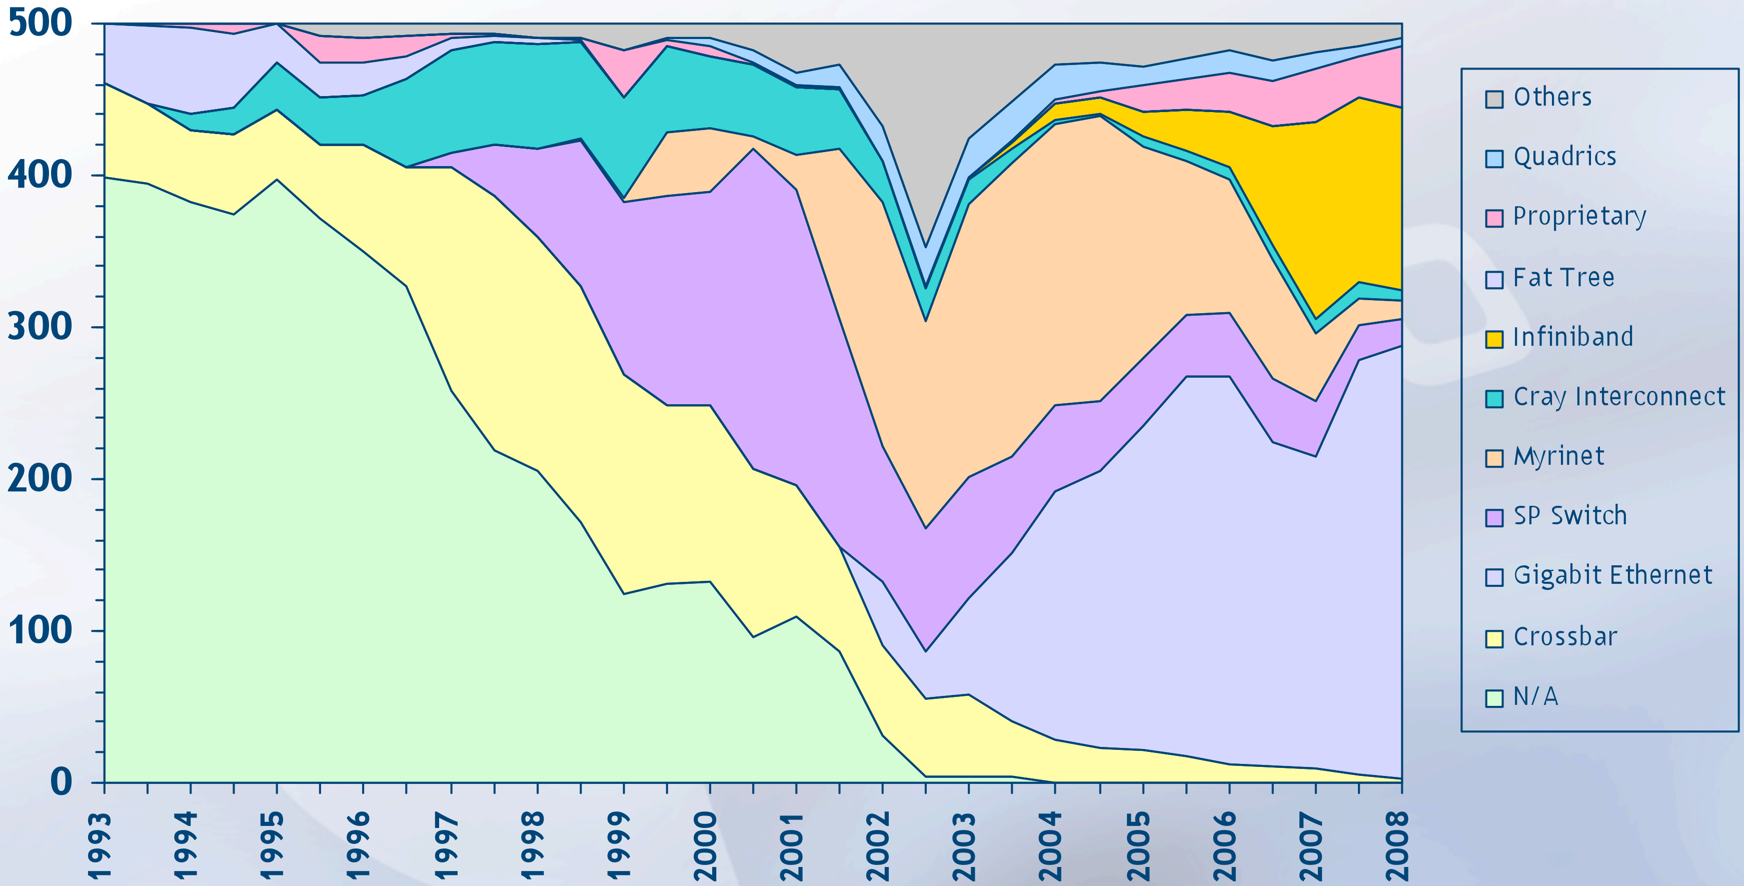

The 31st List as of June 2008

Architectures / Systems

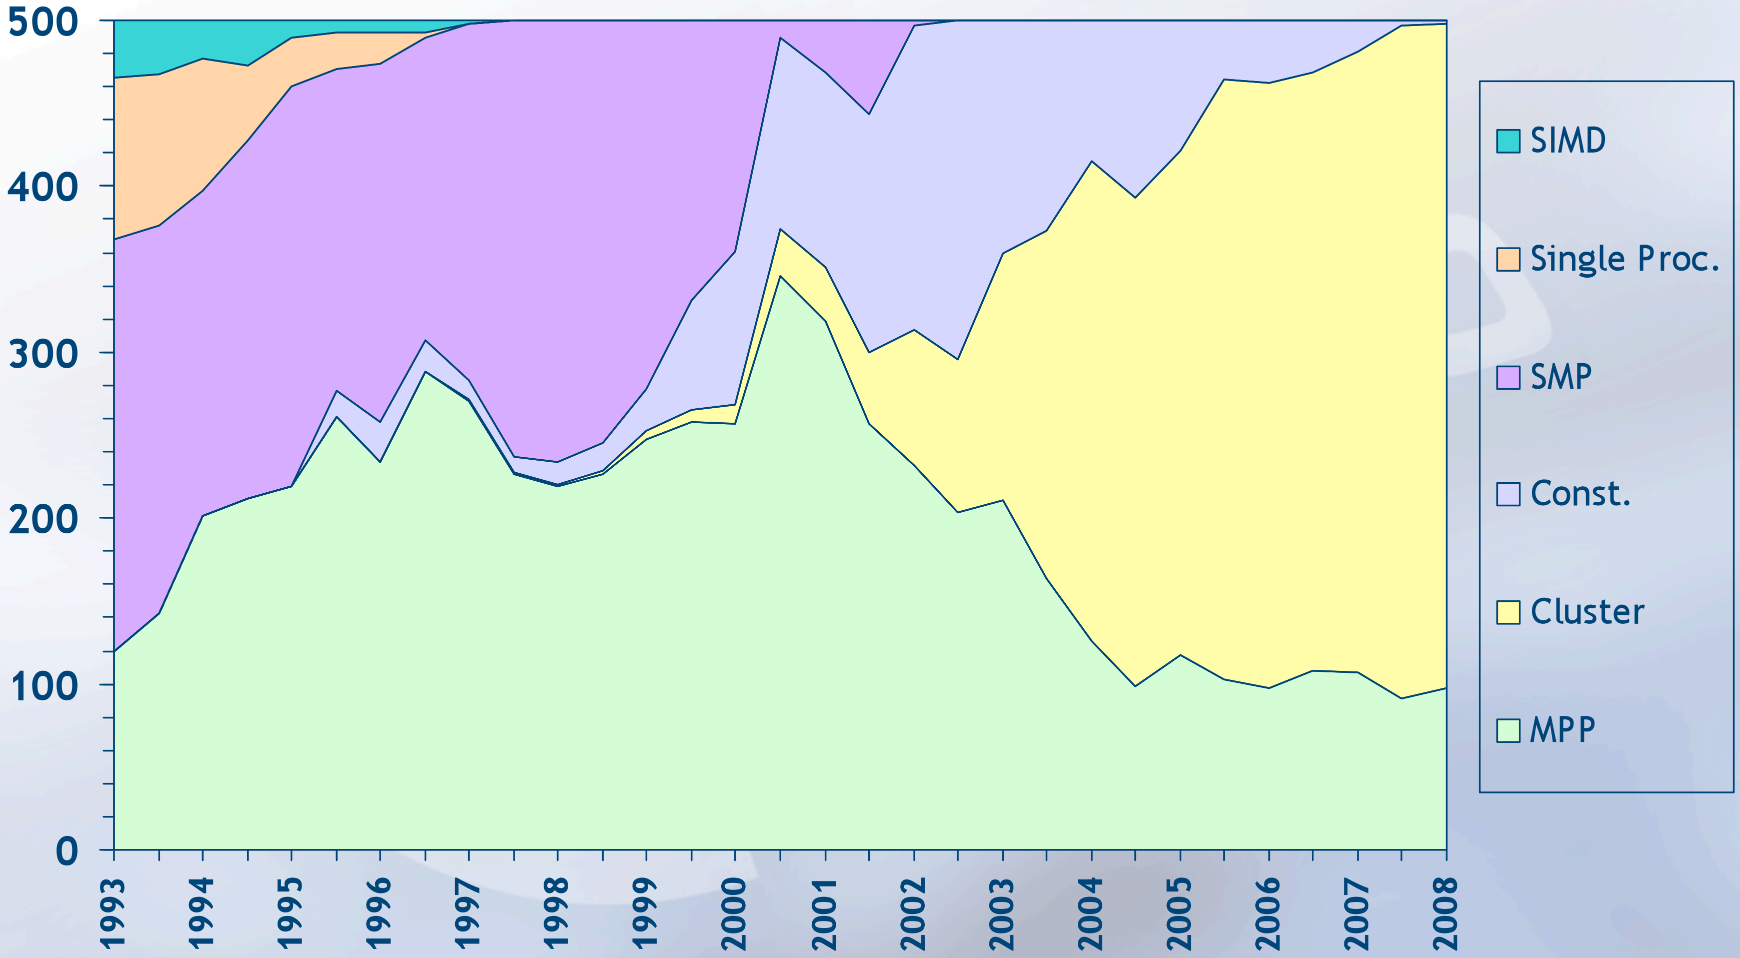

The 31st List as of June 2008

Customer Segments / Systems

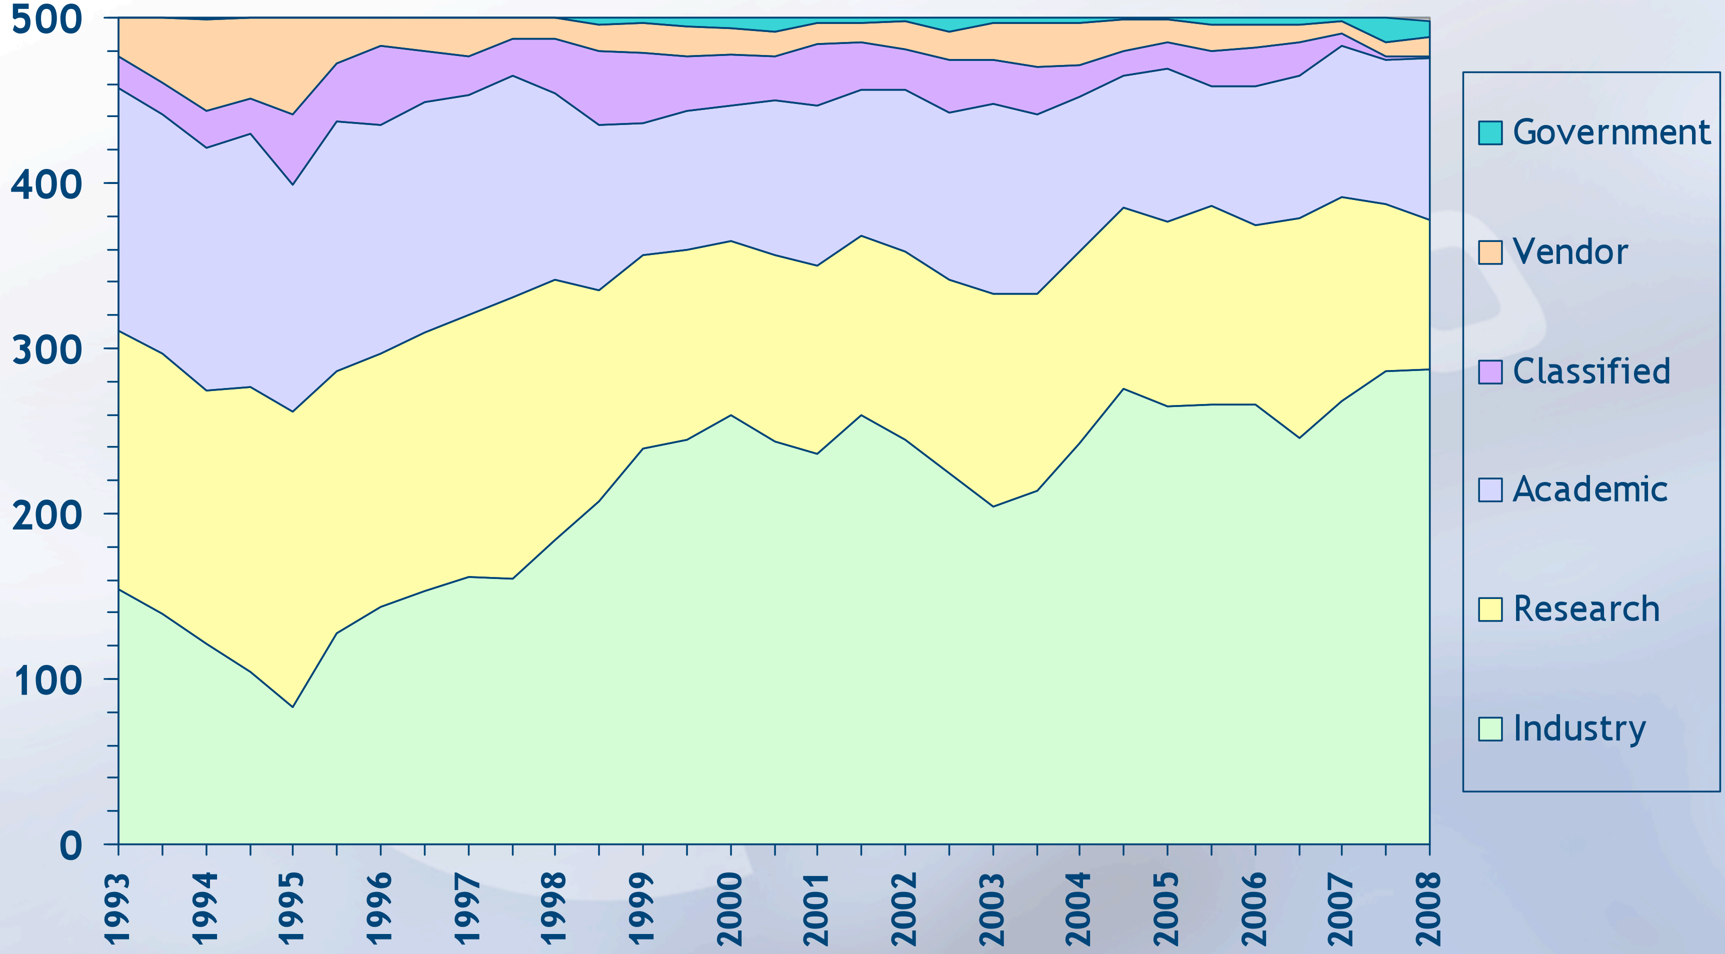

The 31st List as of June 2008

Customer Segments / Performance

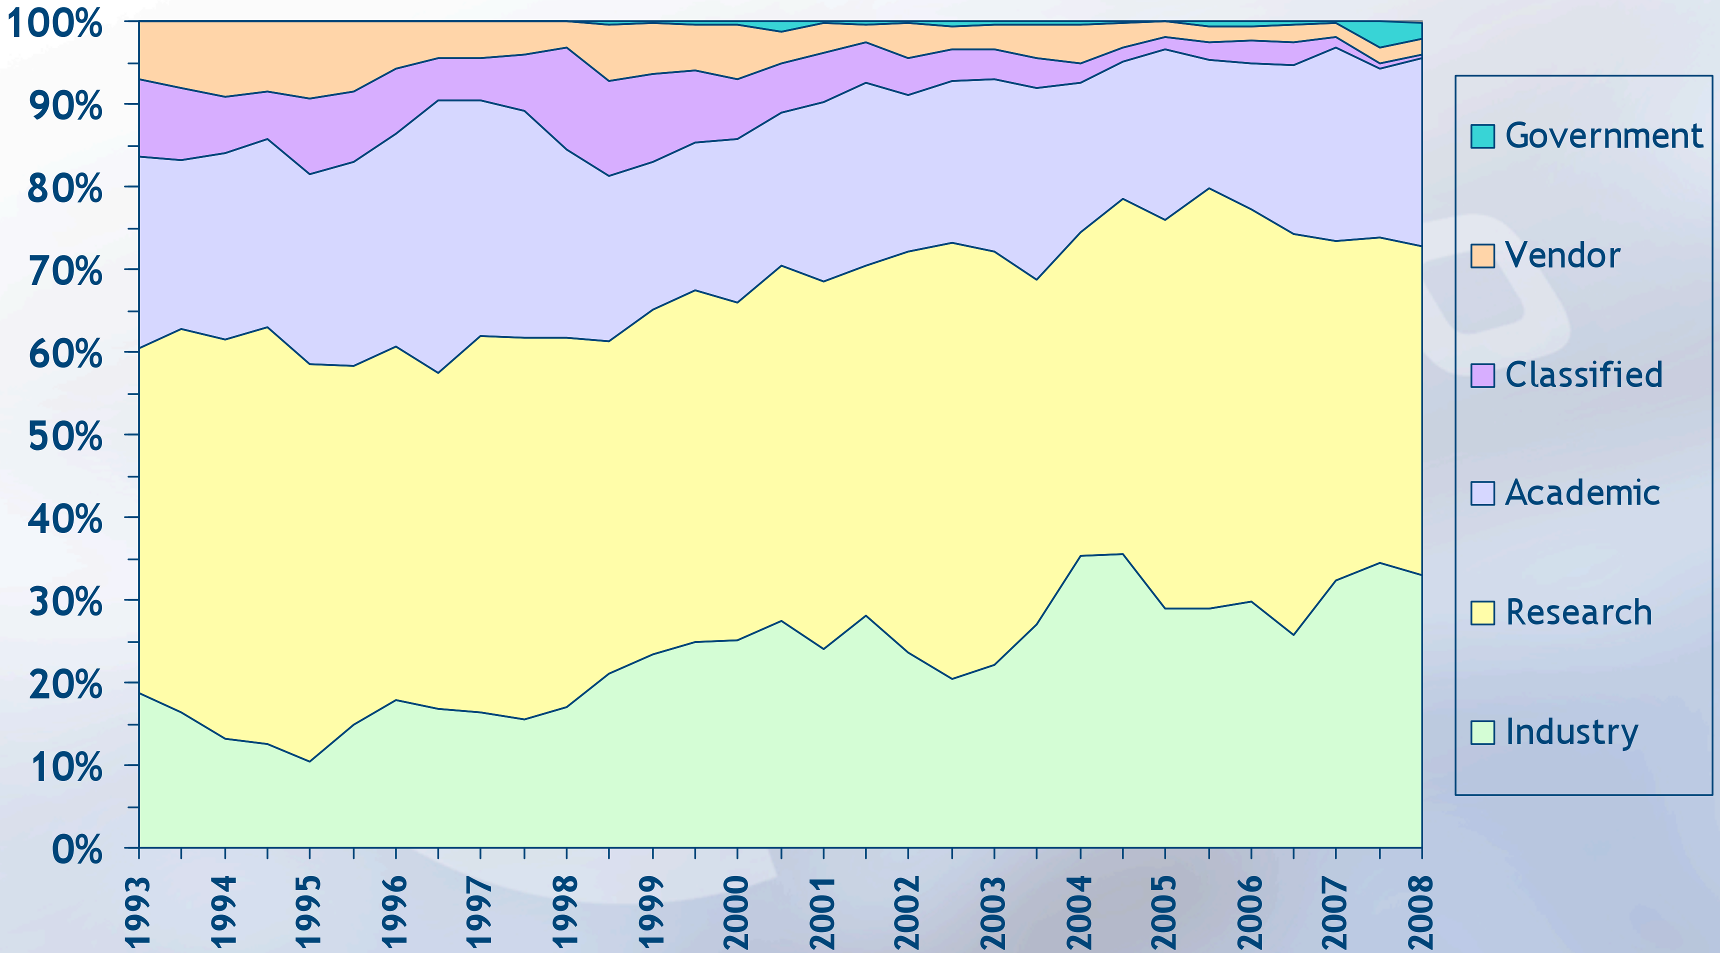

Outline

- Mannheim Supercomputer Statistics & Top500 Project Start in 1993
- Competition between Manufacturers, Countries and Sites
- My Supercomputer Favorite in the Top500 Lists
- The 31st List as of June 2008
- ➔ Performance Development and Projection
- Bell's Law
- Supercomputing, quo vadis?
- Top500, quo vadis?

