Super Computing: Close and far from us.

Roberto Hoyos
april, 2006

Emerging Information Technologies

Abstract:

In this paper we will present briefly the history of supercomputers, some general characteristics and a handful of the most representative machines ever built. The new trends in supercomputing requires technologies to be more consistent with industrial and economical concerns. Cooling, data transmission, as well as memory tricks, are required to improve their performance. Other areas have opportunities as well: the improvement of the user interfaces, a task that has been long-forgotten or set aside, in the primary concern for efficiency.

1. Introduction

As we stated in the previous article, we will focus here on the supercomputing industry not from a technical approach, rather but analyzing their role on the world through history and through the market. Supercomputing, even when in its technical singularities is far beyond the common man conception, it is well established in the fantasies of millions: the human vs. the machines; problems that otherwise would require days or years for a human to perform, are done in minutes by this computers.

We have centered our arguments in the fact that has been military issues the ones that promoted the development of supercomputing in the first place and of the leadership of the United States in that field, because of its emerging military and scientific resources since the Second World War.

Modern society is constructed over the structure given by computers, so we must look close enough to understand the fundamentals of these machines, and their role in its time. The computer that you may use to read this paper is more powerful than the supercomputers fifteen years ago. ¿Can we predict having the super-technology of today for the personal machines fifteen years from now?

2. History

Computers were built, developed, supported and defined by military initiatives. During World War II the *Colossus* was used for code-breaking [1], allowing the key to the *Lorenz* encoding of the Germans; other tasks demanded enormous calculations as well, for the *ENIAC* was to calculate ballistic trajectories, a task which was performed in 30 seconds compared to the 20 hours that a human needed to compute the same result [2]. By the end of the War and during most of the Cold War Era, computers were used to perform the complex calculations required to build and simulate nuclear bombs, bombers trajectories, and missile scenarios, including other military initiatives, like *DARPANet* and the Space Program.

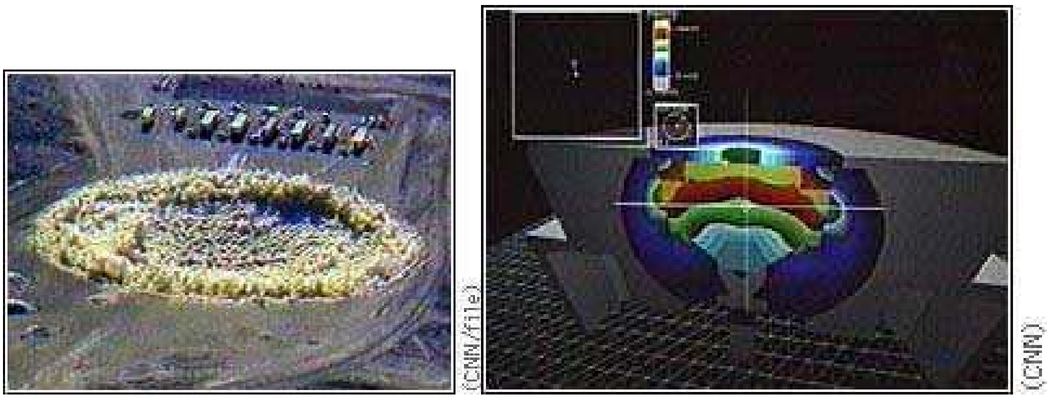

Figs. 1, 2.- Underground testing at Los Alamos vs Computer Simulation. CNN [3]

Apart from the *Colossus*, which was built on England, the United States and the Soviet Union followed parallel paths in developing supercomputers. The Soviet experience is valuable in itself because it was surrounded with the utmost secrecy and it is today that we know little of how advance were Soviet computers [4] (the *M-13* system hold the record in 1984 for being the fastest supercomputer of the day, and even the first computers used in the 1950's to construct the Hydrogen bomb proved to be much more efficient that the American counterpart, with much less processors). Today, the only manufactures are the United States (historically almost the only) and Japan.

The first supercomputers were single-processor architecture, and it was in the sixties that they started to construct machines that used multiple processors, like the *Univac LARC* (1960), which had two processors and could perform 500 kFlops. During the seventies and until the nineties they used vector processing, which required specialized components and were expensive to build, albeit very efficient. Because the end of the Cold War required less investment from the military, supercomputers used instead the multiple processor architecture, which proved to have lower costs –we have to remember how cheap hardware improvements are compared to research on new architectures or software.

