ELSEVIER

Future Generation Computer Systems 12 (1997) 461–474

FGCS
Future Generation Computer Systems

Changing technologies of HPC

J.J. Dongarra[a,b,*], H.W. Meuer[c,1], H.D. Simon[d,2], E. Strohmaier[a,3]

[a] *Computer Science Department, University of Tennessee, Knoxville, TN 37996, USA*
[b] *Mathematical Science Section, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA*
[c] *Computing Center, University of Mannheim, D-68131 Mannheim, Germany*
[d] *NERSC, Lawrence Berkeley National Laboratory, 50B-4230, Berkeley, CA 94720, USA*

Abstract

In 1993 for the first time a list of the top 500 supercomputer sites world-wide was made available. The Top500 list allows a much more detailed and well-founded analysis of the state of high performance computing. Previously data such as the number and geographical distribution of supercomputer installations were difficult to obtain, and only a few analysts undertook the effort to track the press releases by dozens of vendors. With the Top500 report now generally and easily available it is possible to present an analysis of the state of High Performance Computing (HPC). This paper summarizes some of the most important observations about HPC as of late 1996, in particular the continued dominance of the world market in HPC by the US, the market penetration by commodity microprocessor based systems, and the growing industrial use of super-computers.

Keywords: High performance computing; HPC technology; Supercomputer market; Supercomputer technology

1. Introduction

In the past years the field of High Performance Computing (HPC) faced a major move of their building blocks – the computing nodes – away from proprietary designs towards nodes built out of workstation boards. This movement came along with the success of companies like Silicon Graphics and IBM. As the other companies acting in this field are also moving to CMOS as basic technology these two companies along with HP/Convex are building their HPC systems not only with CPUs but with boards "off-the-shelf" from successful workstation families. Major signs for the success of this approach are not only the pure number of systems they can sell, but the percentage of systems they are able to sell to industrial users. We will discuss in this paper the different developments based on the TOP500 lists of supercomputer sites available since June 1993 [1] and which, for the first time, provide a reliable base for a well-founded analysis of the high performance computing field. Reports about the situation in 1993, 1994 and 1995 have been published before [2–5].

* Corresponding author. E-mail: dongarra @cs. utk.edu.
[1] E-mail: meuer@rz. uni-mannheim.de.
[2] E-mail: simon@nersc. gov.
[3] E-mail: erich@cs. utk.edu.

PII S0167-739X(96)00031-3

At the end of 1996 the top position of the TOP500 was held by the CP-PACS system which is a specialized version of Hitachi's SR2201 systems. The list itself underwent quite a change in 1996. Ten new architectures entered the list, but most of them can be seen as first models of a new generation of parallel systems. SGI and Cray, now one company, started to install the Cray based T3E system and the SGI based Origin2000 system which is also a distributed memory system. NEC introduced quite successfully the NEC SX-4 with a cluster approach and with excellent price–performance ratio. Fujitsu continues its VPP architecture in the VPP700 and VPP300 now completely based on CMOS technology. HP/Convex delivered the SPP2000 and SPP1600 as follow-up of the SPP1000/SPP1200 again with virtual shared memory and also with a hierarchical design. Many new systems appear in top positions of the list, many new T3Es from SGI/Cray, and now also the big SP2s as performance values are available for them. Some of the newcomers already are present on the first pages of the TOP500.

Looking at the computing power of the individual machines present in the Top500 and the evolution of the total market size, we plot the performance of the systems at positions 1, 10, 100 and 500 in the list as well as the total accumulated performance. In Fig. 1 the curves of position 100 and 500 show on the average an increase of a factor of two within one year. The curves for position 1, 10 and for the accumulated performance, however, show only a factor of 1.8 increase per year.

We now look at the replacement rate at which new systems enter the list and which systems will be omitted because of their performance being too small. We show in Table 1 the positions of the systems which were at positions 500, 100, and 10 in one of the last eight lists. Looking at Table 1 we see that during the seven revisions of the list, 91% of the entries were removed. This gives on the average a replacement rate of 29% for each new list, thus only the first 71% of the list will be present in the successive issue of the list half a year later. Looking close at the data we see an over proportional replacement rate for the June 1995 list. The explanation for this was the big number of SGI Power Challenge and IBM SP2 systems entering the list at this time. This was the first sign for the change in technology used by the HPC community.

A similar and sometimes even slightly higher replacement rate can be seen for positions 100 and 10 in Table 1. This shows that the replacement is not

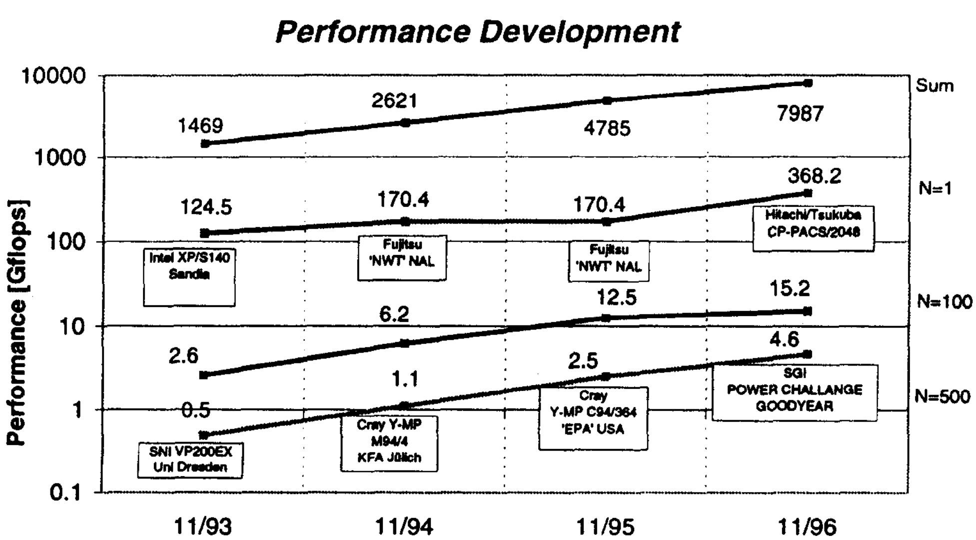

Fig. 1. The performance over time as it can be seen in the Top500.

