Lawrence Berkeley National Laboratory

Recent Work

Title

Supercomputing: What have we learned from the TOP500 project?

Permalink

<https://escholarship.org/uc/item/3212m88g>

Journal

Computing and Visualization in Science, 6

Authors

Strohmaier, Erich
Meuer, Hans W.

Publication Date

2003-12-12

Supercomputing: What have we learned from the TOP500 Project ?

ERICH STROHMAIER

NERSC, Lawrence Berkeley National Laboratory, USA

1 Cyclotron Road, Berkeley, CA 94720, USA

erich@top500.org

HANS W. MEUER

University of Mannheim, Germany

L 15,16; 68131 Mannheim

meuer@rz.uni-mannheim.de

www.top500.org

Abstract

Since 1993 we compile and publish twice a year a list of the most powerful supercomputers in the world. In this article we analyse some major trends in High-Performance Computing (HPC) 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 2002 focusing on the changes over time with respect to the performance levels, the overall dynamics in this market, the manufacturers active in this market-segment, the architectures, and technologies used for these systems, and the major centers using such systems.

Introduction

Beginning of the nineties we started collecting data and publishing statistics about the supercomputer market. A considerable number of companies competed in the HPC market with a large variety of architectures such as vector computer, mini vector computer, SIMD (Singel Instruction on Multiple Data) and MPP (Massive Parallel Processing) systems. A clear and flexible 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, as this was the only benchmark for which results were available for nearly all systems of interest [1].

Since 1993 we publish the TOP500 twice a year using Linpack results¹. Over the years the TOP500 served well as a tool to track and analyze technological, architectural, and other changes in the HPC arena [2]. Table 1 show the top 10 system as of November 2003. The TOP500 shows the Japanese Earth Simulator System clearly as largest supercomputer worldwide since June 2002.

Table 1								
Rank	Manufacturer	Computer	R_{max} [TF/s]	Installation Site	Country	Year	Area of Installation	# Proc
1	NEC	Earth-Simulator	35.86	Earth Simulator Center	Japan	2002	Research	5120
2	HP	<i>ASCI Q</i> AlphaServer SC	13.88	Los Alamos National Laboratory	USA	2002	Research	8192
3	Self-Made	<i>X</i> Apple G5, Mellanox	10.28	Virginia Tech	USA	2003	Academic	2200
4	Dell	<i>Tungsten</i> PowerEdge, Myrinet	9.82	NCSA	USA	2003	Academic	2500
5	HP	<i>Mpp2</i> , Integrity rx2600 Itanium2, Qadrics	8.63	Pacific Northwest National Laboratory	USA	2003	Research	1936
6	Linux Networx	<i>Lightning</i> , Opteron, Myrinet	8.05	Los Alamos National Laboratory	USA	2003	Research	2816
7	Linux Networx/ Quadrics	MCR Cluster	7.63	Lawrence Livermore National Laboratory	USA	2002	Research	2304
8	IBM	<i>ASCI White</i> SP Power3	7.3	Lawrence Livermore National Laboratory	USA	2000	Research	8192
9	IBM	<i>Seaborg</i> SP Power 3	7.3	NERSC Lawrence Berkeley Nat. Lab.	USA	2002	Research	6656
10	IBM/Quadrics	xSeries Cluster Xeon 2.4 GHz	6.59	Lawrence Livermore National Laboratory	USA	2003	Research	1920

Table 1: Top 10 supercomputer systems as of November 2003.

¹ All data from the TOP500 and further analysis are available from our main web site at www.top500.org.

Performance Growth and Dynamic

One trend of major interest to the HPC community is the growth of the performance levels seen in the TOP500. In Figure 1 the evolution of the total installed performance seen in the TOP500. We plot the performance of the first and last systems at positions 1 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 till the end of the decade. 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 then 1 TFlop/s should be able to make the TOP500 any more. Towards the end of the decade we can expect supercomputer systems to reach the performance level of 1 PetaFlop/s.

