

PEE-5793 Commercial Applications in High Performance Computing

Edward David Moreno
<http://www.lsi.usp.br/~edmoreno>

Lec-02 Top-500 Supercomputer Sites (06/99) Jack Dongarra et Al.

<http://www.netlib.org/benchmark/top500.html>

1

Introduction and Objectives

Statistics on high-performance computers are of major interest to manufacturers, users, and potential users. These people wish to know not only the number of systems installed, but also the location of the various supercomputers within the high-performance computing community and the applications for which a computer system is being used.

Such statistics can facilitate the establishment of collaborations, the exchange of data and software, and provide a better understanding of the high-performance computer market.

New statistics are required that reflect the diversification of supercomputers, the enormous performance difference between low-end and high-end models, the increasing availability of **massively parallel processing (MPP) systems**, and the strong increase in computing power of **the high-end models of workstation suppliers (SMP)**.

2

The Linpack Benchmark (1)

In the present list (which we call the TOP500), we list computers ranked by their performance on the **LINPACK Benchmark**.

As a yardstick of performance we are using the 'best' performance as measured by the LINPACK Benchmark. LINPACK was chosen because it is widely used and performance numbers are available for almost all relevant systems.

The benchmark used in the LINPACK Benchmark is to solve a dense system of linear equations. For the TOP500, we used that version of the benchmark that allows the user to scale the size of the problem and to optimize the software in order to achieve the best performance for a given machine.

This performance does not reflect the overall performance of a given system, as no single number ever can. It does, however, reflect the performance of a dedicated system for solving a dense system of linear equations.

Since the problem is very regular, the performance achieved is quite high, and the performance numbers give a good correction of peak performance.

3

The Linpack Benchmark (2)

In an attempt to obtain uniformity across all computers in performance reporting, the algorithm used in solving the system of equations in the benchmark procedure must conform to the standard operation count for LU factorization with partial pivoting.

In particular, the operation count for the algorithm must be $\frac{2}{3}n^3 + O(n^2)$ floating point operations.

This excludes the use of a fast matrix multiply algorithm like "Strassen's Method"

This is done to provide a comparable set of performance numbers across all computers. If in the future a more realistic metric finds widespread usage, so that numbers for all systems in question are available, we may convert to that performance measure.

4

The TOP500 List (1)

The following table shows the 500 most powerful commercially available computer systems known to us. To keep the list as compact as possible, we show only a part of our information here:

• Nworld -	Position within the TOP500 ranking
• Manufacturer -	Manufacturer or vendor
• Computer -	Type indicated by manufacturer or vendor
• Installation Site -	Customer
• Location -	Location and country
• Year -	Year of installation/last major update
• Field of Application	Applications
• #Proc. -	Number of processors
• Rmax -	Maximal LINPACK performance achieved
• Rpeak -	Theoretical peak performance
• Nmax -	Problemsize for achieving Rmax
• N1/2 -	Problem size for achieving half of Rmax

5

The TOP500 List (2)

If Rmax from Table 3 of the LINPACK Report is not available, we use the TPP performance given in Table 1 of the LINPACK Report for solving a system of 1000 equations.

To use a consistent yardstick for all systems, we do not use results achieved by advanced parallel algorithm as defined in [Dongarra,94].

Performance of Various Computers Using Standard Linear Equations
Software. TR-CS-89-85-94, Univ. of Tennessee, U.S.A.

In case of the CRAY T90, C90 and J90 systems we had to use older table 3 or table 1 results.

In a few cases we interpolated between two measured system sizes or we scaled by cycle times.

For models where we did not receive the requested data, the performance of the next smaller system measured is used.

As the TOP500 should provide a basis for statistics on the market of high-performance computers, we limit the number of systems installed at vendor sites.