Table 1.
In order to visualize the replacement rate, the positions over time for entries 10, 100 and 500 are shown for all lists (brackets denote that the kind of system was not present in that list because it was not yet in the market)

System	R_{max}	6/93	11/93	6/94	11/94	6/95	11/95	6/96	11/96
VPP700/46	94.300	(1)	(2)	(3)	(3)	(6)	(6)	(8)	10
SP2/384	66.300	(1)	(2)	(3)	(3)	(6)	3	10	16
CM-5/1056	59.700	1	2	3	3	6	10	13	21
T3D MC512-8	50.800	(3)	(3)	(4)	(4)	10	18	21	31
Hitachi S3800/480	28.400	(5)	(6)	10	10	22	27	37	56
CM-5/256	15.100	7	10	22	23	44	57	72	107
C916/16	13.700	10	15	28	28	50	63	85	125
NEC SX4/8	15.400	(7)	(8)	(20)	(21)	(41)	53	67	100
C916/16	13.700	18	25	40	40	65	78	100	139
SP2/70	12.510	(22)	(28)	(45)	(54)	(81)	100	129	173
VPP500/7	9.650	(26)	(34)	(61)	70	100	173	153	198
Paragon XP/S15	6.250	(34)	48	85	100	137	173	219	322
SP1/64	4.800	(46)	(62)	100	118	172	233	307	486
iPSC/860-128	2.600	80	101	151	200	345	492		
YMP8/8	2.144	100	123	180	235				
SGI PC/18	4.620	(46)	(63)	(105)	(123)	(182)	240	317	500
SPP 1000/32	3.306	(70)	(92)	(136)	(171)	(268)	363	500	
C94/3	2.489	(87)	(110)	160	211	(356)	500		
SGI PC/8	1.955	(127)	(160)	(217)	284	500			
YMP/M4	1.114	251	297	385	500				
XMP4	0.822	319	379	500					
VP200EX	0.472	427	500						
VP200	0.422	500							

only driven by changes in the entry level system market but by more general trends in all market segments.

We now can roughly estimate how long the systems will be present in the Top500. In Table 2 we show the minimal position a system has to have to remain in the Top500 for one to five years. A system that should remain in the Top500 for four years should now be at least around position 30.

Table 2.
The estimate for the minimal position for a system, so that it remains for *n* years in the Top500, based on a replacement rate of 29% each half year

Years	Position
1	252
2	127
3	64
4	32
5	16

2. Geographical distribution

Looking at the Top500 systems installed we see a quite stable distribution over time in Fig. 2. There is an overall upward trend of systems in the US and on the average also a downward trend in Japan. This reflects the fact that Japan is behind in the number of installed SMP systems.

Looking at the total of installed performance in Fig. 3, contrary to the number of systems seen in Fig. 2, Japan is again well ahead of Europe as it was in the past years. This reflects the fact that during the past years several very powerful VPP500 systems were installed in Japan all belonging to the Top500. Taking a closer look at the strong increase of the

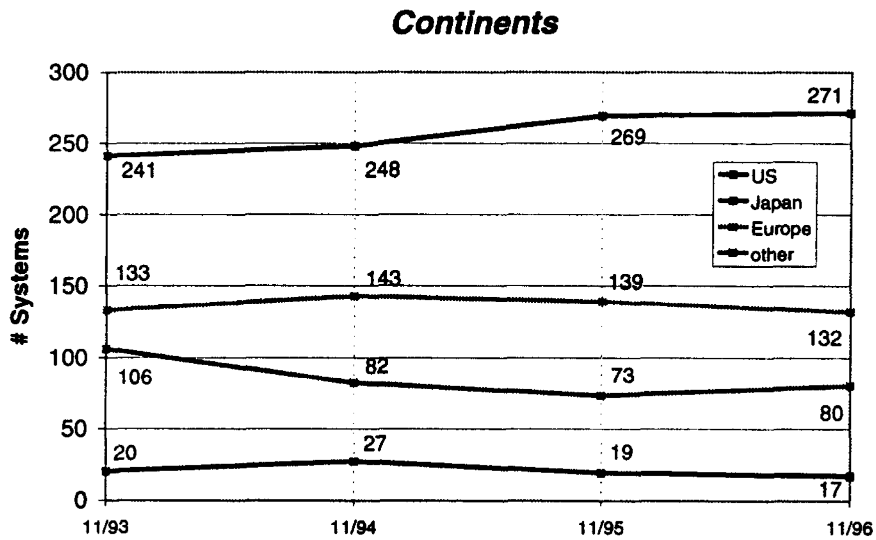

Fig. 2. The geographical distribution of the system counts over time.