Figure 1

Figure 1: Performance growth in the TOP500 and extrapolation till the end of the decade.

The HPC market is by its very nature very dynamic. This is not only reflected by the coming and going of new manufacturers but especially by the need to update and replace systems quite often to keep pace with the general performance increase. This general dynamic of the HPC market is well reflected in the TOP500. We see an average replacement rate of about 160 systems every half year or more than half the list every year. This means that a system which is at position 100 at a given time will fall off the TOP500 within 2-3 years.

Manufacturers

We now look more closely into which companies actually produce the systems seen in the TOP500. In Figure 2 we see that 11 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, and HP have largely taken their place.

Figure 2

Figure 2: Manufacturers of systems in the TOP500.

System Architectures

The changing share of the different system architectures in the HPC market as reflected in the TOP500 is shown in figure 3. Single processor systems and SMPs with shared flat memory are any longer powerful enough to enter the TOP500. For most of the last 11 year MPP systems dominated this market. During the last few years the number of clustered systems grew considerable. Considering the impressive performance dominance of the vector based Earth Simulator System it is an interesting and open question which share of the TOP500 such clusters will be able to capture.

Figure 3

Figure 3: Dominant supercomputer system architectures. Constellations (Const.) are cluster of large SMPs.

Main Supercomputing Sites

Government programs such as ASCI (Accelerated Strategic Computing Initiative) certainly attract a lot of public interest. It is however not clear to which extend 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 analyze this we now look at the combined 11 years history of the TOP500. We define as normalized Linpack performance for a system in a specific TOP500 edition the ratio of its Linpack performance and the sum of the Linpack performances for all the systems on that list. Our so defined normalized performance is therefore not influenced by Moore's Law and can be used for aggregate statistics over all 22 editions of the TOP500 giving equal weight to early lists.

For all centers we sum up the hypothetical normalized Linpack performance all their systems could have delivered over their lifetime. The top 10 entries of this

list of centers assembled in this fashion are shown in Table 2. We see that there are 7 centers from the USA, 3 from Japan and none from Europe. The first three centers are the ASCI centers. The other 7 centers together provided roughly the same number of compute cycles as the 3 ASCI centers. The strong influence of government programs on very large centers can clearly be seen.

Table 2

Site	sum % norm. Rmax	Sum Rmax TF/s	Country
1 Sandia National Laboratories	87.6%	47.60	US
2 Los Alamos National Laboratory	85.2%	90.95	US
3 Lawrence Livermore National Lab.	71.7%	122.64	US
4 NAL	48.8%	16.92	Japan
5 Earth Simulator Center	44.7%	143.44	Japan
6 University of Tokyo	35.1%	22.29	Japan
7 Pittsburgh Supercomputing Center	30.7%	28.30	US
8 NERSC/LBNL	30.4%	35.51	US
9 Oak Ridge National Laboratory	29.4%	23.05	US
10 NASA/Ames Research Center	25.1%	25.11	US

Table 2: Top 10 centers determined by the sum of the per list normalized linpack performance for all systems installed in a center. Classified sites are excluded from this analysis.

The lack of comparable European programs is also clearly reflected by the absence of any European Center in this table. If this situation continues, 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.

Conclusion

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. Tracing 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 11 years the

diversity of architectures and applications in the HPC market has increased substantially. It has to be kept in mind that doing justice to this large variety is certainly not possible with any single benchmark and we are evaluating several approaches to improve this situation. This includes ongoing projects for the creation of new benchmarking metrics such as the DOE SciDAC PERC [3] and the Apex-MAP [4] in the US, and IPACS in Germany [5]

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.
- 3 See <http://perc.nersc.gov> for current details.
- 4 See <http://ftg.lbl.gov> for current details.
- 5 See <http://www.ipacs-benchmark.org/> for current details.

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.