Performance Development and Projection

Performance Development

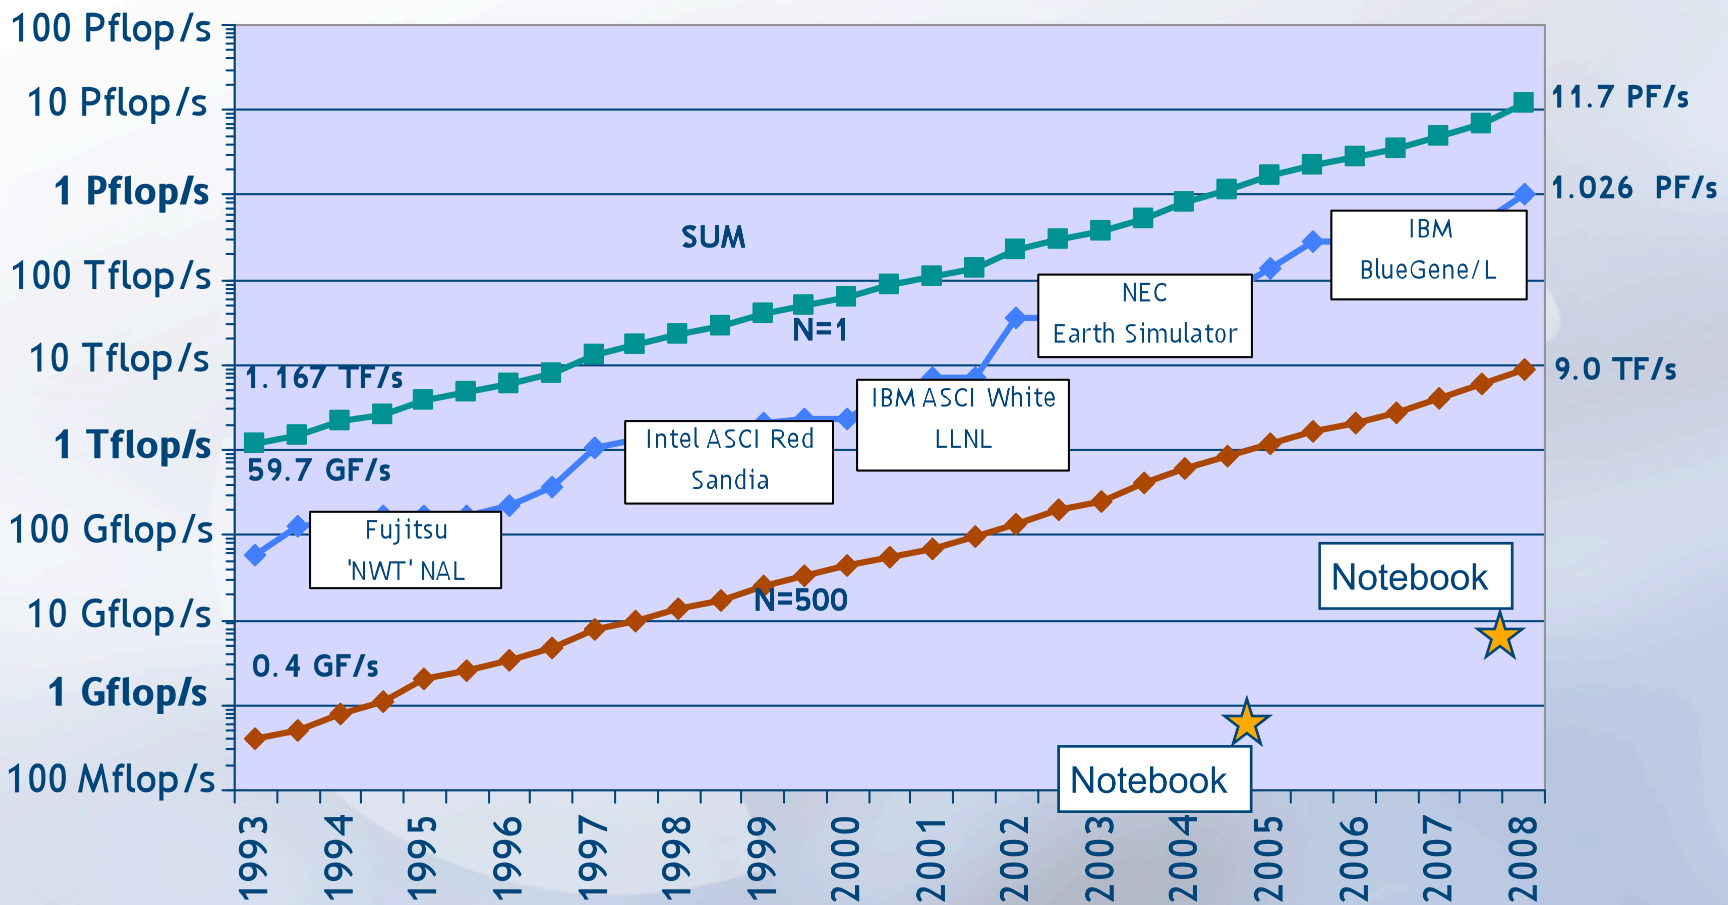

Performance Projection

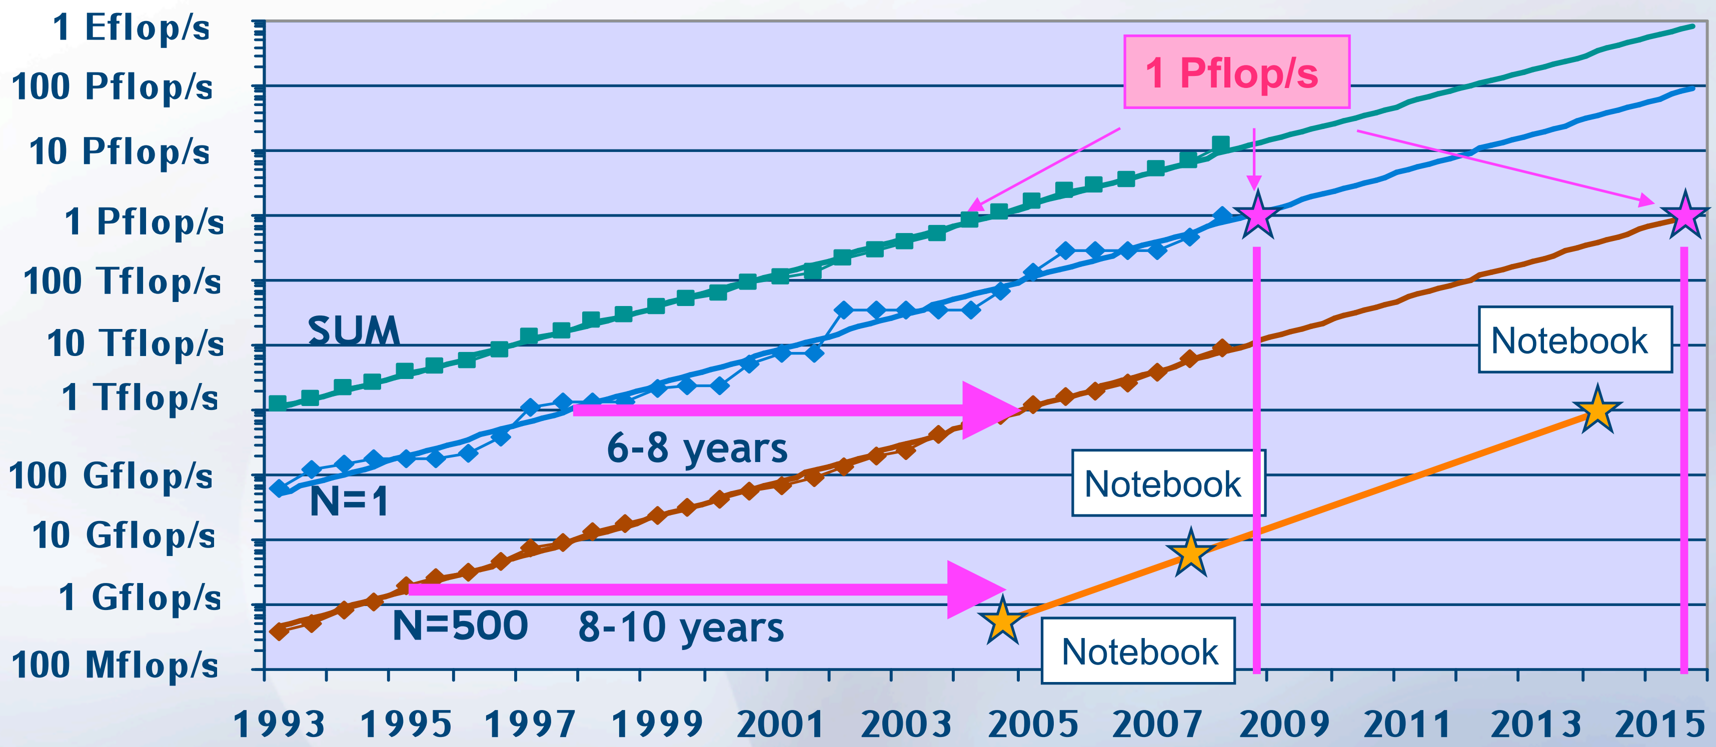

RoadRunner – the First Petaflop/s System in the World and its Impact on Supercomputing

Chair: Prof. Dr. Hans Meuer, Prometeus & University of Mannheim, Germany

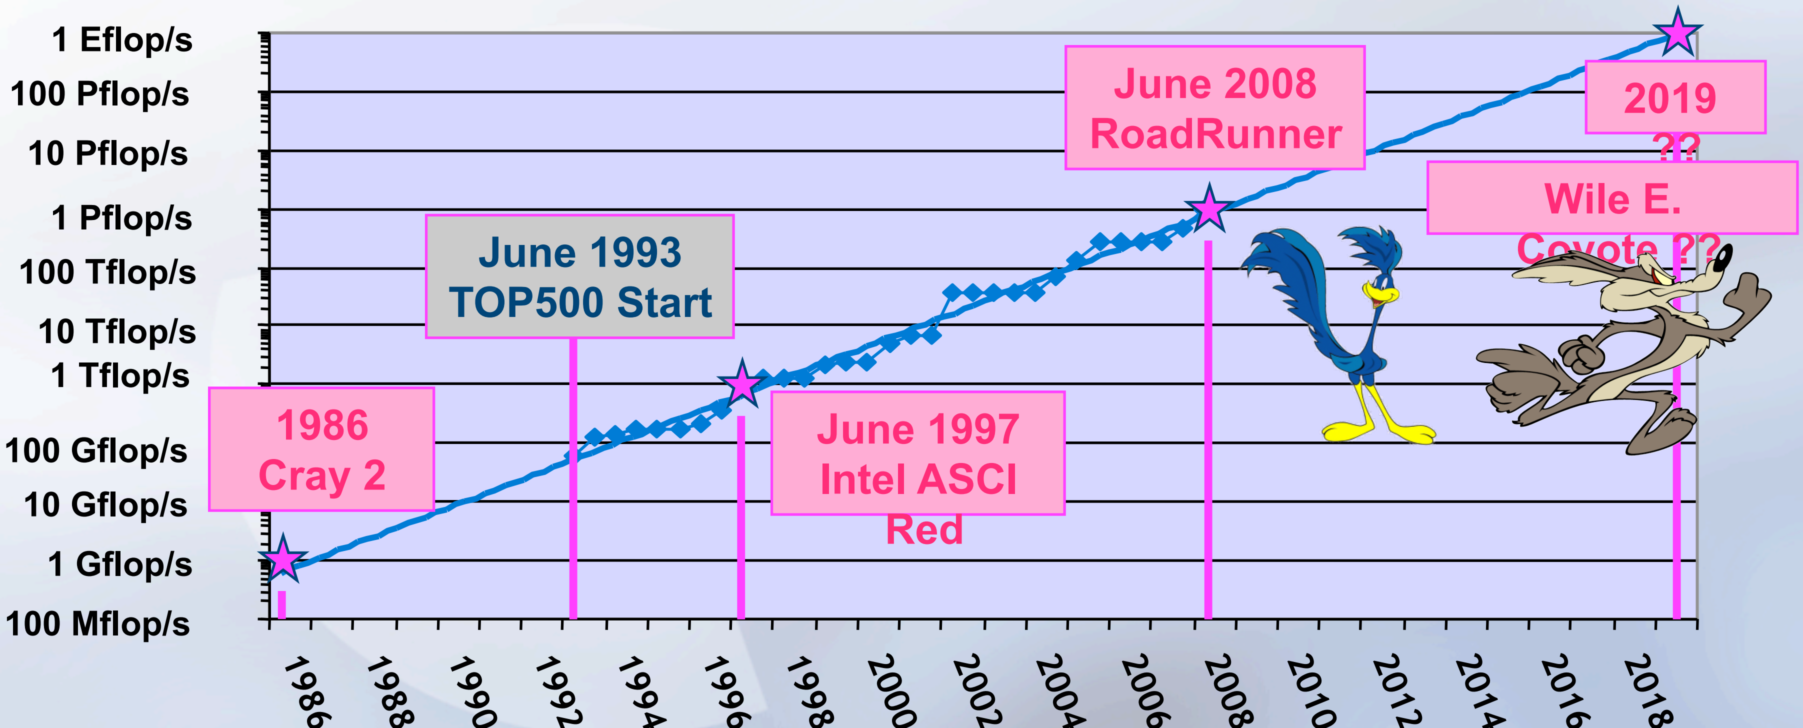

Outline

- Mannheim Supercomputer Statistics & Top500 Project Start in 1993
- Competition between Manufacturers, Countries and Sites
- My Supercomputer Favorite in the Top500 Lists
- The 31st List as of June 2008
- Performance Development and Projection
- ➔ Bell's Law
- Supercomputing, quo vadis?
- Top500, quo vadis?