Examples: The first 10 in the Top-500 List

#	Manufacturer Computer	Rmax	Installation Site	Country	Year	Area of Installation	# Proc	Rpeak	Nmax	N1/2
1	Intel ASCI Red	2121.3	Sandia National Labs Albuquerque	USA	1999	Research	9472	3154	251904	66000
2	SGI ASCI Blue Mountain	1608	Los Alamos National Laboratory	Los Alamos USA	1998	Research	6144	3072	37440	
3	SGI T3E1200	891.5	Government	USA	1998	Classified	1084	1300.8	259200	26400
4	Hitachi SR8000/128	873.6	University of Tokyo	Tokyo Japan	1999	Academic	128	1024	120000	16000
5	SGI T3E900	815.1	Government	USA	1997	Classified	1324	1191.6	134400	26880
6	SGI ORIGIN 2000	250 MHz	690.9	Los Alamos National Laboratory/ACL	Los Alamos USA	1999	Research	2048	10	
7	SGI T3E900	552.92	United Kingdom Meteorological Office	Bracknell UK	1997	Research Weather	876	788.4		
8	IBM SP Silver 547	IBM Poughkeepsie	USA	1998	Vendor Energy	1952	1296	244000	58000	
9	SGI T3E900	515.1	Naval Oceanographic Office (NAVOCEANO)	Bay Saint Louis USA	1999	Research Weather	81			
10	SGI T3E1200	509.9	UK Centre for Science	Manchester UK	1998	Academic	612	734.4		

Data Analysis of Top500 (Nov-1998)

Industry
206


Research
138

Academic
91

Classified
46

Others
19

Division by Applications




Industry Research Academic Classified Others

8

Industrial Sector using HPC

Ano	Industry	Research	Academic	Classified	Others
96	136 (20)	159 (32)	138 (31)	50 (10)	17 (3)
97	156 (31)	177 (35)	118 (24)	26 (5)	23 (5)
98	206 (41)	138 (28)	91 (18)	46 (9)	19 (4)
99	241 (50)	123 (25)	71 (15)	44 (9)	05 (1)


Division by Applications



1
9
9
9

9

Division of Industrial Applications (1998)



Exemplos de Empresas/Industrias Usando HPC



- 1. Tele** 22.8% BCT - British Columbia Telec., Deutsche Telekom AG, Korean Tel., NTT - Nippon Telegraph and Telephone, ATT CFO, Bell Canada, GTE Comm., O-tel-o, Sprint, GTE Comm., France Telecom, Telcel, Viag Interkom, Cincinnati Bell Information Systems (CBIS), Telecom Denmark, MCI, Ameritech, Worldcom, US West Comm.,
- 2. Fina** 14.5% Charles Schwab, Chase Manhattan, Commerzbank, Deutsche Morgan Grenfell, Prudential Insurance, SMVG, KPMG Peat Marwick, Financial Corporation, American Express, Bear Stearns, Dresdner Bank, National Reserve Bank (Moscow), Citicorp, NatWest Markets, Toronto Stock Exchange, Union Bank of Switzerland,
- 3. Elec** 10.2% Chip Manufacturer, Philips Research, Motorola, Lincoln Electric, Samsung, Texas Instruments, Raytheon, FDK Corporation, Micron Technology, E-Systems & Raytheon, Toshiba, Cisco, Lucent Technologies,
- 4. Geop** 10.2% Exxon, P-Petroleum Company, Baker Hughes, Vastar, EP Company, Western Geophysical, Veritas DGC, Ensign, Mobil, Texaco,
- 5. Univ** 09.2% KIST, Universities, Thyssen, ICAM-Alcatel, UKI IS, Montgomery Ward, Origin IT, ATAC, Xerox, Miele GmbH, Moody,