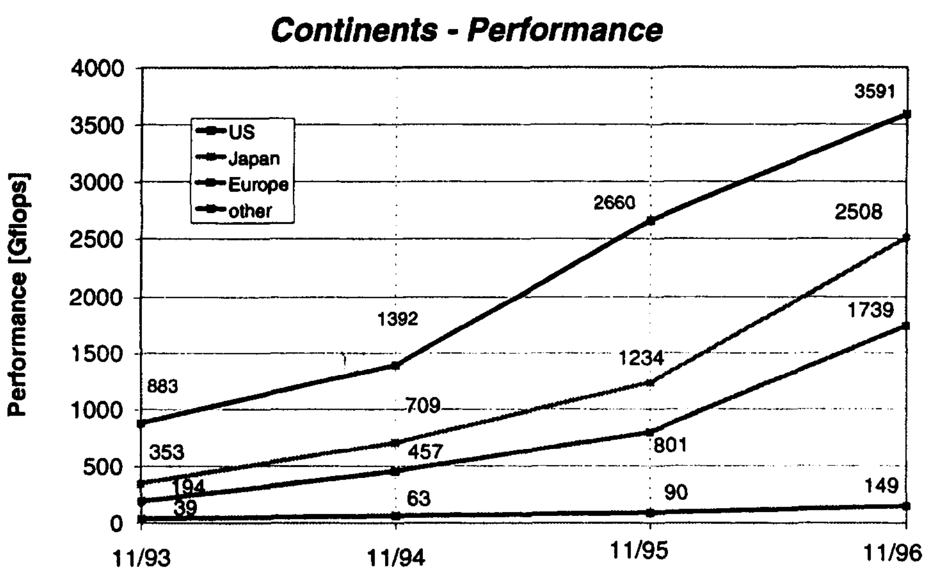

Fig. 3. The geographical distribution of the performance over time.

installed performance in the US during the last year, we find that Cray Research installed 535 Gflop/s, IBM 336, Gflop/s, HP/Convex 54 Gflop/s and SGI only 2 Gflop/s. The share of all other vendors together went down by 6 Gflop/s. In Europe, Cray took a big jump from 281 to 782 Gflop/s while in Japan, Hitachi gained the most installed performance and is now second with 775 Gflop/s behind Fujitsu with 911 Gflop/s. In case of Hitachi, however, the first two of their systems already accumulated 589 Gflop/s.

3. US dominance of the worldwide HPC market

The Top500 continues to demonstrate the dominant position the US assumes in the world both as producer and as consumer of high performance computers. In Table 3 the total number of installed systems in the major world regions is given with respect to the origin of the computers.

If one considers in Table 3 the country of origin then it is striking that 418 out of the Top500 systems are produced in the US, which amounts to 84% of all installed systems. Japan accounts for 14% of the systems, and Europe produces only 2%. The extent of the American dominance of the market is quite surprising, and has been stable from the previous report, when the US share was 85%. For years, in particular in the mid 1980s, there were ominous and ubiquitous warnings that the American supercomputer industry (which was essentially Cray Research at that time) is highly vulnerable to an "attack" by the Japanese vertically integrated computer giants Fujitsu, NEC, and Hitachi. Obviously this has not happened. How much various efforts such as the NSF Supercomputing Initiative in the mid 1980s, or more recently the HPCC Program have contributed to the current vast superiority of the US high performance computing industry, remains to be investigated. It is interesting to note that one view expressed outside the US [6] is that strengthening the US HPC industry and easing the transition to MPP was the *only* rationale for the HPCC Program.

The numbers for Europe are about the same as last year (10 machines in November 1996 versus 15 machines in November 1995). This situation is probably not going to change, since one of the remaining two European vendors (Parsytec) will no longer focus on the HPC market. With lack of immediate access to the newest hardware, and the absence of the close interactions of users with vendors as is prevalent in the US, the best the European High Performance Computing and Networking Initiative can accomplish is maintaining the status quo of Europe as a distant third in high performance computing technologies.

Table 4 is analogous to Table 3, but instead of the number of systems, the aggregate performance in R_{max} (in GFlop/s) is listed. Table 4 shows within the last 12 months an increase in the total number of installed GFlop/s in the US from 2600 GFlop/s in November 1995 to 3591 GFlop/s in November 1996. This is an increase of 35% in one year. At the same time similar growth can be seen in other regions.

In 1995 the same table demonstrates a truly astounding event: within six months the total number of installed GFlop/s in the US increased from 1392 GFlop/s in June to 2660 GFlop/s in November 1995. This is an increase of 92% in only six months. What is more astounding is that this growth did not happen by installing a few very large-machines. Instead a large number of machines were installed, which now occupy medium to lower ranks on the Top500 list. One conclusion from this data is that the HPCC initiative in the US has succeeded in the sense that the infrastructure for HPC is dramatically changing. A large number of institutions now have access to GFlop/s level computing

Table 3.
Geographical distribution where systems are installed and where they are manufactured

Systems manufactured in	Installed in				
	USA	Japan	Europe	Others	Total
USA	261	31	110	16	418
Japan	8	48	15	1	72
Europe	2	1	7		10
Total	271	80	132	17	500

Table 4.
Geographical distribution of the accumulated performance R_{max} (in Gflop/s) showing where it is installed and where it is manufactured

R_{max} (Gflop/s), manufactured in	Installed in				
	USA	Japan	Europe	Others	Total
USA	3464	391	1332	122	5308
Japan	117	2111	365	28	2622
Europe	10	5	42		57
Total	3591	2508	1739	149	7987

Table 5.
Population (in thousands) per Top500 supercomputer

Country	Population (in thousands)	Number of Top500 entries	Population (in thousands), per supercomputer
Switzerland	6813	9	757
USA	255200	266	959
Japan	124500	80	1556
Germany	80250	51	1574
Finland	5008	3	1669
Netherlands	15160	9	1684
Denmark	5158	3	1719
Austria	7776	4	1944
Sweden	8652	3	2884
UK	57700	18	3206
France	57180	17	3364
Australia	17493	5	3490
Canada	27370	5	5474
Spain	38998	5	7800
Italy	57157	6	9526

for machines which cost not much more that $ 1 M. Only five years ago this compute power was accessible only to the elite few institutions being able to spend tens of millions of dollars. We can anticipate exciting times for HPC: more and more people in the US will have access to inexpensive computational modeling tools. It will be worthwhile to examine what this revolution will do to economic productivity measures such as the Gross Domestic Product (GDP) in the US.