Bell's Law

Bell's Law of Computer Class formation was discovered about 1972. It states that technology advances in semiconductors, storage, user interface and networking advance every decade enable a new, usually lower priced computing platform to form. Once formed, each class is maintained as a quite independent industry structure. This explains mainframes, minicomputers, workstations and Personal computers, the web, emerging web services, palm and mobile devices, and ubiquitous interconnected networks. We can expect home and body area networks to follow this path.

From Gordon Bell (2007), http://research.microsoft.com/~GBell/Pubs.htm

Bell's Law

Bell's Law states, that:

- Important classes of computer architectures come in cycles of about 10 years.
- It takes about a decade for each phase :
  - Early research
  - Early adoption and maturation
  - Prime usage
  - Phase out past its prime
- Can we use Bell's Law to classify computer architectures in the TOP500?

Bell‘s Law

Gordon Bell (1972): 10 year cycles for computer classes
Computer classes in HPC based on the TOP500:

- Data Parallel Systems:
  - Vector (Cray Y-MP and X1, NEC SX, …)
  - SIMD (CM-2, …)
- Custom Scalar Systems:
  - MPP (Cray T3E and XT3, IBM SP, …)
  - Scalar SMPs and Constellations (Cluster of big SMPs)
- Commodity Cluster: NOW, PC cluster, Blades, …
- Power-Efficient Systems (BG/L or BG/P as first example of low-power / embedded systems = potential new class ?)
  - Tsubame with Clearspeed, Roadrunner with Cell ?

Bell's Law

Computer Classes / Systems

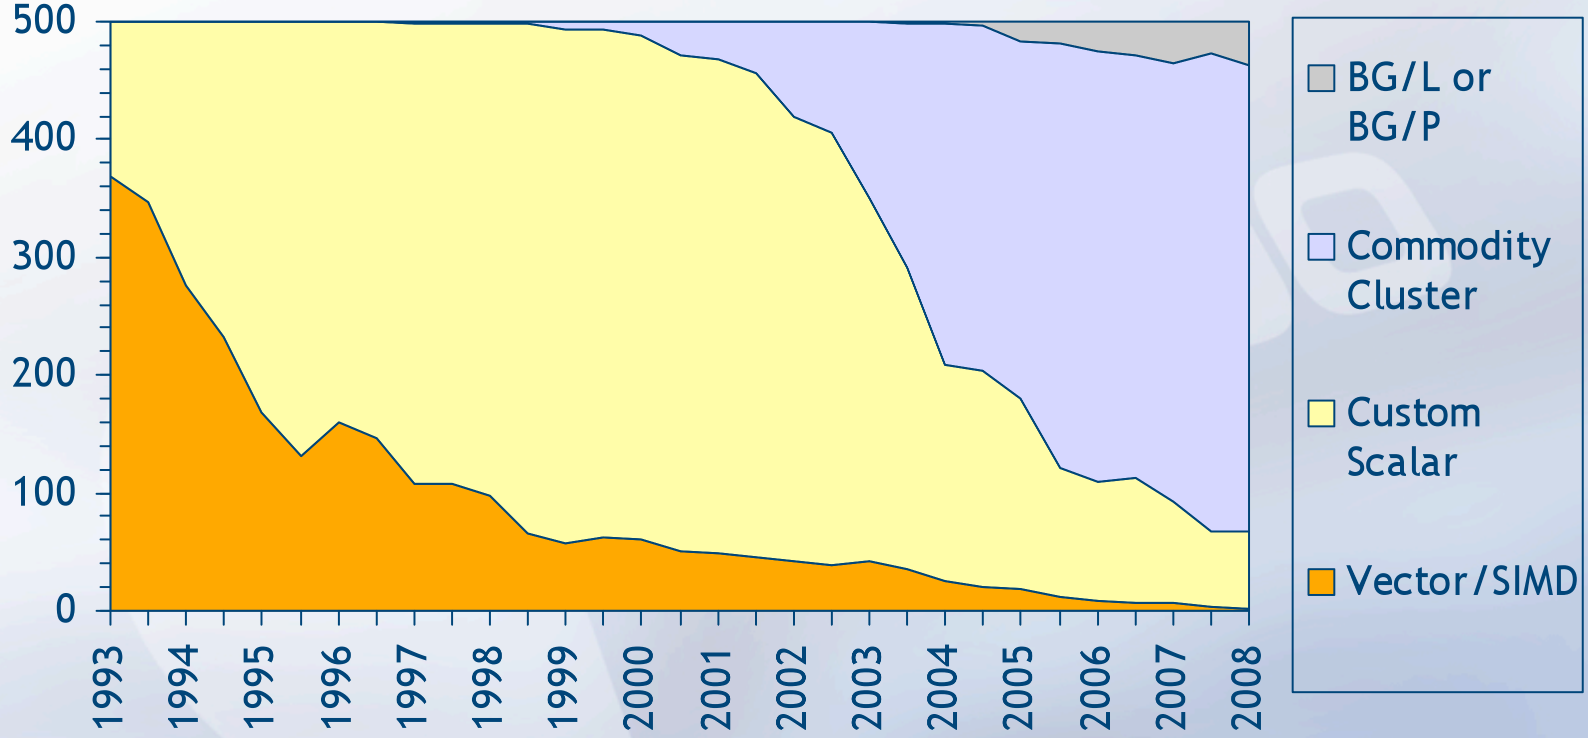

Bell's Law

Computer Classes - refined / Systems

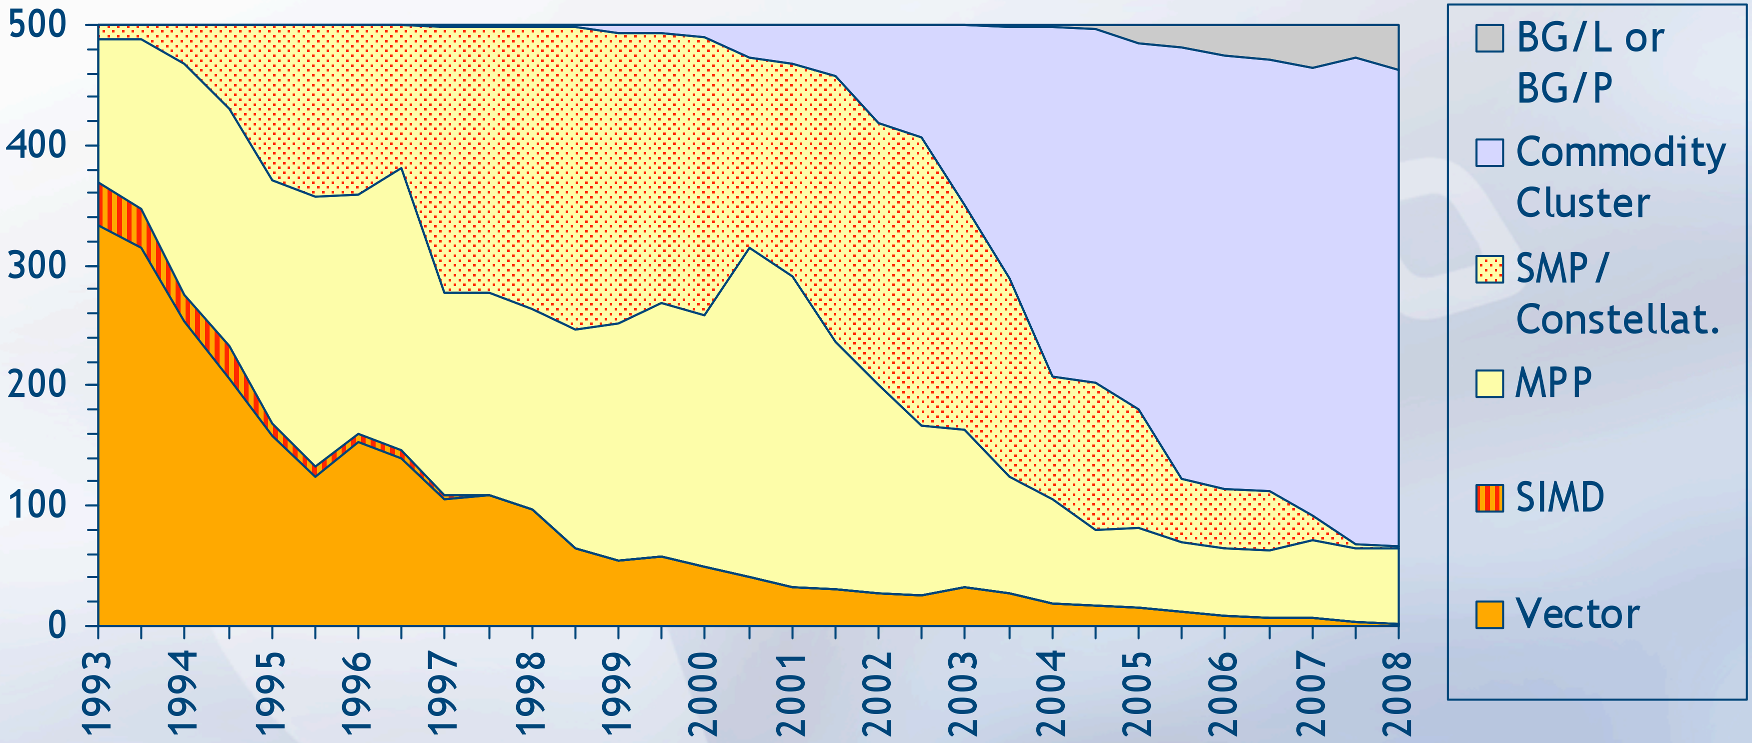

Bell's Law

HPC Computer Classes

Class	Early Adoption starts:	Prime Use starts:	Past Prime Usage starts:
Data Parallel Systems	Mid 70's	Mid 80's	Mid 90's
Custom Scalar Systems	Mid 80's	Mid 90's	Mid 2000's
Commodity Cluster	Mid 90's	Mid 2000's	Mid 2010's ???
BG/L or BG/P	Mid 2000's	Mid 2010's ???	Mid 2020's ???

Outline

- Mannheim Supercomputer Statistics & Top500 Project Start in 1993
- Competition between Manufacturers, Countries and Sites
- My Supercomputer Favorite in the Top500 Lists
- The 31st List as of June 2008
- Performance Development and Projection
- Bell's Law
- ➔ Supercomputing, quo vadis?
- Top500, quo vadis?

Supercomputing, quo vadis?