11

- 6. Data** 06.3% State Farm, Oracle Corporation, IBM, Sears, Allstate Insurance, Daiei Information, Kwuiksav, W.W.Grainger, American Express, Krupp Hoesch Info., Federal Express
- 7. Auto** 05.8% Toyota Central Research Development, Audi AG, TRW, Automotive Manufacturing, EDS/General Motors, BMW, Ford Motor Company, Nissan Motor,
- 8. Trans** 4.0% Deutsche Bahn AG, Lufthansa, American Airlines, EDS de Mexico, Emery Worldwide, Sabre Group, Delta Airlines,
- 9. Aero** 3.4% CSC, Lockheed Martin, Hughes Space Communication,
- 10. Chem** 2.9% BASF, Owens Corning, Bayer AG, Witco, Osaka Gas. Ltd.,
- 11. Manu** 2.4% Allied Signal Federal, Kawasaki Heavy Industrie, Medline Industries, Becton Dickinson,
- 12. Serv** 1.9% Lexis Nexis, Sylvest Management Systems, R.L. Polk, RAG Informatik,
- 13. Phar** 1.4% NRC (Nichols Research Corp.) Vertex Pharm., Government, Eli Lilly and Company, Recruit (www), eBAY (www), Tennessee Valley Authority (NTA)-energy, Bristol-Myers Squibb, Enron Capital (energy),

Energy, Defense, WWW
- 14. Imag.** 1.0% SAP, HTC, MediaOne,
Media, Software


12



	1998	1999	USA	Europe	Japan	Others
1. Telecomm	23	16	31	10	4	2
2. Finance	14	14	11	16	2	1
3. Electronics	11	6.6	17	1	2	1
4. Geophysics	11	6.2	16	5		
5. Univ.-Indus	10	28	08	10		1
6. Database	7.0	5.4	8	3	2	
7. Automotive	6.0	9.5	3	2	6	
8. Transportation	4.0	2.5	5	2		1
9. Aerospace	4.0	2.5	6	1		
10. Chemistry	3.0	3.0	3	2	1	
11. Manufacturing	3.0	1.0	3		2	
12. Service	2.0	2.5	3	1		
13. Pharmaceuticals	1.0	1.7	10			
Energy, Defense, WWW						
14. Image Processing	1.0	1.5	2	1		
Media, Software						

15

Industry by area (1998)




16

Companies & Industry		1998	SGI	IBM	SUN	NEC	HP	FUJ	HI
1. Telecomm		23	1	18	29				
2. Finance		14	2	7	22				
3. Electronics		11	6	6	7		1	1	
4. Geophysics		11	8	9	2	1			
5. Univ.-Indus		10	4	13	2				
6. Database		7.0	1	3	9				
7. Automotive		6.0	8	1		1		1	
8. Transportation		4.0		1	7				
9. Aerospace		4.0	5		2				
10. Chemistry		3.0		3	2			1	
11. Manufacturing		3.0	1	1	3				
12. Service		2.0	1		3				
13. Pharmaceuticals		1.0	3	1	6				
Energy, Defense, WWW									
14. Image Processing		1.0		1	1		1		
Media, Software									

17

Industry by Machines

Division by Machines



18

Machines

SGI	T3E1200/900	T3D-MC-512	SGI-Origin-2000
IBM	ASCI-Blue-Mountain	SP-P2SC	SP-PC604
SUN	HPC-10000		
NEC	SX-4/128M4	SX-4/128H4	
HP	Exemplar X-Class		
FUJITSU	Numerical-Wind-Tunnel	VPP700	
HITACHI	CP-PACS/2048	SR2201/1024	
INTEL	ASCI-Red	XP/S140	

19


HPC in Research

20

Research Sector using HPC

Ano	Industry	Research	Academic	Classified	Others
96	136 (20)	159 (32)	138 (31)	50 (10)	17 (3)
97	156 (31)	177 (35)	118 (24)	26 (5)	23 (5)
98	206 (41)	138 (28)	91 (18)	46 (9)	19 (4)
99	241 (50)	123 (25)	71 (15)	44 (9)	05 (1)

Division by Applications



1
9
9
9


21

Division of Research's Applications (1998)