In an international comparison one should, however, also consider the relative size of countries and their economies. Here we present a new Top500 set of statistics. In Table 5 we list a measure of the supercomputer density by ranking the top ten countries with the highest number of supercomputers per capita. Population data are from the "Interactive 3D Atlas" and date from 1992. Table 5 shows that on an international comparison most industralized countries are providing about one supercomputer per 1–2.5 million inhabitants. The number of US installations is no longer that dramatically different from the rest of the industrialized countries. It should be mentioned that among the major industrialized nations the big anomaly with respect to supercomputing usage is Italy. In Italy there is only one supercomputer per 9.6 million inhabitants, far below the number of all other western European countries.

4. Market shares of vendors

The shake out of the HPC manufacturers culminated 1996 in SGI buying Cray Research. This merger created a strong new market leader in the HPC arena. Together they are dominating the market with a total share of 44% of the installed systems. However, this is only slightly more than Cray Research had on its own (41%) when we started the Top500 in June 1993. In Fig. 4 we see that Cray Research by itself has gained back the pole position from SGI with which it switches positions if we look at the situation in June 1996. Most of the raise of Cray is due to the 23 early T3E installations in the list. IBM is close second to Cray Research with 25% of systems installed. SGI/Cray and IBM hold together 2/3 of the market. The three Japanese companies Fujitsu, NEC and Hitachi have together 72 (14%) systems in the list. Looking at the changes in the accumulated performances of the different vendors

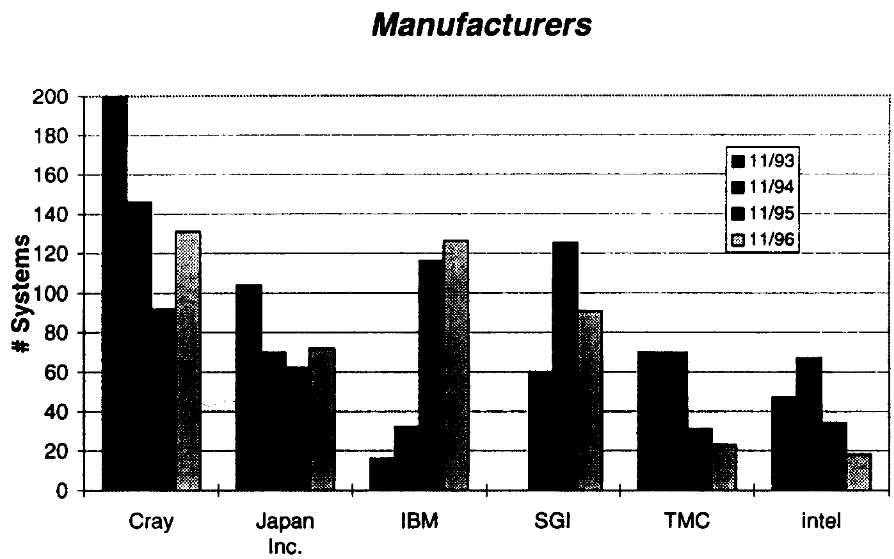

Fig. 4. The market share of the most important vendors over time.

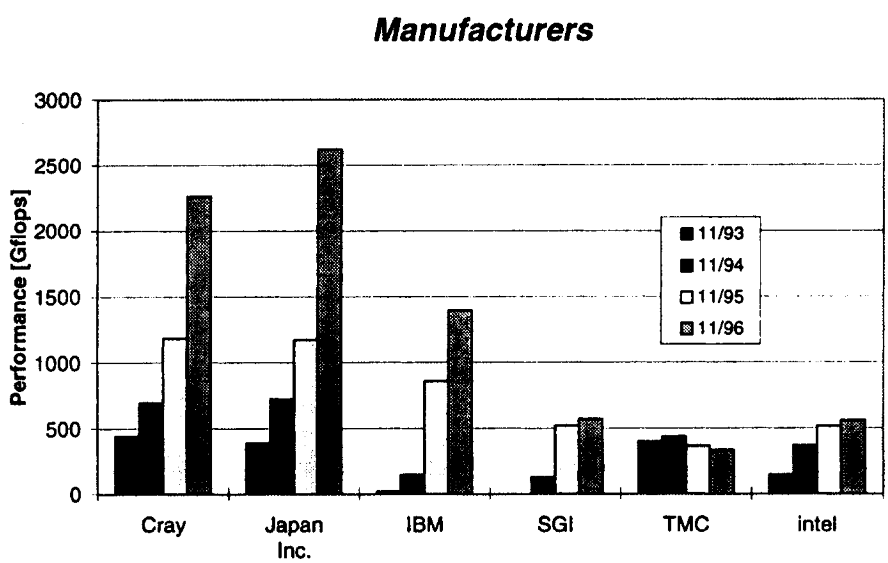

Fig. 5. The market share in performance of the most important vendors over time.

in Fig. 5, we see that the installed performance of Cray made a big jump due to the T3E. The strong increase of the Japanese vendors and IBM is continuing.