European Countries / Systems

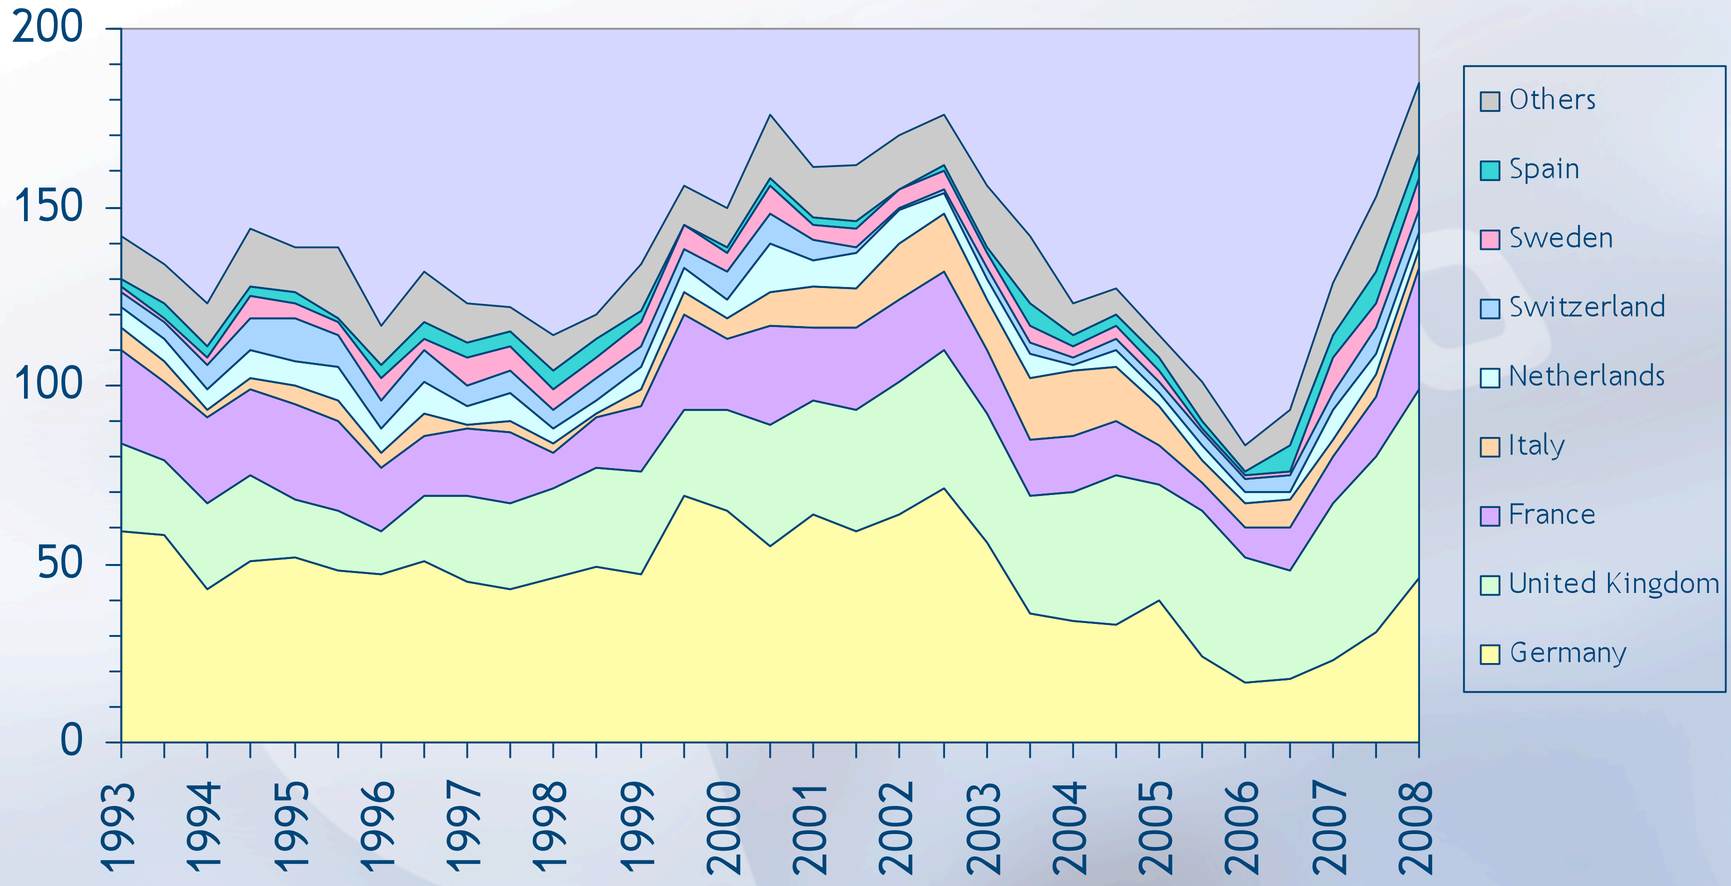

Supercomputing, quo vadis?

Countries/Systems (June 2008)

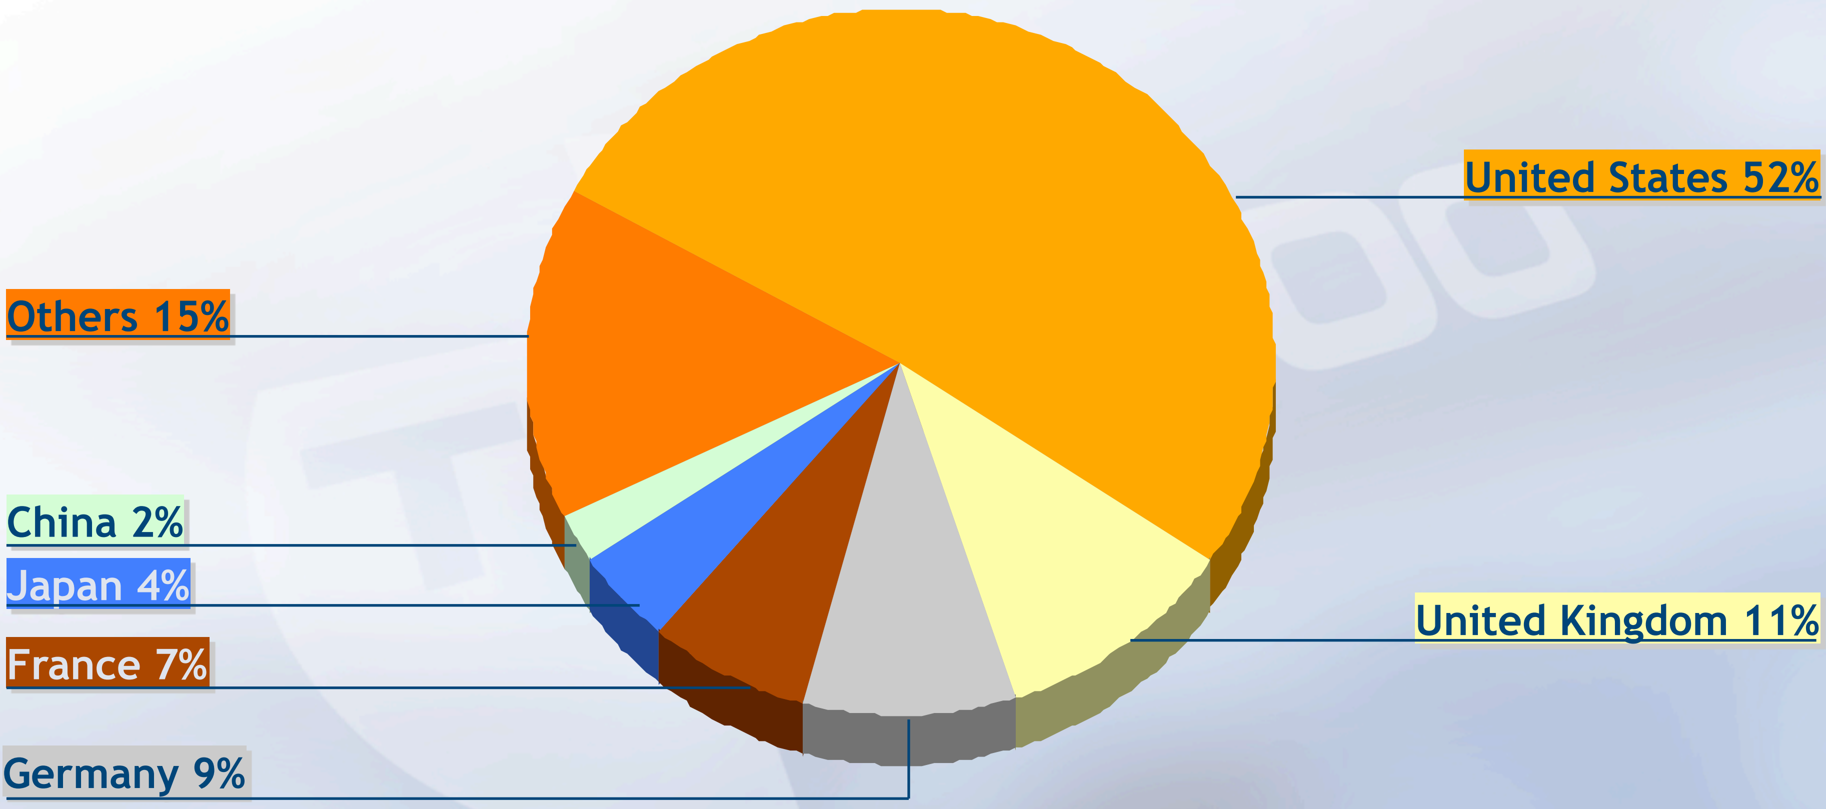

Countries/Performance (June 2008)