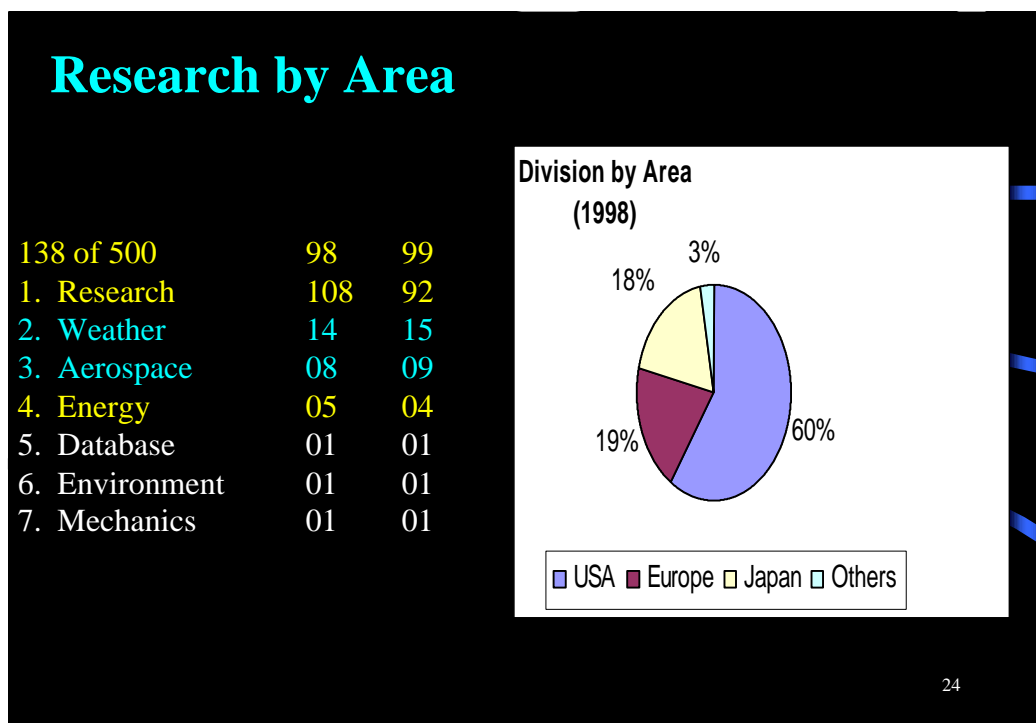
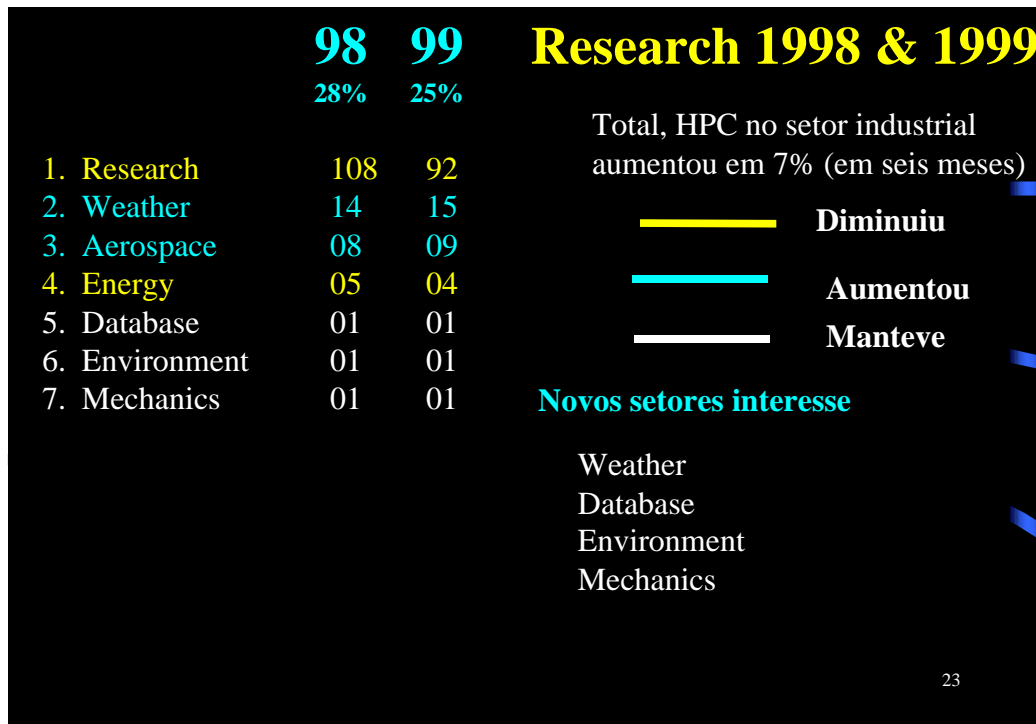
138 of 500 28%