5. Architectural changes

The big increase in the number of installed symmetric multiprocessor workstations (SMP) in 1994

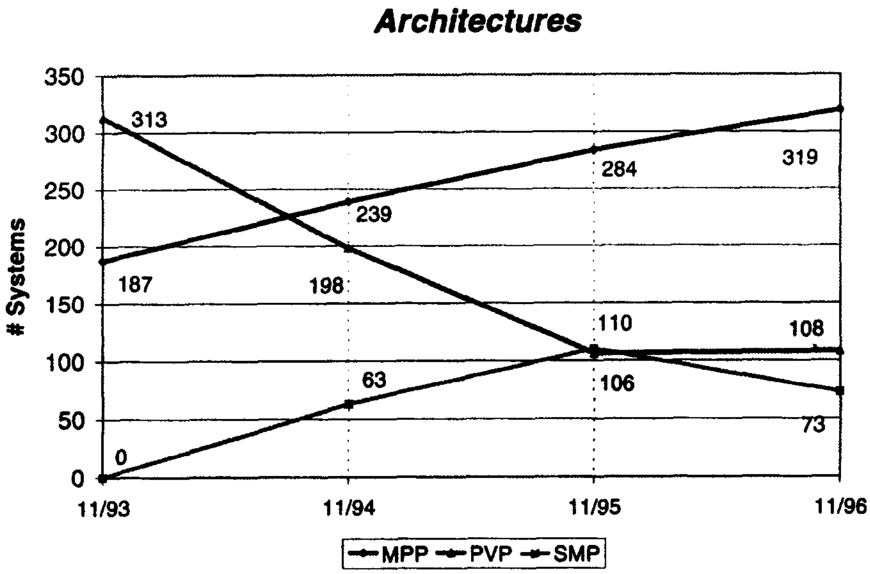

Fig. 6. The evolution of the architectures as it can be seen in the Top500.

and 1995 was the dominating effect with respect to computer architecture. In 1996 SMPs are already on their way out of the Top500 again while the number of MPP systems is still raising. This is reflected in the product announcement of single companies like SGI. They introduced the Origin 2000 series (6 system on the list) which is an MPP system as "follow-up" to their very successful SMP series PowerChallenge. The share of parallel vector processors (PVP) remained stable at a level slightly above 20%. MPP systems are the clearly dominating class of systems in the Top500 with 2/3 of all systems belonging to this class.

In our first report [2] Japan was very much behind with the number of installed MPP systems in 1993. This began to change in 1994 [3]. The number of installed MPP systems in Japan is with 48% now only a little behind the worldwide average of 53%. But like last year almost no SMP systems have been installed in Japan again.

Looking at the average performance of a system in the different classes for the different regions we see in Table 6 that the MPP systems installed in Japan are quite powerful. The average size of systems in Japan measured in GFlop/s is three times as high as in the US or in Europe. Most of the large MPP systems in Japan are produced by Japanese manufacturers. Their architectures are mainly based on distributed memory systems with nodes with vector capabilities. This class of scalable parallel vector processors implemented in CMOS (Fujitsu VPP700/VPP300, NEC SX-4, Hitachi SR2201) does not play an important part outside of Japan yet, but is already entering the European market.

Table 6.
Average system size for the different classes of systems

R_{max} (in GFlop/s)	MPP	PVP	SMP	ALL
USA	16.3	10.3	5.5	13.3
Japan	38.6	22.6	6.0	31.3
Europe	14.3	14.6	5.7	13.1
other	10.2	13.7	5.5	8.8
ALL	18.8	14.6	5.6	16.0

6. Technological changes

Let us now try to analyze the technology used for the processors. With respect to the chip technology we find that the number of systems based on ECL

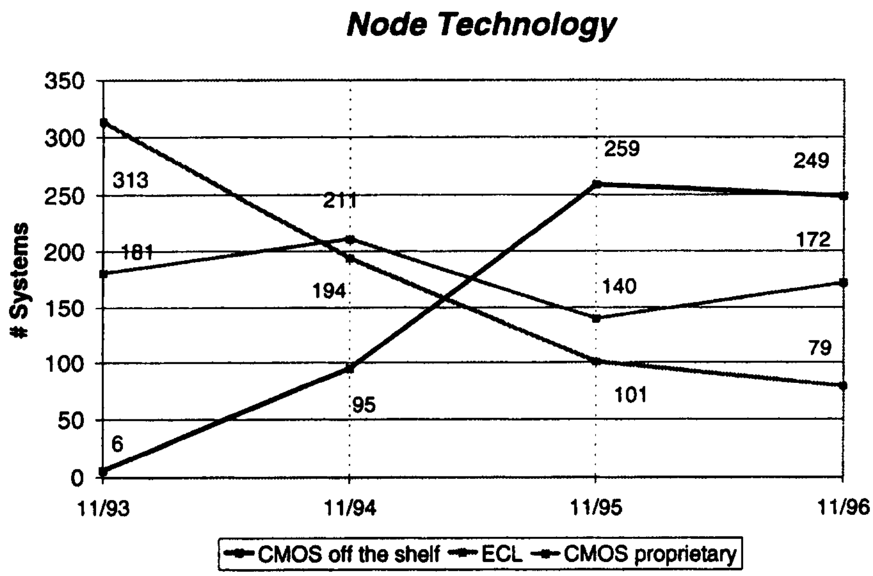

Fig. 7 The usage of different node technologies as can be seen in the Top500. We count for this figure the following systems as CMOS off-the-shelf: Convex SPP, IBM SP1/2, SGI.

chip technology is steadily decreasing from 332 in mid 1993 to 79 by the end of 1996. During the same time the number of systems using proprietary processors with custom chips decreased from 59 to 35 in 1995 and rose again to 60 in November 1996. This increase is due to the vector processors built with CMOS technology such as the SGI/Cray J90, NEC SX-4 and Fujitsu VPP700/VPP300. 342 of the systems in the current list are built by using "off-the-shelf" processors.

In Fig. 7 we see that the number of systems with nodes binary-compatible to workstation systems has increased termendously during 1994 and 1995 and is now stable at 50%. This class of systems includes the ones from Silicon Graphics, the Convex SPP and the IBM SP2. The very strong increase of systems with such a node design indicates a very strong trend in the field of high performance computing. This trend is supported by the advantage of using standard workstation nodes keeping the design costs low. Also all available software for the workstations can immediately be used on the parallel systems, at least on a single processor. This seems to be a big advantage for selling systems to industrial users as can be seen in [4].