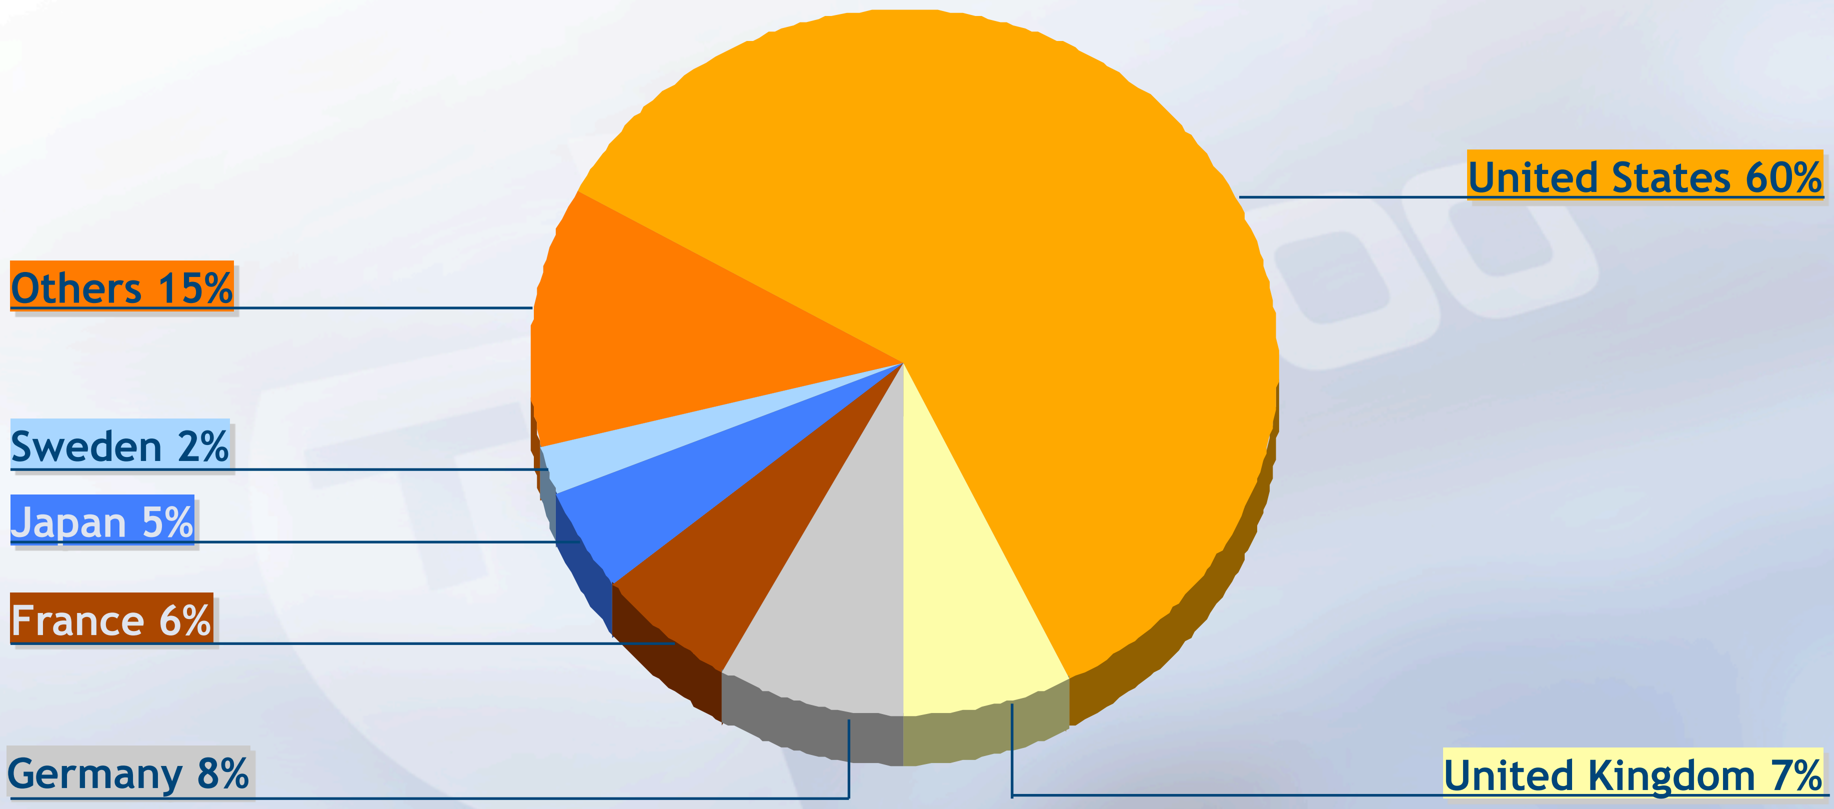

Supercomputing, quo vadis?

Countries/Systems (June 2008)

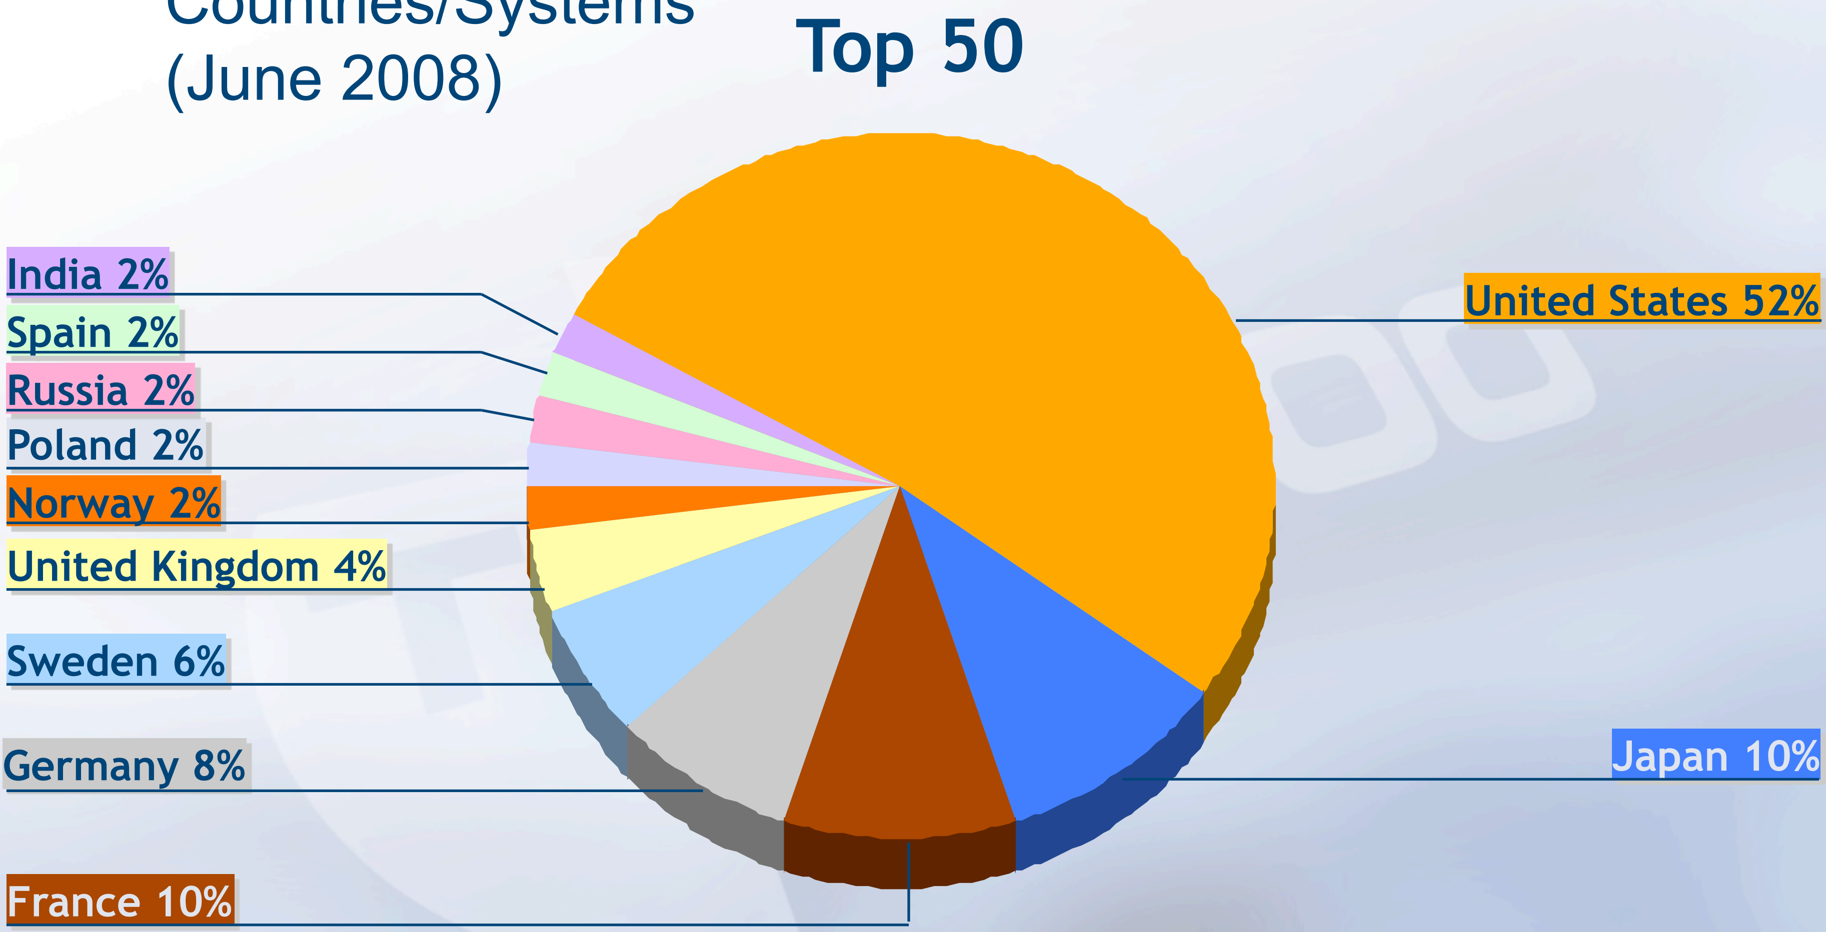

Supercomputing, quo vadis?