3. Programming Languages

Most of the supercomputers today are programmed with *Fortran* (*Formula Translation*), which has a faster compiler than C, and as a specialized programming language it's better suited for scientific tasks (focusing on data rather than instructions, so they perform poorly on other tasks apart from numerical operations). *Fortran* actually is the only language that takes full advantage of the *floating point unit*, specially in matrix and vector calculations [5] (the best mathematical algorithms in the world are written in *Fortran*).

4. Operating Systems

Today, commercial supercomputers projects use Unix variants on its machines. Actually they do use more than one operating system. Every vendor has its own version, generally not known for the general community, because are built to suit specific needs.

Because this operating system focuses on the critical operation demands of supercomputing the graphical interface is left aside, and while modern supercomputers have functional GUI's most of them still wok in the command line mode [6].

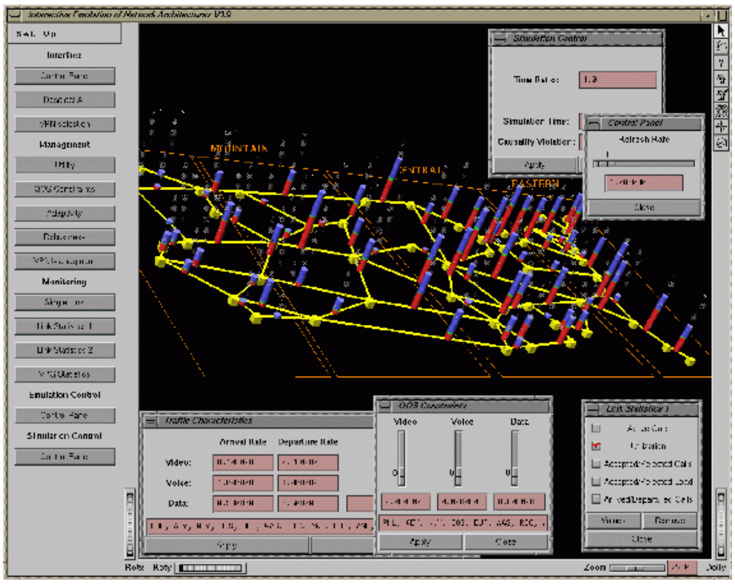

Fig 3. Multimedia network simulation in a IBM SP2. Columbia University [7]

4. Applications

The main source for "who´s who in supercomputing", comes from the semiannually updated list of the 500 more potent supercomputers installed in the world, the *top500.org*. The following charts illustrate some interest features of supercomputing industry.

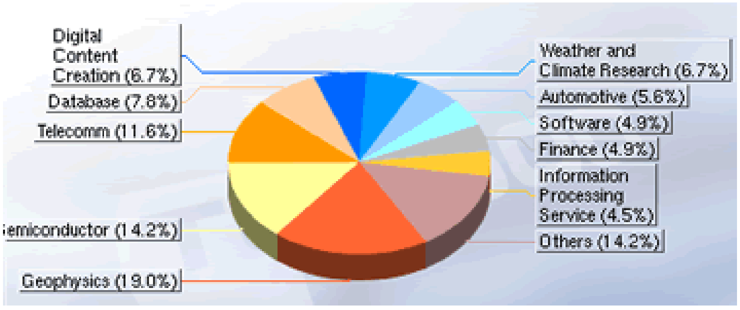

Fig. 4 Supercomputer applications. Top500.org [8]

46.6% of the supercomputer use is invested on scientific tasks, whereas commercial applications has above 50% of use.

This trend is growing though.

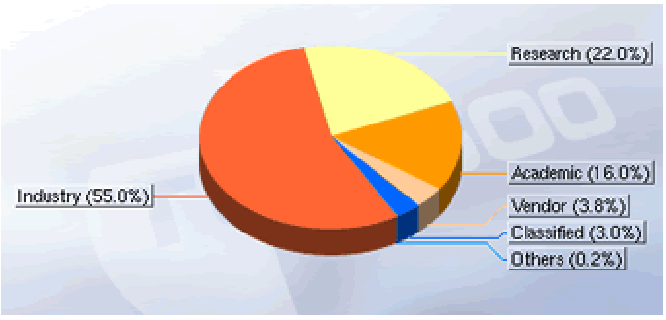

Fig. 5 Costumer Segments. Top500.org

As we see in this chart, industry is the leading costumer segment for supercomputers, followed only by academy and research activities.