7. Application areas

Looking at the different application areas in Figs. 8 & 9 we see an increasing share of industrial installations for 1996 with finally 30% of installed systems and 14.5% of the installed performance after the decreasing share of industrial installations during the past years. If you look at the Top500 in more detail you see that only IBM with 53%, SGI with 38% and HP/Convex with 32% have an over proportional share of industrial installations. This is a very strong indication of the advantage binary compatible nodes might have on the HPC market.

8. List of the TOP25 sites

The Top500 list of supercomputer sites is given in Table 7. This list has been established by simply adding the LINPACK R_{max} performance in GFlop/s of all supercomputers installed at a given site. Generally under a "site" we have combined supercomputers, which are installed in the same geographical location, and belong to the same organizational unit. Thus all machines belonging to a university on

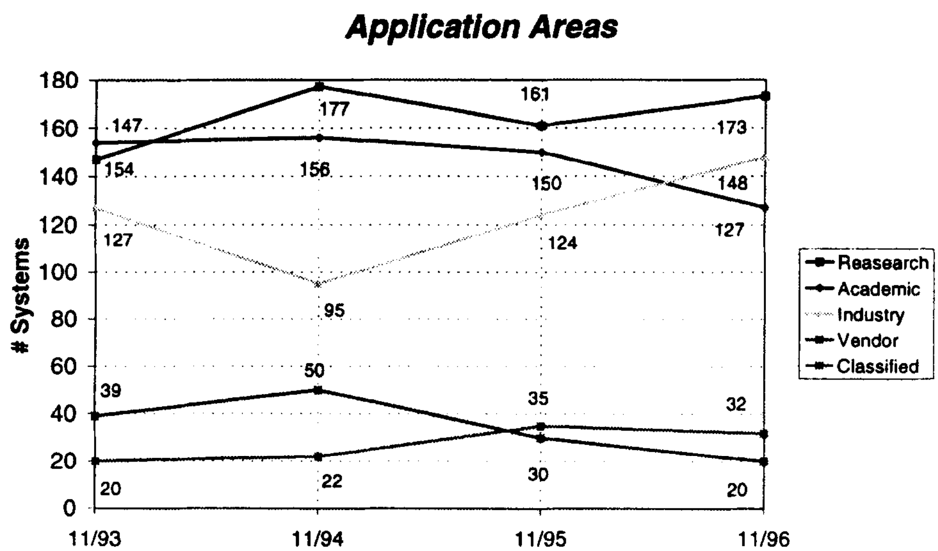

Fig. 8. The distribution of systems on the different application areas over time.

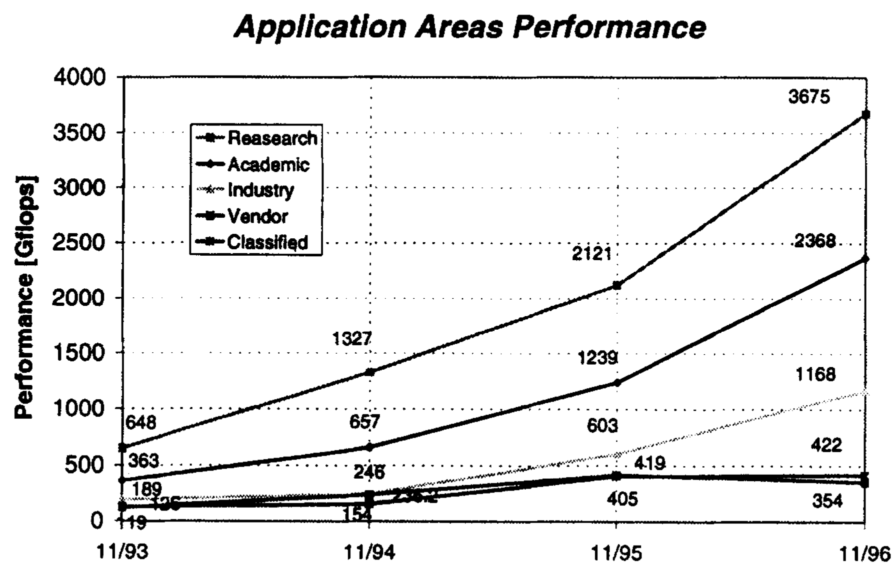

Fig. 9. The distribution of performance on the different applications areas over time.

the same campus were added, even though they might be in different departments. The previous ranking from November 1995 is given in the second column (see [7]).

The list does not contain any of the vendor machines. Most of the supercomputer vendors have substantial compute capabilities, which would make the TOP25 centers list. However, the intent of this list is to give an indication where most compute power in terms of scientific and research applications is concentrated. Therefore we decided to list the vendors separately in Table 9.

In all tables the column "machines" lists the machines whose performance have been added to

Table 7.
Top25 supercomputer sites

	1995	Institution	Machines	Performance
1	24	Tsukuba University	1, 40	408.0
2	9	Tokyo University	3, 27, 58, 180, 485	315.4
3	1	National Aerospace Lab. (NAL), Tokyo	2, 191	239.7
4	11	Japan Atomic Energy Research	6, 22, 46, 118, 365, 391	217.8
5	3	National Security Agency	7, 53, 107, 127, 249, 291, 322, 435, 442, 446	195.5
6	4	Los Alamos National Laboratory	21, 28, 66, 196, 197, 370, 393	166.9
7	13	Pittsburgh Supercomputing Center	13, 31, 143	157.7
8	2	Oak Ridge National Laboratory	5, 104, 178	154.3
9	5	Sandia National Labs, Albuquerque	2	143.4
10	6	University of Minnesota	25, 29, 284, 362, 400, 443	126.4
11		Osaka University	18, 19, 475	126.2
12	19	Lawrence Livermore National Laboratory	32, 38, 65, 438, 441	123.7
13		ECWMF, Reading, UK	10, 128, 164	120.8
14		University of Stuttgart, Germany	20, 36, 334	117.3
15		CNRS/IDRIS, France	11, 160, 278	112.9
16		DOD/CEWES, Vicksburg	12, 124	106.9
17	7	Nat. Lab. High Energy Physics, Japan	8	98.9
18		Kyushu University	9	94.3
19	8	Cornell Theory Center	12	88.4
20	12	Tohuku University	74, 79, 145, 169, 209, 425	85.8
21	22	NCSA, University of Illinois	52, 95, 211, 221, 244, 264, 292	85.0
22		NERSC, Lawrence Berkeley Nat. Lab.	34, 140, 373, 374, 375	81.5
23	10	Maui HPCC	16, 175	78.8
24	15	Atmospheric Env. Serv., Dorval, Canada	48, 73, 82	73.9
25	23	Caltech/JPL	61, 103, 120, 251, 261	69.6
Total			95 Systems	3589.1
Percentage			19.0%	44.9%

