High Performance Computing in the U.S. in 1995 – An Analysis on the Basis of the TOP500 List

Jack J. Dongarra
Computer Science Department
University of Tennessee
Knoxville, TN 37996-1301
and
Mathematical Science Section
Oak Ridge National Laboratory
Oak Ridge, TN 37831-6367
dongarra@cs.utk.edu

and

Horst D. Simon
Advanced Systems Division
Silicon Graphics Computer Systems
Mail Stop 580
Mountain View, CA 94043-1389

November 17, 1995

Abstract

In 1993 for the first time a list of the top 500 supercomputer sites worldwide has been 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) in the U.S. This note summarizes some of the most important observations about HPC in the U.S. as of late 1995, in particular the continued dominance of the world market in HPC by the U.S, the market penetration by commodity microprocessor based systems, and the growing industrial use of supercomputers.

1 Introduction

The rapid transformation of the high performance computing market in the U.S. which began in 1994 continued at an accelerated rate in 1995. With the introduction of powerful new CMOS microprocessors such as the IBM Power2, the MIPS R8000, and the DEC Alpha processors, supercomputing increasingly has moved to CMOS microprocessor based systems. This trend continued in 1995.

1995 was a remarkable year for the TOP500 in two respects: on one hand almost half (246) of all machines on the TOP500 list were installed in 1995, so 1995 was a year of rapid change and turnover of computer equipment. On the other hand, there was a remarkable absence of new product introductions. The new Cray T90, which would have been big supercomputer news a few years ago, made barely an impression on the TOP500 list (only 8 systems installed in 1995). The leading architectures (IBM SP-2, SGI Power Challenge, and Cray C90) are all pre 1995, and had already significant numbers of installations in past editions of the TOP500 list.

Similarly, there is a lack of "massively" parallel machines. The total number of machines with more than 1024 processors on the list is now eleven. All of these eleven are located in the U.S. Eight of these eleven machines are computers made by TMC (one CM-5 and 7 CM-200s), and probably won't be replaced by similar machines, since TMC is out of the hardware business. There are another 10 machines with processors counts between 512 and 1023 processors, and again all of these are in the U.S. as well.

Thus, the HPC situation in the U.S. can be summarized as follows:

- a few massively parallel supercomputers, mainly in research institutes, based on message passing, using 1992 - 1994 technology (Intel Paragon, TMC CM-5, Cray T3D, IBM SP-2);
- a fair, but shrinking number of parallel vector machines, mostly C90's, about evenly spread in universities, research institutes, and industry;
- a large number of moderately parallel (less than 128 processors) supercomputers based on fast commodity microprocessors (IBM SP-2, SGI Power Challenge, Convex SPP);

The world wide market for HPC in 1993 was estimated to be about $2.4 Billion, with overall growth of the market by very modest aggregate rate of only 1.4% in five years until 1998. These projections imply a very fierce competition also in the future, since the number of vendors in this relatively small market continues to be too large. This can seen from the list of vendors in Table 1, which is updated from Smaby [11]. The consequences of such a large number of vendors competing for such a small (but highly visible and important) market are widely discussed [10]. Compared to a year ago this table has three companies

Table 1: **Commerical HPC Vendors in the U.S (late 1995)**

Status	Vendors
Out of business	Alliant, American Supercomputer, Ametek, Culler, Cydrome, Cray Computer Corp., Denelcor, Elexsi, Kendall Square, Multiflow, Myrias, Prevec, Prisma, Saxpy, SCS, SSI(2), Trilogy, Wavetracer
Division closed	Astronautics, BBN, CDC/ETA Systems, E&S, FPS Goodyear, Gould, Loral, Vitesse
Merged	Celerity, FPS, Key, Supertek, Ardent/Steller
Down, not out	AMT(Cambridge), CHoPP, Encore, Stardent/Kubota, Thinking Machines
Currently active	Convex/HP, Cray Research, Fujitsu, IBM, Intel, nCUBE, Meiko, NEC, Parsytec, SGI, Tera

less in the "Currently active" category, and no serious newcomers. The other major change was the acquisition of Convex by HP.

At the same time the federal High Performance Computing and Communications Program (HPCCP) is winding down. After considerable progress has been made as documented in the famous "Blue Book" [4] the focus of federal programs has shifted more towards the NII (National Information Infrastructure). The discussion about HPC in the commercial and in the government market place continues to be based on beliefs and impressions, and often lacks hard data. Claims in the early years of the HPCC that a Teraflop/s performance on significant applications will come to pass by 1996, is almost certainly not going to happen. However, this was the wrong metric to pursue from the very beginning. It continues to surprise that a field such as HPC that is deemed so critically important to the national agenda lacks almost completely any quantitative assessment of its progress.