We mentioned in the past article how important were computers to tackle the twentieth century challenges for science. The DNA sequence was finally obtained in 2003, and there are undergoing programs devoted to *protein folding*. But these problems *do not* represent even the half of the applications in supercomputing: other fundamental problems as well are in the cutting-edge of this technology: financial modeling, database and server applications, digital content creation, multimedia, and information technologies are problems that require specialized supercomputers.

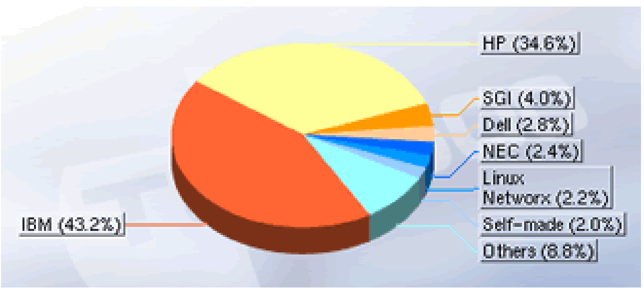

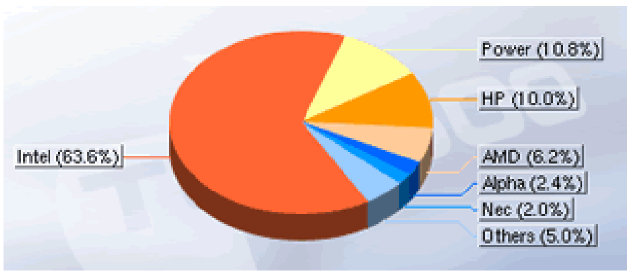

Figs. 6, Manufacturers, and 7, Family Processors. Top500.org

In figures 6 and 7 (above) we see the leading *manufactures* and creators (not sellers) of this technology. Even though Japanese industries are leading in the sales market, the technologies for supercomputing are almost exclusively *U.S. territory*. While it was *Cray Research* who first created supercomputers (today Cray is synonym of supercomputer), IBM and HP bought several companies in the 80's and right now are, together, manufacturing almost 80% of the fastest machines today. Their leadership is unchallenged.

IBM requires special attention: From the list of 500 most powerful computers, IBM provided 43% (216 systems); is in the top 10 with 4 machines; it has created the fastest in Europe (*Mare Nostrum*) and the fastest machine in the world, *Blue Gene*. IBM has been the leading company in information technology for the last 80 years, and historically has dominated the supercomputing world by offering the fastest machines since the 1960's. IBM's role is undeniable, and if this is just not enough, let us remember that it was John Backus (working in IBM) who invented *FORTRAN*, the preferred programming language for supercomputers [9].

In the case of the processors (figure 7) we see that Intel is the biggest supplier of processors, followed by PowerPC and HP as well.

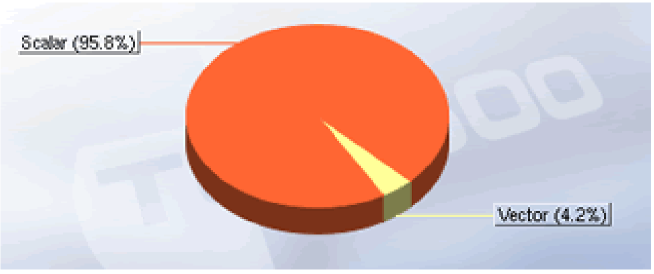

Just to remember how the trend is doing: 95.8% of the fastest systems are in scalar architecture, against a 4.2% of efficient vector systems.

Fig. 8 Architectures. Top500.org

Now let us look briefly to the five fastest computers in the world, and then we will study other group of special purpose machines, like Deep Blue and Deep Crack.

5. Top 5 Supercomputers and other Special Purpose Machines

Thunder, ranked 5.- Uses 4096 processors and reach 19940 GFlops. Processor speed is 1.4 Ghertz. It was manufactured in 2004 by the California Digital Corporation, Quadrix and Intel, and it's right now on *Lawrence Livermore National Laboratory*, California (the Department of Energy's Laboratory). "The Fastest Linux SuperComputer ever built", according to its manufacturers. Its applications: materials science, structural mechanics, electromagnetics, atmospheric science, seismology, biology [10].