reach the total performance for a site. The integers refer to the ranking of these supercomputers on the Top500 list. The performance column lists the aggregate performance of all the machines at the site in LINPACK R_{max} (GFlop/s). An overview of many of the supercomputers in use is given in [8].

There are several intriguing observations one can make from Table 7. In order to qualify as a top supercomputer site, an installation must have at least a machine with about 70 Gflop/s performance. This is almost twice the cut-off one year ago, which was about 35 Gflop/s. Three years ago the cut-off was only 13.7 Gflop/s, and 70 Gflop/s would have placed an institution on rank two. There has been a tremendous acceleration of available cycles at the top supercomputer centers. In 1996 again the number of machines at Top25 sites and their share of the total performance in Gflop/s increased slightly.

Table 8.
Geographical distribution

Region	1995	1996
USA/Canada	15	14
Japan	10	8
Europe	0	3

Another significant change is in the geographical distribution. In 1996 the most important change was that there were three European centers which entered the Top25 list. In 1995 there were no European sites among the Top25. Table 8 shows the change in the geographical distribution of the centers.

Table 9.
Top Vendor sites

	Institution	No. of machines	Performance
1	Cray Res.	8	96.1
2	IBM	1	88.4
3	NEC	1	60.7
4	Hitachi	2	42.6
5	Fujitsu	3	27.2
6	SGI	2	22.4
7	HP/Convex	2	18.3
8	Digital	1	6.7
Total			362.4

The list also shows how much US government spending dominates the supercomputing world. All 13 US sites directly or indirectly are funded by the US government. There are nine US government laboratories/centers (five Departments of Energy, one classified, one NASA, two Departments of Defense), and the five US universities receive their support for supercomputers from the NSF or DoD (Minnesota). However, the foreign sites are also all falling into the same category, and are government institutions in their countries.

8.1. Vendor sites

Most of the supercomputer vendors maintain substantial benchmarking capabilities. These are usually distributed worldwide. Since the vendor centers are geared towards benchmarking and internal software development, we in 1995 decided not to list them in the same list as the Top25 supercomputers centers, which are geared towards research. In Table 9, we list the all vendor sites. Only the first two, Cray and IBM, would have made it to the Top25 list. However, we believe that the vendors no longer report benchmarking machines for the Top500 since there is a limit to the number of vendor machines which can be reported.

9. Conclusions

From the present eight releases of the Top500 we see:

- For positions in the range of 100–500 the performance of the individual systems is increasing by a factor of 2 every year while the total installed performance is increasing by a factor of 1.8 every year.
- The new number one for both releases of the Top500 in 1996 have been Japanese systems and not the announced systems from US manufacturers.
- Ths US is the clear world leader both as producer and as consumer of high performance computers. This leadership has been even more strengthened in 1996.
- The US and Japanese vendors are dominating their home markets, while European manufacturers are playing no role at all, not even in Europe.
- The shake out of the HPC manufacturers culminated in SGI buying Cray Research.
- SGI/Cry and IBM are leading the list with respect to the number of installed systems and with respect to installed performance.
- Microprocessor based supercomputers have brought a major change in the accessibility and affordability of supercomputers. The installed base of supercomputer GFlop/s almost tripled in the last 18 months worldwide of 1995 in the US.
- MPP systems are the dominating architecture, while the number of SMP systems started to go down in the Top500.
- The number of ECL based systems is strongly decreasing all the time, and by the end of 1995 about 84% of the systems in the Top500 were built with CMOS technology.
- In the Top500 a strong trend to nodes being binary-compatible to major workstation families can be seen since 1995.
- Vendors using such "off-the-shelf" nodes (IBM, SGI and Convex) are in the position to sell over proportionally many systems to industrial customers.
- IBM is leader in the industrial market place with 67 systems installed even ahead of the team SGI/Cray with 58 systems.
- The USA is the clear leader in the industrial usage of HPC technology.

With the Top500 project going into its fifth year, many trends and evolutions of the HPC market could be made quite transparent. This has proven the Top500 to be a very valuable tool. Some of the trends mentioned can surely be stated and anticipated without the Top500 while many others are certainly surprising and could not be visualized without it. Future releases of the Top500 list should enable the HPC community to track important developments much more accurately than in the past.

References

[1] J.J. Dongarra, H.W. Meuer and E. Strohmaier, Top500 supercomputer sites, Technical Report 33, 34, 38, 40, 41, 42, University of Mannheim, Germany, June 1993, November 1993, June 1994, November 1994, June 1995, November 1995.

[2] J.J. Dongarra, H.W. Meuer and E. Strohmaier, eds. *TOP500 Report 1993*, University of Mannheim, 1994.

[3] J.J. Dongarra, H.W. Meuer and E. Strohmaier, eds., *TOP500 Report 1994, Supercomputer 60/61* 11(2/3) (1995).

[4] J.J. Dongarra, H.W. Meuer and E. Strohmaier, eds., *TOP500 Report 1995, Supercomputer 63* 12(1) (1996).