There is a possibility that the Tflop/s level will be reached on the Linpack benchmark in 1996. A cooperative agreement between Sandia National Laboratories and Intel will result in the installation of machine with more than 9000 processors in 1996.

Before investigating some of the data in [7] in more detail, it is important to understand the limitations of the TOP500 study. These limitations can be summarized in the past. In spite of these inherent limitations, the TOP500 report can provide extremely useful information, and valuable insights. It is more accurate than many marketing studies, and the possible sources of error discussed above are probably statistically insignificant, if we consider only summary statistics, and not individual data. All Mflop/s or Gflop/s performance figures here refer to performance in terms of Linpack $Rmax$.

In the analysis of geographical distribution, machines in Canada have been included in the figures for the U.S., and the figures for Europe include all European countries, not just EC members. The other country category includes mostly countries of the Pacific Rim with the exclusion of Japan, and a few Latin American Countries.

2 U.S. Dominance of the World Wide HPC Market

The TOP500 continues to demonstrate the dominant position the U.S. assumes in the world both as producer and as consumer of high performance computers. In Table 2 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 2 the country of origin then it is striking that 423 out of the TOP500 systems are produced in the U.S., which amounts to 85% of all installed systems. Japan accounts for 12% of the systems, and Europe produces only 3%. The extent of the American dominance of the market is quite surprising, and has been even increasing from the previous report, when the U.S. share was 84%. For years, in particular in the mid 80's, 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 80's, or more recently the HPCC Program have contributed to the current vast superiority of the U.S. high performance computing industry, remains to be investigated. It is interesting to note that one view expressed outside the U.S. [13] is that strengthening the U.S. HPC industry and easing the transition to MPP was the *only* rationale for the HPCC Program.

The numbers for Europe are actually better than last year (15 machines in Nov. 95 versus 12 machines in June '95). 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 interaction of users with vendors as is prevalent in the U.S., 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 3 is analogous to Table 2, but instead of the number of systems, the aggregate performance in $Rmax$-Gflop/s is listed. Table 3 demonstrates a truly astounding event in 1995: within six months the total number of installed

Table 2: **US Share of Total Number of Installed TOP500 Systems**

Systems Manufactured In	Systems Installed In U.S.	Japan	Europe	Other	Total
U.S.	262	27	116	18	423
Japan	5	45	11	1	62
Europe	2	1	12	0	15
Total	269	73	139	19	500

Table 3: **US Share of Total Rmax (in Gflop/s) of Installed TOP500 Systems.**

Systems Manufactured In	Systems Installed In U.S.	Japan	Europe	Other	Total
U.S.	2581	198	675	88	3542
Japan	69	1030	72	3	1174
Europe	10	5	55	0	70
Total	2660	1234	801	91	4786

Gflop/s in the U.S. increased from 1392 Gflop/s in June to 2660 Gflop/s in November 1995. This is an increase of 92% in only six months. At the same time growth in other regions was substantial but not quite as high: from 709 Gflop/s to 1234 Gflop/s or 74% in Japan, and from 457 Gflop/s to 801 Gflop/s or 75% in Europe. 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 U.S. has succeeded in the sense that the infrastructure for HPC is dramatically changing. A large number of institutions now has access to Gflop/s level computing for machines which cost not much more than $ 1M. 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 U.S. 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 GDP in the U.S.

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 4 we list the a measure of the supercomputer density by ranking the top ten countries with the highest number of supercomputer per

Table 4: **Population (in thousands) per TOP500 supercomputer.**

Country	Population (in thousands)	Number of TOP500 entries	Population (in thousands) per Supercomputer
Switzerland	6,813	9	757
Singapore	2,769	3	923
USA	255,200	261	978
Denmark	5,158	4	1290
Norway	4,288	3	1429
Finland	5,008	3	1669
Germany	80,250	48	1672
Netherlands	15,160	9	1684
Japan	124,500	73	1705
Hong Kong	5,800	3	1933
Sweden	8,652	4	2163
France	57,180	25	2287
Austria	7,776	3	2592
UK	57,700	17	3394
Canada	27,370	8	3421

capita. Population date are from the "Interactive 3D Atlas" and date from 1992.

Table 4 shows that on an international comparison most industrialized 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 industrialized countries. It should be mentioned that the 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.

3 Market Penetration by Technology and Architecture

The penetration of the supercomputer market by microprocessor based supercomputers and the increased use of SMPs and arrays of SMP is another often debated trend. The trend towards commodity CMOS is now firmly established. In Table 5 we present the number of installations for the different machine architectures used among the TOP500, both world-wide and in the US/Canada.

