

High Performance Computing in Europe and USA: A Comparison

Erich Strohmaier

NERSC, Lawrence Berkeley National Laboratory, USA

Hans W. Meuer

University of Mannheim, Germany

erich@top500.org

meuer@rz.uni-mannheim.de

1 Cyclotron Road, Berkeley, CA 94720, USA

L 15,16; 68131 Mannheim

www.top500.org

Keywords: High-Performance Computing market, Moore's Law, performance evaluation, Supercomputing

Abstract: Since 1993 we compile and publish twice a year a list of the most powerful supercomputers in the world. In this article we compare the situation of High-Performance Computing (HPC) in Europe and the USA based on the quantitative data gathered over the years in this TOP500 project (see www.top500.org for a complete access to all data). We start with an overview of the project, including the motivation and rationale behind it. We present the results of the TOP500 project as of November 2000 focusing on the changes over time with respect to the manufacturers active in this market-segment, the architectures, and technologies used for these systems, the applications HPC systems are used for. In our analysis we focus especially on the situation in Europe and the USA. As conclusion we will present plans for the future of the TOP500 project.

Introduction

In the middle of the eighties one author of this paper¹ started to collect and publish statistics about the supercomputer market. At that time it was rather simple to decide which system qualified as supercomputer. The performance gap between vector system and regular mainframes was too large to leave any doubts. Beginning of the nineties this situation had gradually changed. A considerable number of companies competed in the HPC market with a large variety of architectures such as vector computer, mini vector computer, SIMD (Single Instruction on Multiple Data) and MPP (Massive Parallel Processing) systems. A new definition was needed to decide which of these systems was a „supercomputer“. This definition needed to be architecture independent. Because of Moore's Law this definition also had to be dynamic in nature to deal with the constant increase in computer performance.

Consequentially in early 1993 we developed the TOP500 idea in Mannheim. The basic idea was to list the 500 most powerful computer systems installed at some place twice a year and to call these systems supercomputer. The problem was to define how powerful a computer system is. For this we decided to use the performance results of the Linpack benchmark from Jack Dongarra [1], as this was the only benchmark for which results were available for nearly all systems of interest.

Since 1993 we publish the TOP500 twice a year using Linpack results². Over the years the TOP500 served well as a tool to track and

analyze technological, architectural, and other changes in the HPC arena. As one example of such analysis we show in Figure 1 the evolution of the total installed performance seen in the TOP500. We plot the performance of the systems at positions 1, 10, 100, and 500 in the list as well as the total accumulated performance of all 500 systems. Fitting an exponential curve to the observed data points we make an extrapolation 5 years into the future. We see that our data validate the exponential growth of Moore's Law very well even though we use Linpack performance numbers and not peak performance values. Based on the extrapolation from these fits we can expect to have the first 100 TFlop/s system by 2005. At that time also no system smaller than 1 TFlop/s should be able to make the TOP500 any more.

Producers and Consumers of HPC Systems

One of the main differences in the computing industry between Europe and the USA is - of course - the strong dominance of US manufacturers on a global scale. European manufacturers of PCs or workstation equipment are traditionally having a hard time to gain market share outside of their home-markets. If we look at the situation at the high performance end of the computing market place we have to recognize that this difference becomes extreme. Despite government subsidized projects such as the almost forgotten Suprenum computer, European supercomputer manufacturers never gained large market shares even in their home-markets. Within the TOP500 project we tracked European companies such as Parsytec or Meiko. Looking at the overall market share of European manufacturers in Figure 2 we have to realize there simply are no longer any European companies active in this market segment. Most of them died and a few managed to survive by finding highly specialized niche markets which they serve with equally highly specialized systems which are unusable for the broader need of the general HPC community.

The total lack of any production capability for such technology in Europe immediately raises the question if there is a similar difference in the usage of such HPC systems. Looking at the „consumption“ of HPC systems as seen by the TOP500 project in Figure 3 we see that

¹ Hans W. Meuer from the University of Mannheim in Germany.

² All data from the TOP500 and further analysis are available from our main web site at www.top500.org.

High Performance Computing in Europe and USA: A Comparison

this is not the case. Europe managed to increase its share of high-end installations slightly over time. However it should be clear that the USA dominate the usage of HPC technology very strongly as well. The slight relative increase in HPC usage in Europe corresponds mainly to a similar decrease of usage in Japan.