Mare Nostrum, ranked 4.- Using 3564 Power PC processors, this machine, located in Barcelona, can perform 20, 530 GFlops, and it is the fastest computer in Europe. Its manufacturer: IBM, of course. *Mare Nostrum* was born by the Spanish Government initiative in 2004 to support research in areas like bioinformatics, genomics, information based medicine, quantum chemistry, computational fluid dynamics, astronomy, optics, environmental sciences and financial engineering. All these activities are regulated by a consortium, which establishes priorities [11].

NEC's Earth Simulator, ranked 3.- Installed in 2002, in Yokohama, Japan, hold the record of the 'fastest' for two years. It was manufactured by NEC (Nippon Electric Co.), has 5,120 processors, and reaches 35,860 GFlops. Built with the purpose to enable accurate prediction of the future by modeling present conditions based on data about the past, can simulate the *entire earth system* [12]. The most notable feature of this machine is that is one of the few vector computers that appears in the list.

SGI Altix Columbia, ranked 2.- With 10,160 processors and 51,870 GFlops, this computer, installed in 2004, Mountain View, a NASA facility in the U.S., will bring discoveries to support agency missions. One of the tasks of this machine is to calculate safer trajectories for shuttles and ships, as well as hurricane tracking, for example. This is not the first time NASA has one if the fastest in its time, from 1975 to 1981 the *AMES NASA computer* hold the record of the fastest (150 MFlops) [13].

Blue Gene, ranked 1.- Installed in 2004, with 32,768 processors, and 70 GFlops, this machine is IBM's fastest machine. It is to note that uses IBM's *PowerPC processors* (based on *RISC technology*, in contrast to Intel's *Pentium CISC technology*), because *RISC* uses short and fixed sized instructions, better suited to work in supercomputing [14]. It was originally created as a result of a five-year project that cost $100 million in research alone. It started life as a research tool for protein-folding studies [15]. Today there are a number of labs and universities working on a list of applications for Blue Gene in the areas of life sciences, financial modeling, hydrodynamics, quantum chemistry, molecular dynamics, astronomy and climate modeling [16], in the list of applications to be considered are Alzheimer and Mad Cows dynamics, as well as simulations for the nuclear stockpile of the U.S [17].

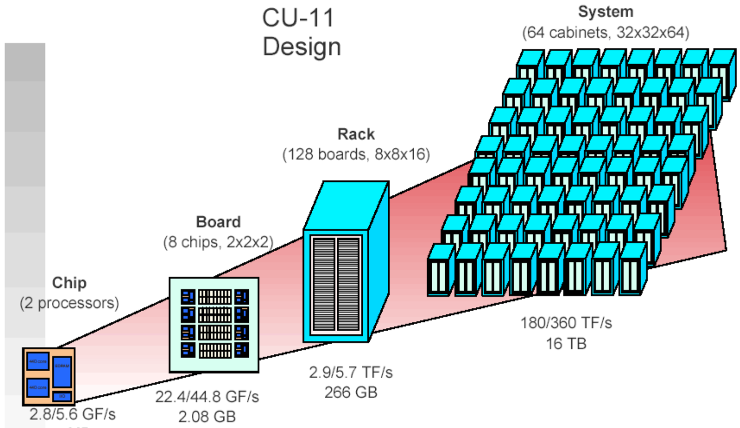

Fig 9. The Blue Gene. IBM Research [18]

Deep Blue.- Also uses a brute force approach to calculate chess positions. Originally started at *Carnegie Mellon University*, was the first computer to beat a human in chess (Garry Kasparov) and holds that record after beating Kasparov in the 1997 rematch. “The brute force of calculation”, said Kasparov [21]. Deep Blue is not the same as Deep Junior, which also draw with Kasparov in 2003 (a computer programmed by Israeli people, but perform 3 million operations per second, compared to the 300 million combinations per second of Deep Blue).

But Deep Blue not only plays chess: with its technology it can clean up toxic waste sites, forecast weather, model financial data, design cars, and develop drug therapies [22].

6. New Trends and Challenges

Improve user interfaces.- Because software for supercomputers focuses on hardware optimization, the less essential parts of the machines have a very limited budget, making user interface not very friendly. To this end there has been companies like Sillicon Graphics and Nvidia, which have been developing technology for supercomputers.

Cooling.- Supercomputers consume large amounts of energy and produces heat. To this end, there are many methods to create less energy comsuming machines as well as better cooling devices, so that the safety temperature of work is not exceeded, causing malcfuncionts.