Continents/Systems

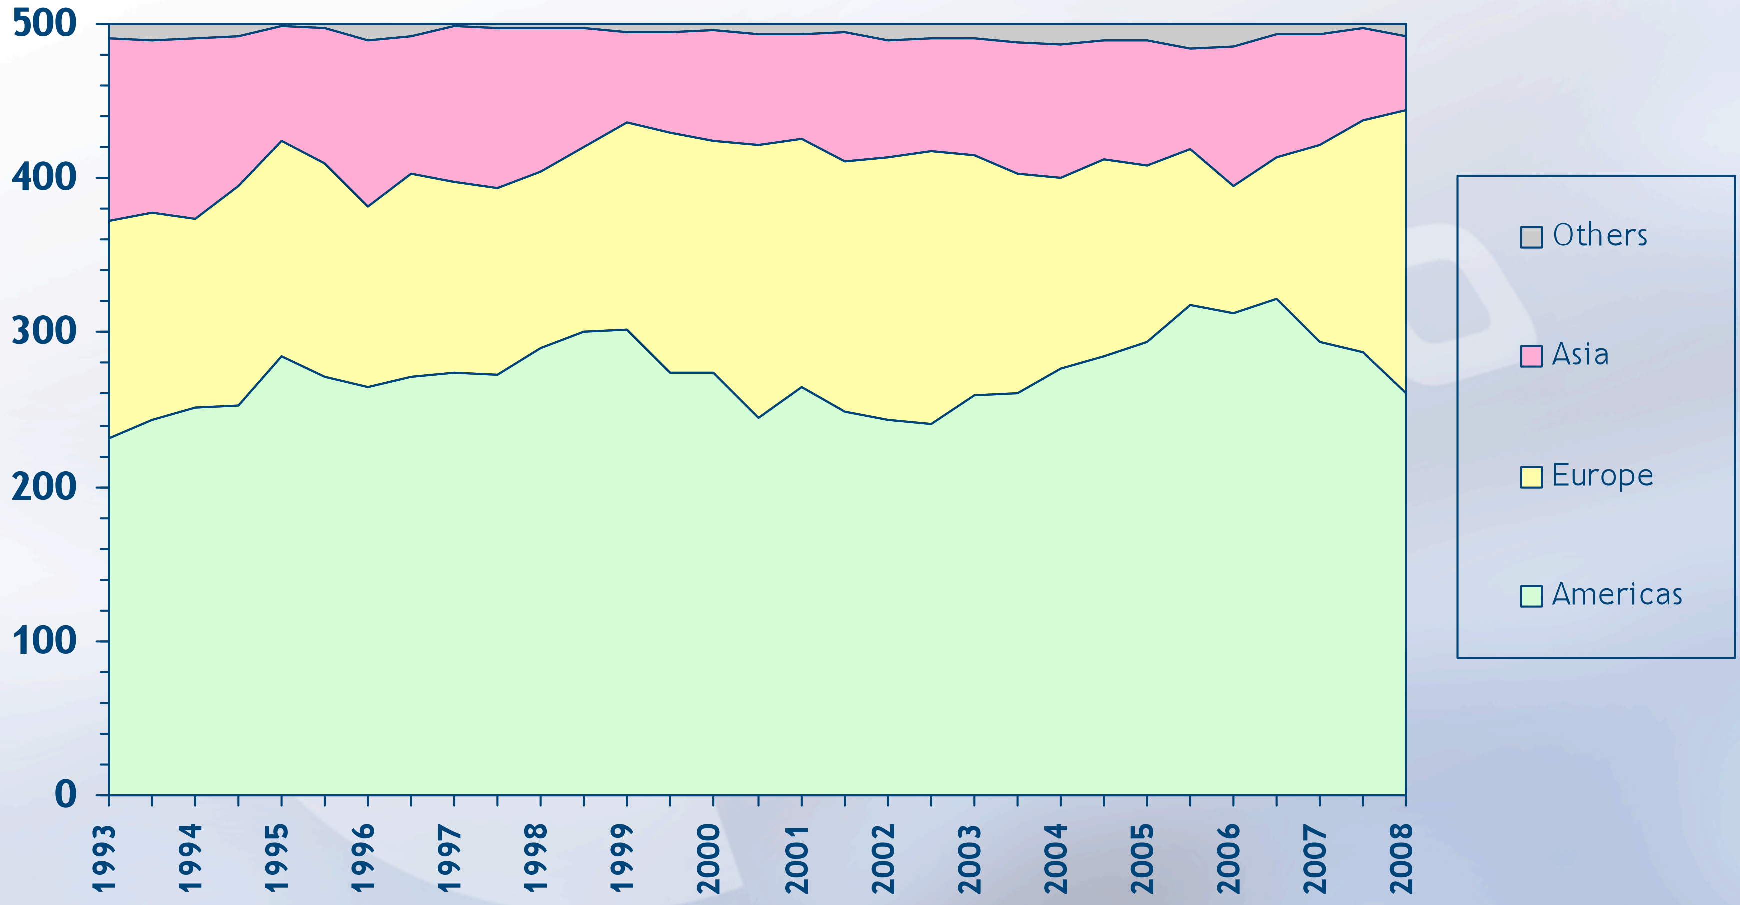

Asian Countries / Systems

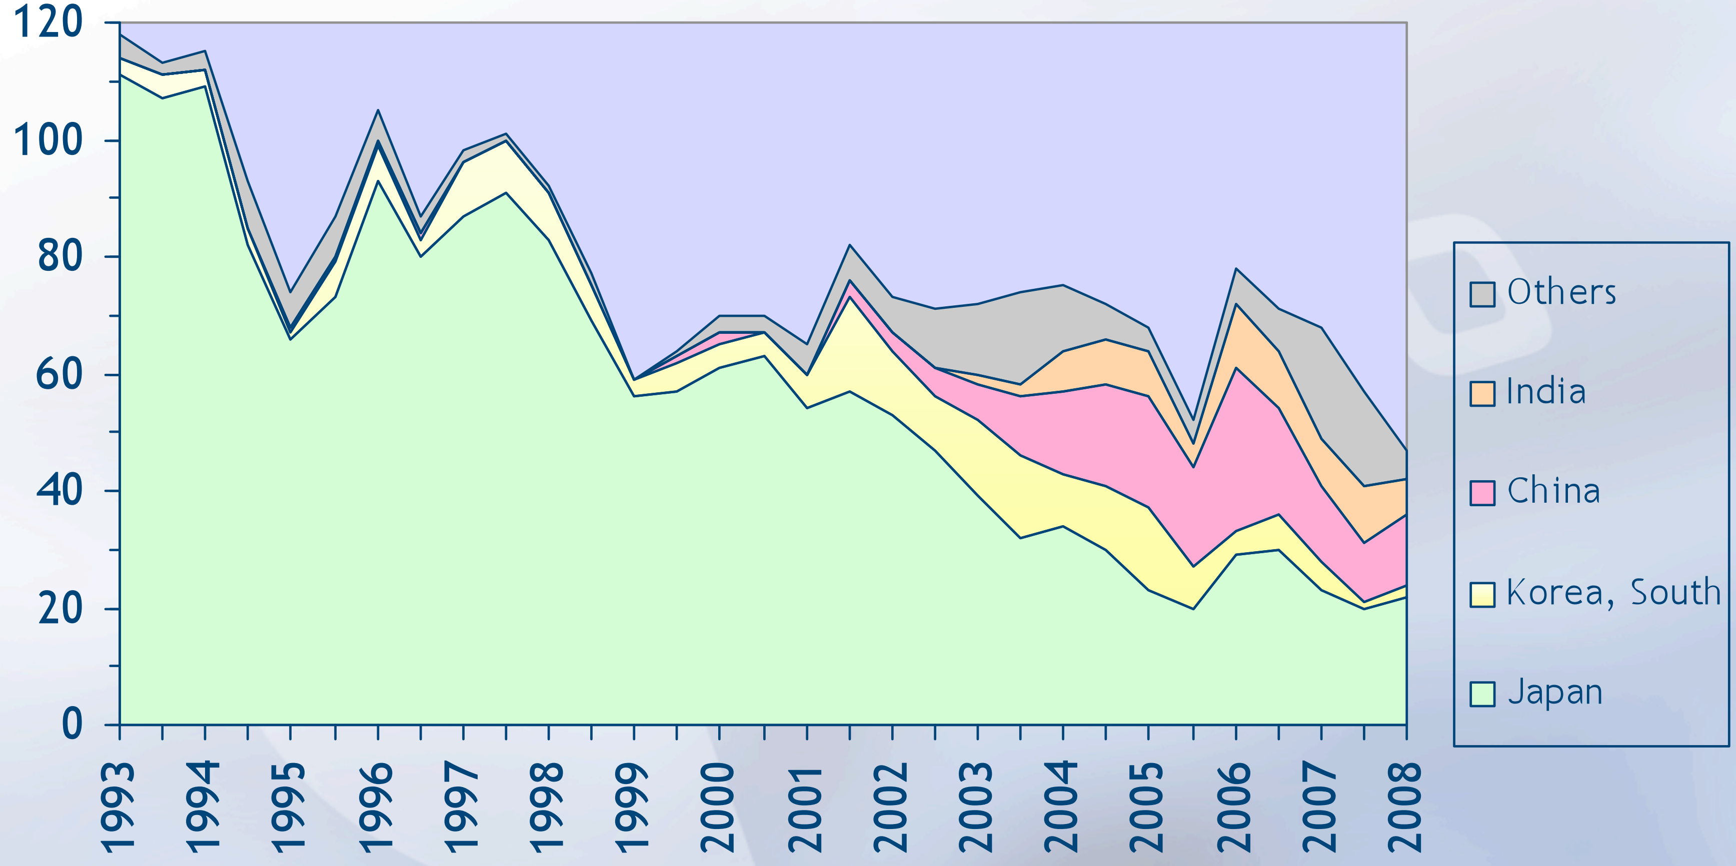

Supercomputing, quo vadis?

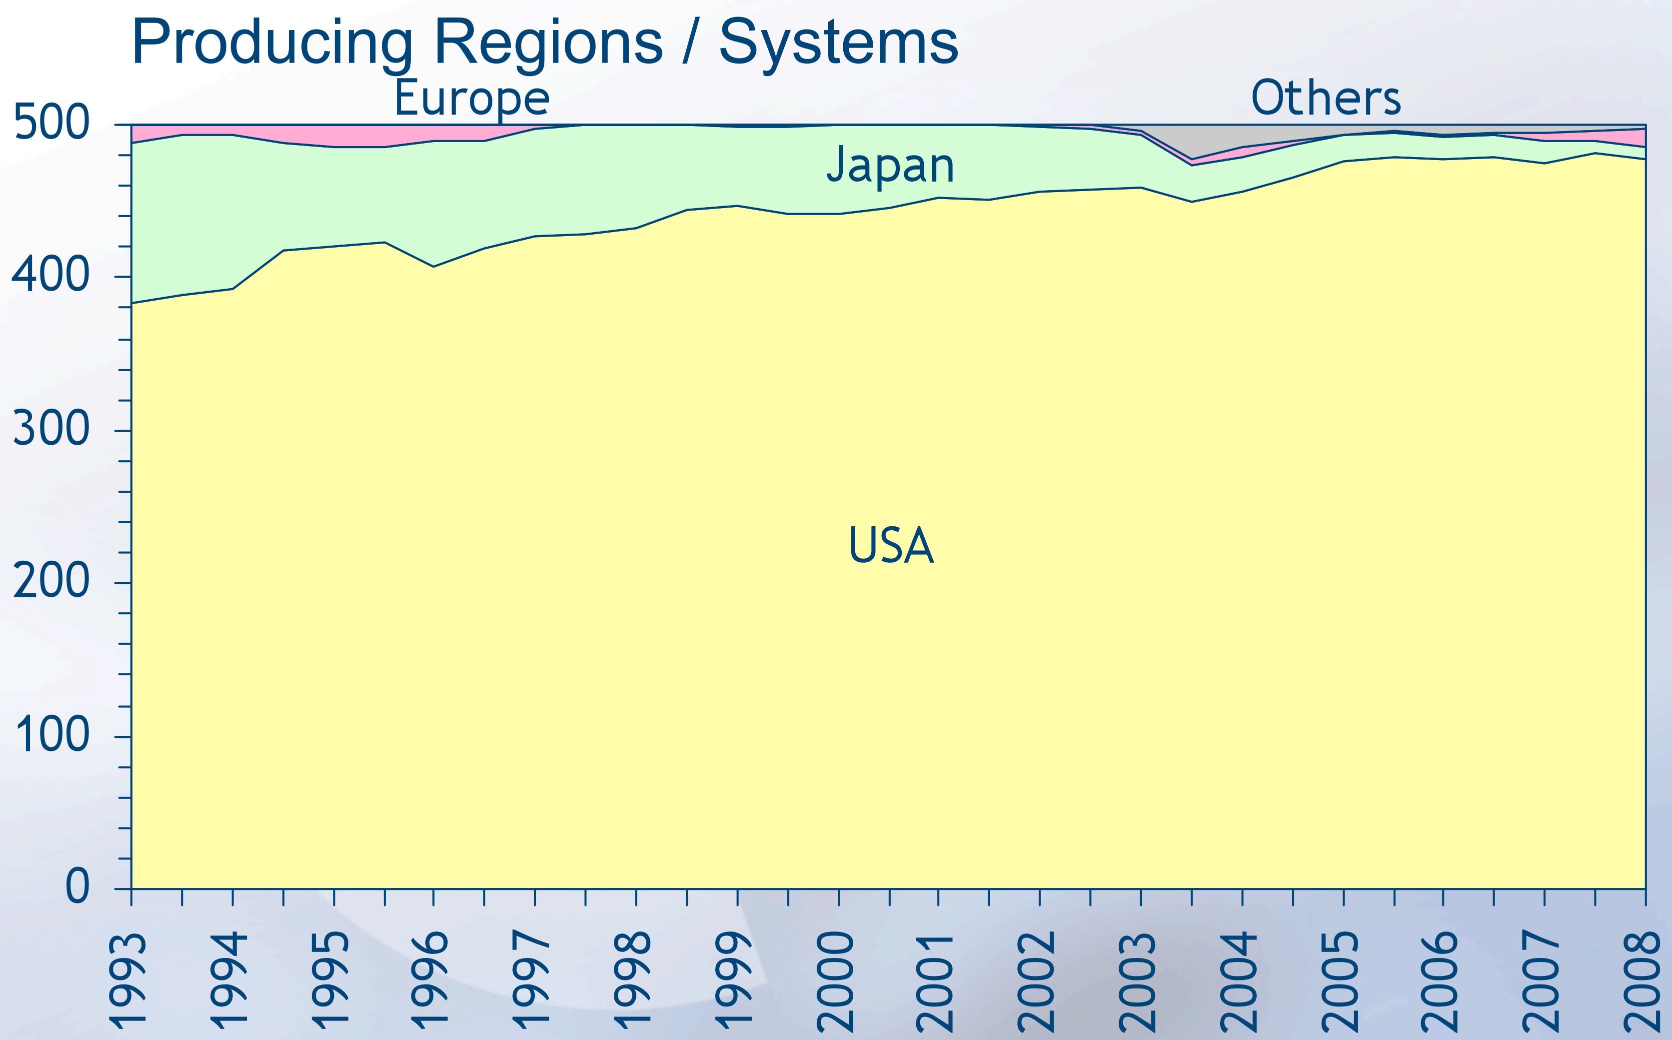

Outline

- Mannheim Supercomputer Statistics & Top500 Project Start in 1993
- Competition between Manufacturers, Countries and Sites
- My Supercomputer Favorite in the Top500 Lists
- The 31st List as of June 2008
- Performance Development and Projection
- Bell's Law
- Supercomputing, quo vadis?
- ➔ Top500, quo vadis?

