Joint International Topical Meeting on Mathematics & Computation and Supercomputing in Nuclear Applications (M&C + SNA 2007)
Monterey, California, April 15-19, 2007, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2007)

EUROPE BACK IN HPC WORLD
In the Framework of the New European Program for Research (PCRD 7), an Example of Synergy Between Research, Industry and Defense: The CEA HPC Program for 2010

Jean Gonnord
Program Director for numerical simulation & computer sciences
CEA/DAM
France

ABSTRACT

On June 26 2006 CEA/DAM and BULL received from TOP500 the awards for the number one machine in Europe. For the first time since the creation of TOP500 in 1993 the most powerful machine in Europe was design and fabricated by a European company. Six months after, the EU announced, as part of PCRD 7, an important program aiming to give European Research world class computing power and to masteries the necessary technologies. This paper describes the long way which has taken back Europe in the high power computing race.

Key Words: Europe, high-performance computing

1. GENESIS

When in 1996, after the signature by France of the Comprehensive Test Ban Treaty (CTBT), CEA/DAM set up, in order to warrant the safety and the reliability of the French nuclear weapons stockpile, the Simulation program, we faced the ruins field of the European HPC industry.

This was the direct consequence of the failure of the large ‘computational projects’ at the beginning of the 1990s. The European intensive computing industry disappeared and only a few businesses survived. This was the case for Meiko in Great Britain for example, which after its financial collapse was bought by the Italian firm Finmeccanica. Renamed Quadrics, this company has been producing the ‘Rolls Royce’ of networks. In France, after a long period in the desert, Bull came back to the forefront with the Tera-10 machine.

With the almost non-existent industrial framework and the lack of any strategy, European countries have been using a ‘cost base’ policy in intensive computing. With such a policy, HPC is no more than a simple tool used in a few disciplines. Each of them has to invest using its basic research funding. So the aim is to get the cheapest machine as possible. This has strong odd effects : users are self-regulating and depend totally on non European vendors to define what will be the future of HPC and this make the problem bigger…Last but not least, the idea of “free computing”, like is seen the Ethernet model, lead a lot of research agency not to invest in HPC thinking they had an other, as powerful and much cheaper, solution with the “Grid”.

To answer for 2010 the gigantic challenge in high power computing involved by the simulation program, that is 100000 times the power we had (40 times other the Moore Law) CEA/DAM set up the TERA program with a milestone of 100 sustained teraflops to be reach in 2010. Such a challenge seemed at that time impossible.

2. THE FUNDAMENTAL CHOICES

The central part of the Simulation program is a weapon functioning simulator. Around 100 computing engineers and mathematicians have been working on this simulator for almost ten years. They write software, millions of lines of code, developed from ‘models’ established by an equal number of physicists and validated in detail by referring back to past experiments. This colossal task is ongoing and increasingly sophisticated models are being included in the simulator. Considering the cost of development of the simulator, and the fact that computer life time do not exceed four years the choice of a single software architecture for the simulator in 1997: a parallel architecture allowing to use computing power much over the Moore law with two level of parallelism giving us a degree of freedom for machine hardware architecture between pure MPP to very large SMP. With these constraints CEA/DAM launched in 1998 the TERA program with the goal of delivering for the operational use of the simulator in 2010 a sustained computing power of a few hundreds teraflops which means a petaflop class machine.

3. THE CEA TERA PROJECT

The TERA project is the way, established by CEA/DAM, to have this power available on time. To follow the development of the simulator and the capacity of the market we established in 1998 a road map with 3 milestones : 1 sustained Teraflops in 2001, 10 sustained in 2005 and 100 teraflops in 2009. This very ambitious project was and is still a gigantic challenge. It implied a jump on the computing capability of a factor 1000 in 2001 essentially obtained by the change of architecture from vectorial to parallel, then a constant increase of a factor 10 every 4 years which is 60% over the Moore Law.