Non-Uniform Memory Access-NUMA.- It is the next step from Symmetric Processing (a series of processors connected to a bus sharing a logical memory) and it works like a cluster, but tighly coupled in a cache-shared architecture. Because they use massive amounts of data they need to retrieve and store this data quickly

Redundant Array of Independent Disks.- Combine multiple drives into a single logical unit (like NUMA) allowing data integrity and faster memory access instead of a single hard-drive.

Vector Processing.- It is a more efficient technology once implemented, but requieres specialized hardware which is to this day, expensive.

Parallel File Systems.- Can access different or same files simoultaneously. Usually goes with RAID and NUMA.

7. Conclusions:

Although early applications were primarily military ones , by the 1960's there was a growing supercomputer industry with non-military applications. The United States has been in the leader innovator since 1945, the only serious challenge is Japan, but even their NEC technology is originated from U.S. research.

As mainstream computing has become ubiquitous, supercomputing applications are using more often commodity components and applications (*NVdia*, *Sillicon Graphics* for GUI's).

Supercomputers play a fundamental role in scientific discovery. At the present time supercomputing is used to tackle stockpile reduction, climate and earthquake prediction, manufacturing, and other areas in industry. But they also inspire traditional postures (like man vs. machine challenges) and culture.

Right now cluster array is the best technology so far, and its caching techniques perform well, however in the future they will need to address faster memory access to reduce the gap in CPU speed and memory access.

Because of the use of commodity technologies, there has been very little innovation in the past decade. This was supported by a lowering in the cost of hardware, but Supercomputing will find it hard to sustain such a framework in the next ten years.

8. References

[1] Tony Sale, Curator of Bletchley Park Museum [online] Big Oxford Computer, http://www.codesandciphers.org.uk/lorenz/colossus.htm [2005, Feb 16]

[2] Kevin W. Richey [online] Virginia Tech/Norfolk State University http://ei.cs.vt.edu/~history/ENIAC.Richey.HTML [2005, Feb 16]

[3] Ann Kellan [online] CNN news. http://www.cnn.com/TECH/computing/9805/14/india.nuke.computers/ [2005, feb16]

[4] International Charity Foundation
for History and Development of Computer Science and Technique
ICFCST http://www.icfcst.kiev.ua/museum/Missile.html [online] [2005, Feb 16]

[5] Cunningham & Cunningham, Inc. [online] http://c2.com/cgi/wiki?FortranLanguage [2005, Feb 17]

[6] http://www.intel.com/technology/itj/q11998/articles/art_1.htm

[7] Choon, Eggenberger et al. [online] Columbia University http://comet.ctr.columbia.edu/research/architectures/emulation.html [2005, Feb 16]

[8] Top500.org [online] http://www.top500.org/lists/2004/11/charts.php?c=0# [2005, Feb 16]

[9] Frank Da Cruz [online] Columbia University, New Cork
http://www.columbia.edu/acis/history/backus.html [2005, Feb 16]

[10] LinuxWorld [online] http://www.linuxworld.com/story/44799.htm [2005, Feb 16]

[11] Barcelona Supercomputing Center [online] http://www.bsc.org.es/ [2005, Feb 16]

[12] NEC Computers [online] http://www.es.jamstec.go.jp/esc/eng/ES/birth.html [2005, Feb 16]

[13] NASA [online] http://www.nas.nasa.gov/News/Science/2004/10-28-04_fvgcm.html [2005, Feb 16]

[14] William Stallings. Organización y Arquitectura de Computadoras, México, 2002.

[15] IBM [online] http://www.research.ibm.com/bluegene/ [2005, Feb 16]

[16] ComputerWorld [online] International Data Group
http://www.computerworld.com/hardwaretopics/hardware/story/0,10801,97347,00.html [2005, Feb 16]

[17] Linux Lebanon [online] http://linux.com.lb/wiki/index.pl?node=DEEP%20INTO%20Blue%20Gene
[2005, Feb 16]

[18] IBM Research [online]
http://www.research.ibm.com/bluegene/BG_External_Presentation_January_2002.pdf

[2005, Feb 16]

[19] Electronic Frontier Foundation [online] http://www.eff.org/Privacy/Crypto/Crypto_misc/DESCracker/
[2005, Feb 16]

[20] Electronic Frontier Foundation, Idem.

[21] CNN [online] http://www.cnn.com/2003/TECH/fun.games/02/08/cnna.kasparov/ [2005, Feb 16]

[22] IBM Research [online] http://www.research.ibm.com/deepblue/learn/html/e.shtml [2005, Feb 16]