Top500, quo vadis?

www.top500.org

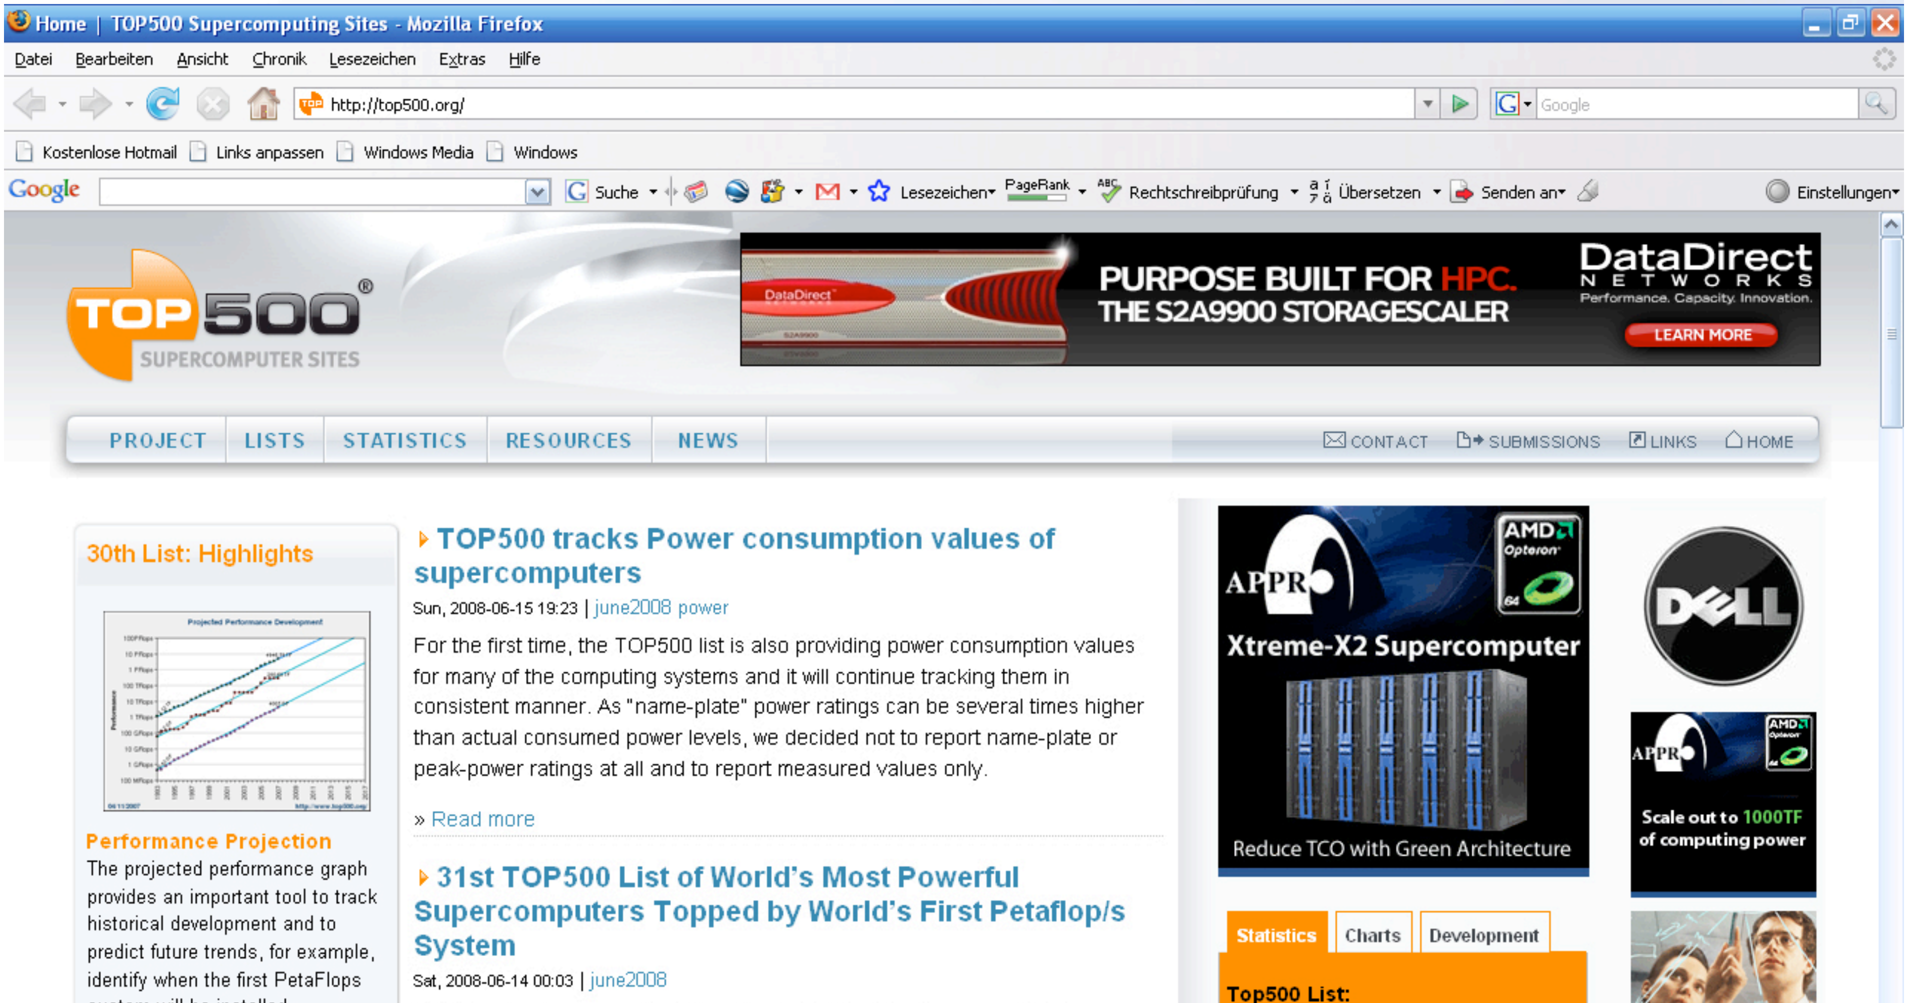

Top500, quo vadis?

Summary after Fifteen Years of Experience

- TOP500 proofed itself by correcting the Mannheim Supercomputer Statistics.
- It is simplistic, but (or because of it) it gets trends right.
- It does not easily allow to track market size
- It's inventory based.
  - This smoothes seasonal fluctuation.
  - Turn-over very high, so it still reflects recent developments.

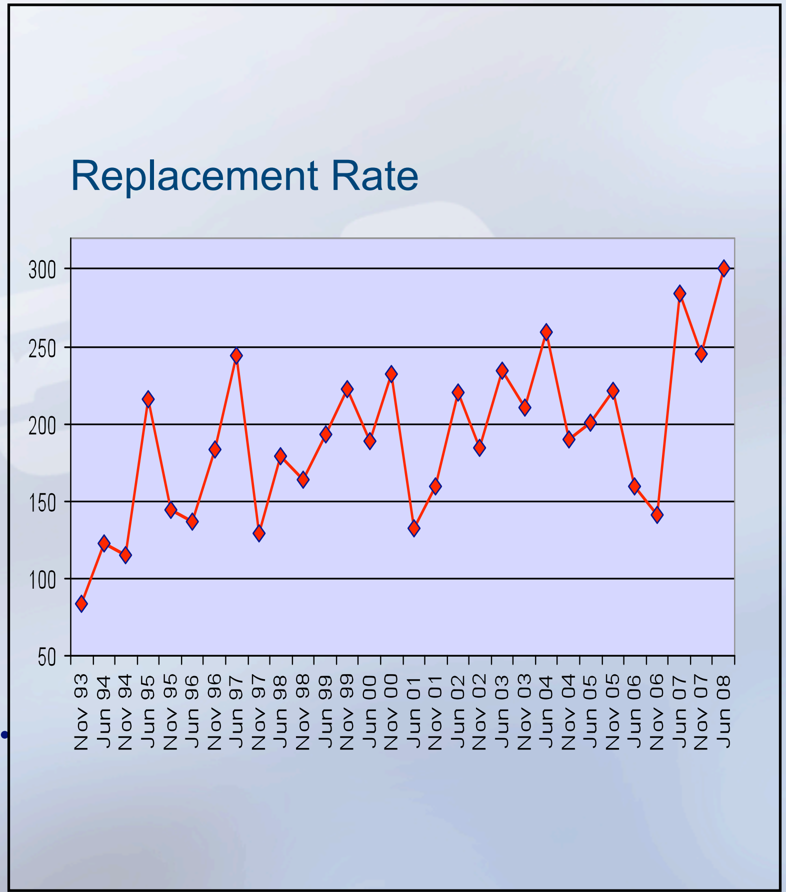

Motivation for Additional Benchmarks

Linpack Benchmark

- **Pros**
  - One number
  - Simple to define and rank
  - Allows problem size to change with machine and over time
  - Allowing Competitions
- **Cons**
  - Emphasizes only "peak" CPU speed and number of CPUs
  - Does not stress local bandwidth
  - Does not stress the network
  - No single number can reflect overall performance

- Clearly need something more than Linpack
- HPC Challenge Benchmark and others

HPC Challenge Benchmarks and the TOP500

Jack Dongarra (UTK/ORNL)

Presented at ISC'06
June 27-30 2006

Office of Science
U.S. Department of Energy

1

Top500, quo vadis?

HPC Challenge Benchmark

Consists of basically 7 benchmarks;
- Think of it as a framework or harness for adding benchmarks of interest.

1. HPL (LINPACK) — MPI Global (Ax = b)
2. STREAM — Local; single CPU
   *STREAM — Embarrassingly parallel
3. PTRANS ($A \leftarrow A + B^T$) — MPI Global
4. RandomAccess — Local; single CPU
   *RandomAccess — Embarrassingly parallel
   RandomAccess — MPI Global
5. BW and Latency – MPI
6. FFT - Global, single CPU, and EP
7. Matrix Multiply – single CPU and EP

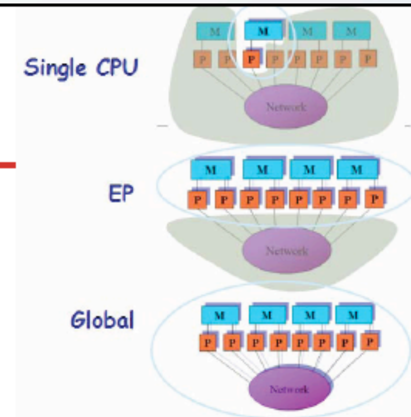

name	kernel	bytes/iter	FLOPS/iter
COPY:	a(i) = b(i)	16	0
SCALE:	a(i) = q*b(i)	16	1
SUM:	a(i) = b(i) + c(i)	24	1
TRIAD:	a(i) = b(i) + q*c(i)	24	2

Random integer read; update; & write

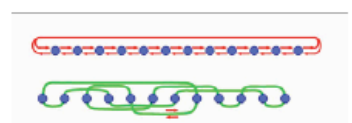

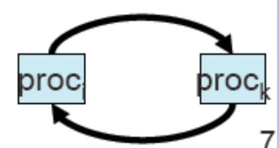

7

Vielen Dank für Ihr Aufmerksamkeit!

Infos/Folien über:

hans@meuer.de