[5] J.J. Dongarra, H.W. Meuer and E. Strohmaier, eds., *TOP500 Report 1996, Supercomputer* 13 (1997), to appear.

[6] Hans-Martin Wacker, The European HPC initiative, in: *Supercomputer'92*, H.-W. Meuer, (Springer, Berlin, 1992) 239.

[7] H.D. Simon, The Top25 supercomputer sites, *Supercomputer 63* 12(1) (1996) 38–47.

[8] K.E. Gates and W.P. Petersen, A technical description of some parallel computers, Technical Report IPS 93-15, IPS, ETH Zuerich, Switzerland, IPS Zentrum, CH 8092 Zuerich, Switzerland, November 1993, (to appear in *Int. J. High Speed Comp.*, 1994).

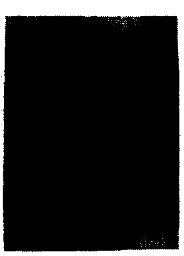

Hans-Werner Meuer is the Director of the Computing Center and Professor of Computer Science at the University of Mannheim since 1974. He received a doctorate in Mathematics from RWTH, Aachen. From 1962 until 1973 he was scientist at the Research Center Juelich, his final position was acting director of the Central Institute of Applied Mathematics.

For over 35 years he has been involved in data processing, starting with a ZUSE Z23 at Giessen University. Since the beginning of the eighties he works intensively in the supercomputing field, he is the founder and organizer of the well known Mannheim Supercomputer Seminar, taking place yearly since 1986. In 1993 he started the Top500 initiative together with Jack Dongarra and Erich Strohmaier. Under his direction many projects in the field of Computer Algebra, Computer Networking, Parallel Computing and Internet/WWW have been successfully completed.

He has published numerous articles in all these fields since 1962 till now. Since more than ten years he is editor-in-chief of the journal PIK- Praxis der Informationsverarbeitung und Kommunikation (KG Saur Publisher, Munich). He works actively in a variety of computer and computer science organizations, including ACM, GI, IEEE, ZKI and SAVE.

Jack Dongarra holds a joint appointment as Distinguished Professor of Computer Science in the Computer Science Department at the University of Tennessee (UT) and as Distinguished Scientist in the Mathematical Sciences Section at Oak Ridge National Laboratory (ORNL) under the UT/ORNL Science Alliance Program. He specializes in numerical algorithms in linear algebra, parallel computing, use of advanced-computer architectures, programming methodology, and tools for parallel computers. Other current research involves the development, testing and documentation of high quality mathematical software. He was involved in the design and implementation of the software packages EISPACK, LINPACK, the BLAS, LAPACK, ScaLAPACK, Netlib/XNetlib, PVM/HeNCE, MPI and the National High Performance Software Exchange; and is currently involved in the design of algorithms and techniques for high performance computer architectures.

Professional activities include membership in
the Society for Industrial and Applied Mathematics (SIAM),
the Institute of Electrical and Electronics Engineers (IEEE),
the Association for Computing Machinery (ACM), and
a Fellow of the American Association for the Advancement of Science (AAAS).

He has published numerous articles, papers, reports and technical memoranda, and has given many presentations on his research interests.

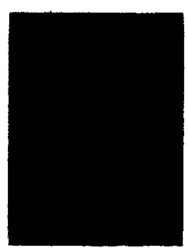

Horst D. Simon is Director of the NERSC (National Energy Research Scientific Computing) Division located at Lawrence Berkeley National Laboratory in Berkeley, CA since February 1996. ERSC is the principal supplier of production high- performance computing services to the nationwide energy research community.

Previously, from 1994 to 1996 Simon was with the Advanced System Division of Silicon Graphics in Mountain View, California, where he managed SGI's university and research laboratory programs. From 1987 to 1994 he was with Computer Sciences Corporation at the NAS Division at NASA Ames Research Center, Moffett Field, California, leading a research department with groups in parallel applications, scientific visualization, and numerical grid generation.

Dr. Simon's research interests are in the development of high performance algorithms for vector and parallel machines. Particular areas of interest are sparse matrix algorithms, algorithms for large scale eigenvalue problems, and domain decomposition algorithms for unstructured domains for parallel processing. Dr. Simon's algorithm research efforts were honored with the 1988 Gordon Bell Award for parallel processing research. He participated significantly in the development of the NAS Parallel Benchmarks.

Simon previously taught at SUNY Stony Brook and worked in a research group at Boeing Computer Services. He holds a Diploma in Mathematik from the TU Berlin, Germany (1978) and a Ph.D. in mathematics from the University of California (1982), Berkeley, CA.

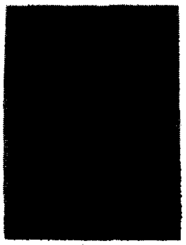

Erich Strohmaier is postdoctoral research associate in the Computer Science Department of the University of Tennessee in Knoxville, where he is leading the Top500 project for the analysis of the HPC market. He is also coordinating the Park Bench project for performance evaluation of parallel systems. His research interests are in the field of high performance computer architectures, parallel algorithms, performance evaluation, performance analysis and benchmarking of parallel computer systems with shared and distributed memory systems.

Previously, from 1990 to 1995 Erich Strohmaier was with the Computing Center of the University of Mannheim, where he was leading the parallel computing group. He was responsible for support and operation of the parallel systems of the University.

During his studies at the University of Heidelberg, Germany he specialized in numerical methods for Theoretical Physics. He holds a Diploma (1987) and a Doctorate (1990) in Theoretical Physics from the University of Heidelberg, Germany.

Professional activities include membership in the Association for Computing Machinery (ACM), the Gesellschaft fuer Informatik (GI), and the Deutsche Physikalische Gesellschaft (DPG).