Performance Development

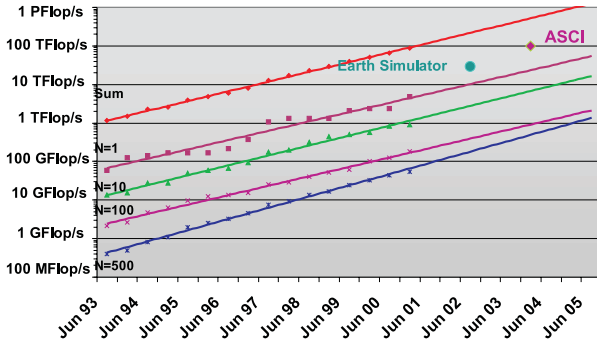


Figure 1: Evolution of system performance as seen in the TOP500.

Producers




Figure 2: Geographic distribution of manufacturers by number of systems in the TOP500.

Continents




Figure 3: Geographic distribution of customers by number of systems in the TOP500.

Companies producing HPC Systems

We now look more closely into which companies actually produce the systems seen in the TOP500. In Figure 4 we see that 8 years ago the specialized HPC companies such as Cray Research, Thinking Machines (TMC), Intel with their hypercube based iPSC systems, and the Japanese vector system manufacturer Fujitsu, NEC, and Hitachi dominated this market. This situation has clearly changed. Nowadays mainstream computer manufacturers from the Workstation and PC segment such as IBM, SGI, Sun, HP, and DEC have largely taken their place.

Manufacturer

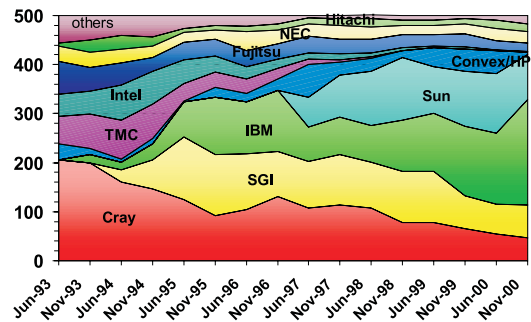


Figure 4: Manufacturers of systems in the TOP500.

Architectures and Technologies

The most noticeable difference to the situation in the overall computer market is, that Intel microprocessor based systems still play only a minor role in the HPC arena. This market is still dominated by super-scalar RISC workstation processors as clearly seen in Figure 5. One reason for the absence of Intel in this market is certainly, that Intel dropped its own ambitions and plans for HPC system in the middle of the nineties. The advent of PC clusters and their slow appearance in the TOP500 might finally increase the number of Intel based supercomputers again.

Chip Technology




Figure 5: Chip technology of systems in the TOP500.

Applications

One of the strong trends seen in the HPC market over the last years was the strong increase in use of HPC systems for new commercial applications (see [2]). Again the USA did clearly lead this change ahead of Europe and Japan. This trend began in 1994 when companies such as SGI, Digital, and Sun started to sell symmetrical multiprocessor (SMP) models of their major workstation families. From the very beginning, these systems were popular with industrial customers because of the maturity of these architectures and their superior price/performance ratio. At the same time, IBM SP2 systems started to appear at a reasonable number of industrial sites. While the SP initially was sold for numerically intensive applications, the system began selling successfully to a larger market, including database applications, in the second half of 1995.

Subsequently, the number of industrial customers listed in the TOP500 increased from 85, or 17%, in June 1995 to about 241, or 48.2%, in June 1999 and is stable since then. This increase shown in Figure 6 appears to be the combination of several factors.

- 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 (Sun, 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.

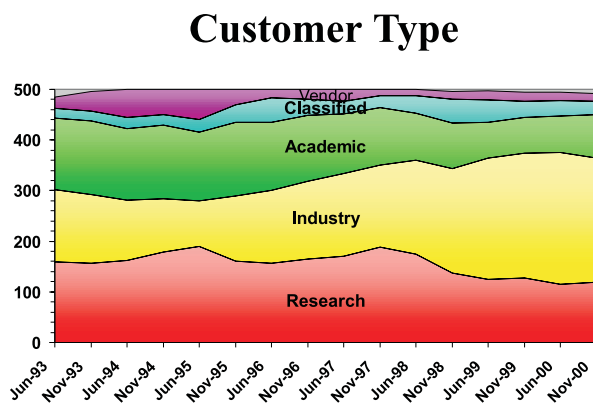


Figure 6: Type of customer seen in the TOP500.