1. Research 108
2. Weather 14
3. Aerospace 08
4. Energy 05
5. Database 01
6. Environment 01
7. Mechanics 01

Apps in Research




22



Research & Companies (1998)

Division by
Machines




25

Academic & HPC (area)

1998: 18.2%

1999: 14.2%

Division by Area
(1998)



USA Europe Japan Others


26

Academic & HPC (Companies)

1998: 18.2%

1999: 14.2%

Division by
Machines




27

Data Information (1997)


28

Performance & Development 1997




29

Architectures & Systems (1997)




30

CPU Technology & Systems(1997)




31

Node technology & Systems(1997)




32

Manufactures & Systems (1997)




33

Manufactures & Performance (1997)




34

Continents & Systems (1997)




35

Systems Installed (1997)



36


R max Installed (1997)



1999:

57.69 Tflops (R_{peak}) & 39.10 Tflops (R_{max})

Applications areas & Systems (1997)




98

99

R: 138
I: 206
A: 91
C: 46
O: 19

123
241
71
44
5

Application Areas & Performance (1997)




98 99

?

39

Systems Installed (1997)



98


99

R: 138
I: 206
A: 91
C: 46
O: 19

123
241
71
44
5

40

Rmax Installed (1997)



1998 1999


?????

??????

????

41

Number of Systems Installed (1997)



42

Conclusions

43

- The architectures installed at industrial sites changed from vector systems to a substantial number of MPP systems. This change reflects the fact that parallel systems are ready for commercial use and environments.
- The most successful companies (IBM and SGI) are selling well to industrial customers. Their success is built on the fact that they are using standard workstation technologies for their MPP nodes. This approach provides a smooth migration path for applications from workstations up to parallel machines.
- The maturity of these advanced systems and the availability of key applications for them make the systems appealing to commercial customers. Especially important are database applications, since these can use highly parallel systems with more than 128 processors.

44

In 1996,
the applications typically were numerically intensive applications,
for example,

- automotive applications,
- aerospace studies,
- chemical and pharmaceutical studies,
- electronics,
- energy research,
- geophysics and oil applications, and
- weather prediction.

45

**Recently, however, industrial systems in the TOP500 have been used
for new application areas.**

These include

- database applications,**
- finance applications,**
- image processing, and**
- WWW servers.**

46

From the present ten releases of the TOP500 we see the following trends:

- The number of industrial customers in the TOP500 has risen steadily since June 1995.
- The most successful companies (IBM and SGI) are selling disproportionately well in the industrial market.
- The average system size at industrial sites is increasing strongly.
- Database applications are the most important and most successful new application area for supercomputers.
- Distributed-memory systems are being installed at industrial sites in reasonable numbers.
- The United States is the clear world leader in the industrial use of HPC systems.

47

Therefore:

The strong increase in the number of processors in systems at industrial sites is another major reason for the rise of industrial sites in TOP500.

The industry is ready to use bigger parallel systems than in the past.

Then, Commercial and Industrial Applications are a good and Interesting area to work
(MONEY ??? Knowledge ++++ Books & Papers +++++⁴⁸