In 1994 MPPs moved ahead of PVP and are now clearly the largest architectural category, both in the US and worldwide. SGI made a very strong showing

Table 5: **Architecture (in number of installations).**

Region	MPP	SMP	PVP
U.S./Canada	155	74	40
share	58%	27%	15%
Worldwide	284	110	106
share	57%	22%	21%

in the SMP category with installing 54 new machines in 1994 alone. In 1995 SMP systems for the first time surpassed also PVP worldwide. This trend is even more clearly visible in the US. The share of SMP systems in the US is 5% higher than world wide. With other vendors considering or announcing to enter the SMP market (DEC and HP), this appears to be the next significant trend. Over the next year we can expect CMOS based SMP systems to replace a significant number of older PVP systems.

4 Conclusions

The analysis of the data provided by the TOP500 report has led us to a number of conclusions concerning the state of HPC in the U.S. at the end of 1995. Some of these conclusions are:

- The U.S. is the clear world leader both as producer and as consumer of high performance computers. This leadership position has been even more strengthened in 1995.
- Microprocessor based supercomputers are about to bring a major change in the accessibility and affordability of supercomputers. The installed base of supercomputer Gflop/s almost doubled in the last six months of 1995 in the US. This increase is due to a large number of medium to small installations of machines based on the IBM Power2 and the SGI/MIPS R8000 processors.
- MPPs now account for more that half of all installed supercomputers worldwide and in the US. Market penetration by MPPs worldwide is now at the same level as the US. SMP system are used more frequently in the US than worldwide. This may be an indication of a trend towards replacing older PVP systems with SMP systems.

Generally the TOP500 list has proven itself to be an extremely valuable tool for evaluating trends in the HPC market. Future releases of this report should enable the HPC community to track important developments much more accurately than in the past.

References

[1] J. Almond, M. Heib, S. Jarre, E. Krause, R. Rühle, H. Simon, and H.- M. Wacker. Erleben die Vectorrechner eine Renaisssance? pages 218 - 238, Supercomputer '94, Editor H.- W. Meuer, K. G. Saur Verlag, Munich, New Providence, London, Paris, 1994.

[2] D. H. Bailey, E. Barszcz, L. Dagum, and H. D. Simon. The NAS Parallel Benchmarks: Review and Current Results pages 164 - 182, Supercomputer '94, Editor H.- W. Meuer, K. G. Saur Verlag, Munich, New Providence, London, Paris, 1994.

[3] D. H. Bailey, J. Barton, T. Lasinski, and H. Simon (editors). The NAS parallel benchmarks. Technical Report 103863, NASA Ames Research Center, Moffett Field, CA 94035, July 1993.

[4] Committee on Physical, Mathematical, and Engineering Sciences. *High Performance Computing and Communications: Toward a National Information Infrastructure*. Office of Science And Technology Policy, Washington, D.C., 1994. (available via Mosaic from the National Coordination Office for HPCC, URL `http://www.hpcc.gov/`).

[5] Congress of the United States, CBO. Promoting high performance computing and communications. Technical Report ISBN 0-16-041828-3, U.S. Government Printing Office, Mail Stop SSOP, Washington D.C. 20402-9328, June 1993.

[6] J. J. Dongarra. Performance of various computers using standard linear equations software. Technical Report CS-89-85, Department of Computer Science, Univ. of Tennessee, November 1993.

[7] J. J. Dongarra, H. W. Meuer, and E. Strohmeier. TOP 500 supercomputer sites. Technical Report RUM 40/94, Computing Centre, Univ. of Mannheim, 68131 Mannheim, Germany, November 1994.

[8] General Accounting Office. High performance computing: Advanced Research Projects Agency should do more to foster program goals. Technical Report GAO-IMTEC-93-24, United States GAO, Washington, D.C. 20548, May 1993.

[9] Alan H. Karp, Don Heller, and Horst Simon. 1994 Gordon Bell Prize winners. *IEEE Computer*, 28(1), January 1995.

[10] H. D. Simon. High performance computing: Architecture, software, algorithms. Technical Report RNR-93-018, NASA Ames Research Center, Moffett Field, CA 94035, December 1993.

[11] G. P. Smaby. Crystal balls and rearview mirrors: A wall street view on supercomputing. Parallel Computing in the 90s Seminar presented at NASA Ames Research Center, January 1992.

[12] B. Traverstat and D. McNab. Acceptance test for the Intel Paragon XP/S15. Technical Report RND-94-004, NASA Ames Research Center, Moffett Field, CA 94035, February 1994.

[13] Hans-Martin Wacker. The European HPC initiative. In H.-W. Meuer, editor, *Supercomputer '92*, pages 239 – 239, Berlin, Heidelberg, New York, 1992. Springer Verlag.