Government Programs and Large HPC Centers

A comparison of the HPC situation between the USA and Europe would not be complete without looking at the potential influence of government programs. Programs such as ASCI (Accelerated Strategic Computing Initiative) certainly attract a lot of public interest. It is

however not clear to which extent these programs are actually capable of influencing the market directly in the short term as they only represent isolated (but large) business opportunities, which are still small compared to the overall market size.

In the long term the USA government programs however do certainly provide an environment for HPC system users and producers to establish, defend and increase their competitive advantage. To demonstrate this we now look over all 16 editions of the TOP500. For all centers we sum up the hypothetical Linpack performance all their systems could have delivered over their lifetime. In Figure 7 we notice that there are 6 centers from the USA, 3 from Europe and 1 from Japan. The first three centers are the ASCI centers, which provided together a total of 21.9 TFY (Tera-Flop-Years). The other 7 centers together provided with 16.7 TFY less than the 3 ASCI centers. In this list the strong influence of government programs on very large centers can clearly be seen.

Two of the three European centers in this list serve only a specialized community, climate research and weather forecasting. This leaves the Forschungszentrum Juelich (FZJ) as only European TOP10 Center serving a broad scientific community. From the European point of view such a situation has to be carefully evaluated and observed. Considering that there are no European HPC manufacturers and no concerted European government initiatives of comparable size to the ASCI program, European scientists might find themselves in a position having access to compute resources which are only a magnitude of order smaller than in the USA.

TOP10 Centers - "integrated"

Rank	Site	$\sum t^*R_{max}$ [TFYears]	Current Best	Country
1	Sandia National Laboratory	8.0	2, 408	USA
2	Lawrence Livermore National Laboratory	7.5	1, 3, 32, 36	USA
3	Los Alamos National Laboratory	6.4	4, 19, 401	USA
4	University of Tokyo	3.3	13, 85	Japan
5	Naval Oceanographic Office (NAVOCEANO)	2.9	5, 21, 326	USA
6	United Kingdom Meteorological Office	2.7	26, 29	UK
7	Forschungszentrum Juelich (FZJ)	2.1	40, 83	Germany
8	NERSC/LBNL	2.0	42, 60	USA
9	Oak Ridge National Laboratory	1.9	17, 70, 402	USA
10	ECMWF	1.8	12, 90, 214	UK

Figure 7: Top10 centers worldwide with respect to Linpack performance delivered over the last 8 years. Purely classified sites and manufacturers are excluded. „Current Best“ refers to the positions of the individual systems of each center in the November 2000 edition of the TOP500.

Conclusions

The HPC market was always dominated by a very rapid change of technologies and architectures. The speed of this change is ultimately coupled to Moore's Law, which states an exponential growth of our computing capabilities by roughly a factor of 2 every 18 months. Tracking the evolution of such a dynamic market place is a challenge and the tools and methods used for this have to be re-evaluated on a constant basis. This is no different for the TOP500 project. In 1993 we decided to switch from our old form of HPC market statistics to the TOP500 in its current form and it has served us well since then.

In the last 8 years the diversity of architectures and applications in the HPC market has increased substantially. Doing justice to this

High Performance Computing in Europe and USA: A Comparison

large variety is certainly not possible with any single benchmark. Nor is it any longer possible to clearly define whether some of the very large but only loosely coupled computing complexes, which we nowadays see in certain application domains, are indeed supercomputers. These systems perform like a "supercomputer" for the type of applications they are designed for but they might miserably fail for other traditional supercomputer applications. Ultimately the increased variety in architectures and applications in HPC needs to be reflected in a more domain specific definition of what constitutes a "Supercomputer" for a certain application domain and measured with domain specific benchmarks. By adapting such a more differentiated approach we hope that we will be able to continue collecting information about the evolution of computer systems on the highest levels of performance.

Bibliography

- 1 Dongarra, J., *Performance of Various Computers Using Standard Linear Equations Software*. 2000, University of Tennessee, Computer Science Tech Report CS-89-85: Knoxville.
- 2 Erich Strohmaier, Jack J. Dongarra, Hans-Werner Meuer, and Simon. Horst D. *The Marketplace of HPC*. Parallel Computing, 25th anniversary edition, North Holland, 25():1517-1544, 1999.