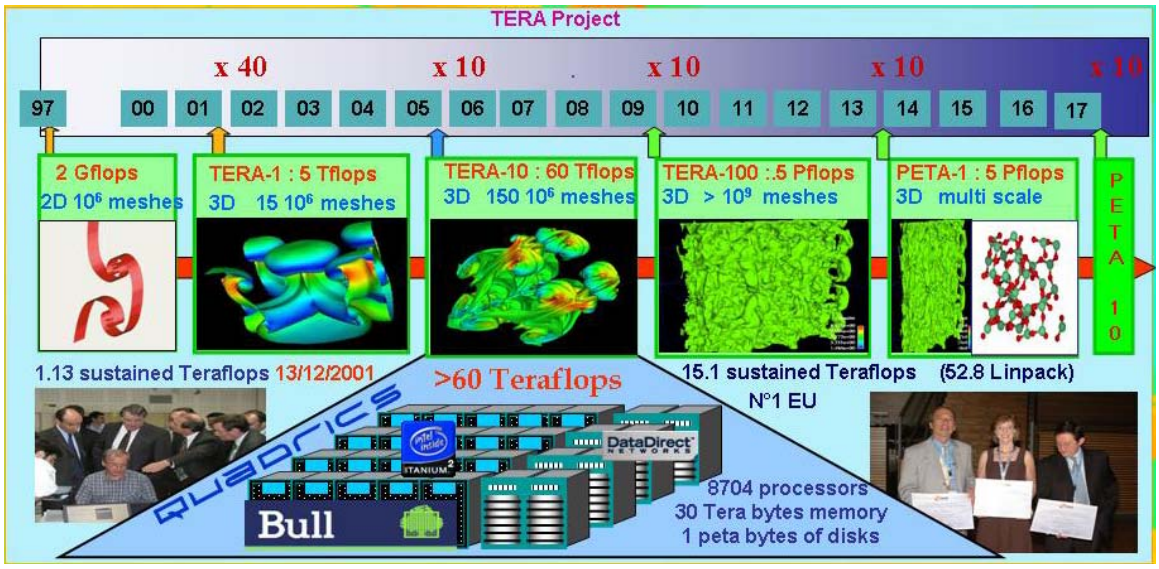

The first step, which was also the demonstration of the feasibility of the chosen architecture "Cluster of SMP" was a success. We obtained 1.13 sustained teraflops from the COMPAQ machine with more than 2500 alpha processors.

We learnt a lot from this machine:

- First, that such a monster was not only running very fast but was also gigantic producer of data. TERA-1 had been producing every day of the year several terabytes leading to petabytes archival…Solving the I/Os problem becomes the main challenges for TERA-10.

- Second, that in front of a so complex and unique machine, we could not rely on a vendor to solve problems that we were the first to discover. Using Open source software had been on other challenge for TERA-10.

This second step was demonstrated on time at the end of year 2005 with the BULL TERA-10 machine: a more than 60 teraflops peak computer based on 4352 INTEL Montecito chips, QUADRICS ELAN-4 networks and DDN disks. TERA-10 which rated 53.8 Linpack teraflops is to day, for one year, in full production and is a big success. Besides the classified production, several open scientific challenges have been solved using the machine like the first complete simulation of a PET scan exam on a human body leading to better prevention of cancer, or the simulation of the Dassault Falcon 7X, the first business aircraft totally design using numerical simulation…

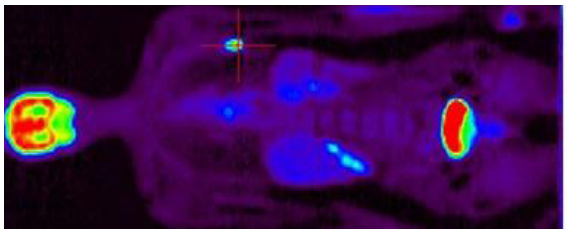

Simulation of PETSCAN Tomography

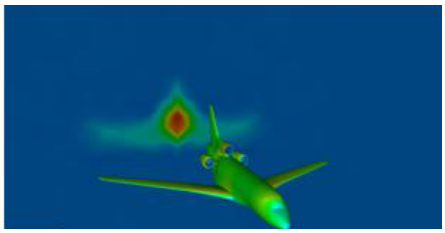

Simulation Dassault Falco 7X

We are now working on the third step: a petaflop class machine. This time the challenge will be electrical power consumption and cooling.

4. PREACHING IN THE DESERT

From the beginning of the TERA project we knew that sooner or later we will have to face the problem of dependency on vendors that we won't be able to influence. To answer this question we realized that we could not act alone and decided to establish a strategy in order first to promote numerical simulation among research and industry, second to develop the necessary technologies. The idea was to put together, around the CEA team of experts, people from research labs, university and industry, and use the synergy between research, industry and defense programs. This took nearly five years – a very long time- . At this time, in France and Europe, most of industrials were not ready to buy big machine thinking that the return on investment was too far away and most of the researchers, thinking that the "computing budget" would be taken on their own research budget, didn't push. CEA had to show the way alone and decide to put on a single site all his computing capacity and experts teams. In 2003 the "CEA scientific computing complex" was created at Bruyères le Châtel.

By creating this complex, the CEA wanted to get the most out of the synergy in its defence-industry-research programmes and generate spin-offs from the numerical simulation programme. Nearly one hundred and fifty CEA engineers and researchers are now working there. The complex is made up of the Defence Computation Centre with the Tera machine, the CCRT (Centre de Calcul Recherche et Technologie) which is open to all, and finally, a centre for experimentation where our experts work with people from university and industry to develop and test technologies for HPC. The complex is managed by the Ministry of Defence for the numerical simulation programme and by a committee on which each partner is represented in proportion to their investment for the CCRT part. Today the CEA has a little over one half of the CCRT shares. The remainder belongs to large corporations (EDF, SAFRAN…) or laboratories like Onera. With Tera-10 and the new BULL CCRT machine the overall capacity of the complex will be over the100 teraflop mark in 2007.

One year after, the CEA scientific computing complex becomes the hub of a much wider operation: Ter@tec. The association aims to unify all parties interested in numerical simulation around the scientific computing complex: researchers, industrials, and technology users and suppliers. Partners are sharing the spin-off of CEA defence programme and from here will bring Europe back up to the top level in high-performance computing. To day more than fifty industrial and academic labs are belonging to Ter@tec which has become a powerful lobby to promote numerical simulation and HPC technologies.

5. UNDER THE SUNLIGHT

5.1. On the French Side

In 2004 Medias, deciders and politicians suddenly discover how far away France, and Europe in general, were from leading country in HPC like USA and Japan. Audits and reports were asked by the French government leading to two new entities.

- The ministry of Economy profiting of the announcement of its competitiveness policy create in 2005 around Paris SYSTEMATIC: Competitiveness Center for Complex System. SYSTEMATIC which focus on three markets: transport, telecommunication and technologies for homeland security is funding developments on software and hardware technologies. The funds are provided at 50% by the participating industrials and 50% shared between ministry of industry and regions. Three big projects issued from Ter@tec and directly interesting HPC have been decided :
  - the development of a server of next generation: FAME2,
  - the development of the technologies for a petaflop machine : POPS,
  - and the installation of a 50 Gigabit network between the partners of Paris region.
- The ministry of research created in mid 2006 an agency : GENCI (National HPC infrastructure) with the goal of
  - Regrouping all French research computing centre under the same management,
  - Doubling the funds allocated for HPC,
  - Being the unique French negotiator in HPC with the UE.

5.2. On the European Side

Europe : a dwarf in HPC

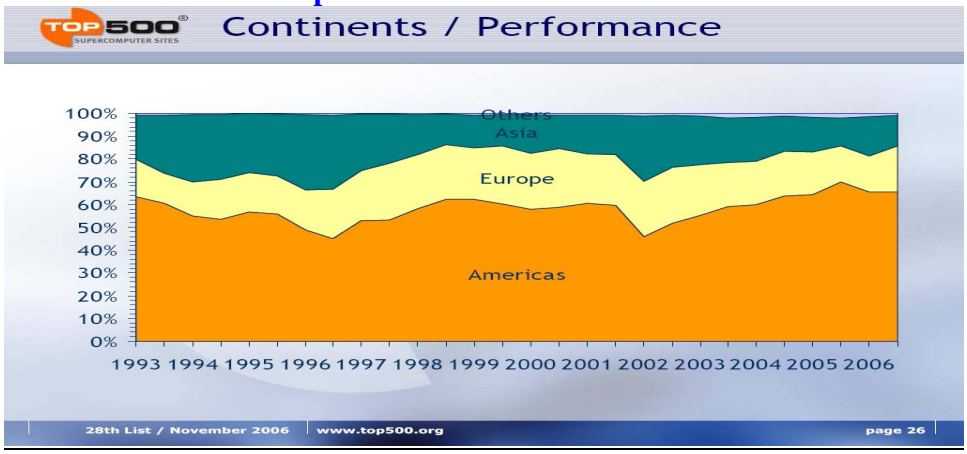

Evolution of Top500 list from its creation to 2006

In parallel, three countries UK, Germany and France created in 2005 “HPC Eur” with the goal of promoting HPC and lobbying UE in the scope of the preparation of the European Program for Research and Development PCRD7 (2007-2013). Joined by Spain in 2006, they issue a complete report mid 2006 the “Scientific case” showing the strategic importance of HPC in all fields of sciences, the more and more important gap between Europe, USA, Japan and soon China, and ask for a dramatic effort for HPC in PCRD 7.

In June 2006 they constituted with seven other countries the “HPC European Taskforce” (HET) aiming to put in common resources necessary to provide their scientific communities with high computing power.

The result of this intense lobbying was the decision by EU, end 2006, to include HPC in its infrastructure roadmap and to issue on the beginning of this year a call for a project of “HPC European Agency”. The eleven countries of HET are now writing an answer due on may 2nd. They should propose to federate in one entity all country efforts for Tier 0 computing machines (petaflops scale in 2009/2013) and to build a sustainable HPC ecosystem. Three to four Tier 0 computing centre should be installed between 2009 and 2012 allowing Europe to be permanently on Top 5 of the Top500 list coordinated and fund by this entity. The new agency will not only have the responsibility of keeping the computing power for European research at a world class level, she will have to ensure that the necessary technologies to keep this level are mastered in Europe. Pre-commercial procurement will be used and vendors will have to show that a large amount of the development (80%) and fabrication(40%) for these “Tier 0” machines is made in Europe.

If confirm by UE, the consortium of countries answering to the call will have two years (2008/2010) to:

- set up the Agency that is defining
  - the governance,
  - the budget,
  - the functioning rules,
  - the location of “Tier 0” sites,
  - the way they will distribute to all European research the computing power.

- Prepare the first calls for the machines to be ordered from end 2009 by the new agency
  - Defining and testing prototype for new technologies,
  - Specifying the machines from user needs and capacity of using new technologies,
  - Preparing the applications software to use these type of machines
  - Developing in conjunction with other European programs (DEISA) the middleware necessary to use these machines anywhere in Europe.

Such a decision will change totally the HPC landscape in Europe.

6. CONCLUSION

CEA has been and is a major actor of this evolution as the largest research and technology agency in Europe. Its 200 expert’s team, set up for the Simulation Program, has been and will be the kernel of all developments in HPC technology in Europe. After creating the Ter@tec ecosystem, CEA has launched in last January the TALOS (Technologies for Advanced Large scale Open Supercomputing) alliance which gather BULL, Quadrics, INTEL, HLRS (Stuttgart computing centre) and CEA with the goals of:

- Mastering the technical challenges in all aspects of HPC and large scale data management,
- Providing European research and industry with world class tool and expertise.

This Alliance is open to any partners sharing these goals and is one of the key element to help the European Union responding to the challenge of supercomputing.

A Bronze giant symbol of new technologies offered by Zeus to Europa