

# Supercomputers – Prestige Objects or Crucial Tools for Science and Industry?

Hans W. Meuer <sup>a 1</sup>, Horst Gietl <sup>b 2</sup>

<sup>a</sup> *University of Mannheim & Prometheus GmbH, 68131 Mannheim, Germany;*

<sup>b</sup> *Prometheus GmbH, 81245 Munich, Germany;*

---

*This paper is the revised and extended version of the Lorraine King Memorial Lecture Hans Werner Meuer was invited by Lord Laird of Artigarvan to give at the House of Lords, London, on April 18, 2012.*

**Keywords:** TOP500, High Performance Computing, HPC, Supercomputing, HPC Technology, Supercomputer Market, Supercomputer Architecture, Supercomputer Applications, Supercomputer Technology, Supercomputer Performance, Supercomputer Future.

---

<sup>1</sup> e-mail: [hans.meuer@isc-events.com](mailto:hans.meuer@isc-events.com)

<sup>2</sup> e-mail: [horst.gietl@isc-events.com](mailto:horst.gietl@isc-events.com)

## Content

1	Introduction.....	3
2	The TOP500 Supercomputer Project.....	3
2.1	The LINPACK Benchmark.....	4
2.2	TOP500 Authors .....	4
2.3	The 39 <sup>th</sup> TOP500 List since 1993.....	5
2.4	The 39 <sup>th</sup> TOP10 List since 1993.....	5
2.5	Competition between Countries in 39 <sup>th</sup> TOP500 List .....	6
2.6	Competition between Manufacturers in 39 <sup>th</sup> TOP500 List .....	7
2.7	Competition between Sites .....	7
2.8	Performance Development .....	8
3	Supercomputer Applications in Science and Industry.....	10
3.1	The Oil and Gas Challenge .....	10
3.2	Weather Forecast .....	11
3.3	Numerical Simulation in Aircraft Aerodynamics .....	13
3.4	Simulating the Galaxy Population in Dark Energy Universes: The Millennium-XXL Project .	14
3.5	Big Data .....	15
3.6	The Blue Brain Project .....	16
4	Alternative Performance Metrics.....	17
4.1	HPC Challenge Benchmark .....	18
4.2	The GREEN500 List .....	18
4.3	The GRAPH500 List .....	19
5	The Future of Supercomputing .....	19
6	Acknowledgement.....	21
7	References.....	21

## 1 Introduction

This paper is intended for people interested in High Performance Computing (HPC) in general, in the performance development of HPC systems from the beginning in the 1970s and, above all, in HPC applications in the past, today and tomorrow. Readers do not need to be supercomputer experts.

In wiseGEEK [1], a supercomputer is defined as follows: *A supercomputer is a computer which performs at a rate of speed which is far above that of other computers. Given the constantly changing world of computing, it should come as no surprise to learn that most supercomputers bear their superlative titles for a few years, at best. Computer programmers are fond of saying that today's supercomputer will become tomorrow's workstation.*

To give an example: In 1986, the most powerful supercomputer in the world was the famous Cray2, available for roughly US\$ 22 million at that time. Last year, Apple released their iPad2 tablet computer, the performance of which is two-thirds of Cray2's performance, but the price of which is only US\$ 500. This means that, in 25 years, the price has gone down by a factor of 44,000, so today's supercomputer is tomorrow's (tablet) computer.

The father of supercomputers is Seymour Cray, former chief designer at Control Data Corporation (CDC). In the 1960s, he had already developed powerful mainframe computers like the CDC 6600/7600. In 1972, Seymour Cray left CDC and founded his own (supercomputer) company, Cray Research. The first product was the Cray1, the very first supercomputer in the world; the first system was delivered to Los Alamos National Laboratory in the U.S.

In this paper, we will first discuss the TOP500 supercomputer project, which was launched at the University of Mannheim, Germany, in 1993; it is the only project in the world that has been successfully tracking and evaluating the supercomputer market for 20 years. Two TOP500 lists are published per year, one at the International Supercomputing Conference in Germany in June and one at the SC Conference in the U.S. in November. This paper will focus on the results of the 39<sup>th</sup> TOP500 list published in June 2012. In chapter 3, we will present six examples of HPC applications in science and industry. After looking at new alternative performance metrics, we will discuss briefly the future of supercomputing at the end of this paper.

## 2 The TOP500 Supercomputer Project

From 1986 through 1992, Hans Werner Meuer presented the Mannheim supercomputer statistics at the opening session of the Supercomputer Seminars at Mannheim University. In 2001, these Seminars were renamed the 'International Supercomputing Conference – ISC'. We noticed an increasing interest in these statistics from year to year. In 1992, we released the last Mannheim statistics, with 530 installed supercomputers worldwide. The statistics simply counted the vector computers installed in the U.S., Japan and Europe, since a supercomputer was synonymous with a vector supercomputer at that time. The statistics of 1986-1992 are listed in [2]. Cray Research was the clear leader, with a constant share of 60%. The U.S. led clearly in 1986 with a share of 50%, but this dropped to 35% in 1992. In Japan, the development was the other way around. Japan started with a share of about 20% in 1986 and went up to 40% in 1992, thus surpassing the U.S. At this time, there was a lot of concern in the U.S. about the 'Japanese supercomputer danger'.

At the beginning of the 1990s, statistics based merely on the name of the manufacturer were no longer useful, though. New statistics were required to reflect the diversification of supercomputers,

the enormous performance differences between low-end and high-end models, the increasing availability of massively parallel processing (MPP) systems, and the strong increase in the computing power of the high-end models of workstation suppliers; symmetric multiprocessor (SMP).

To provide a new statistical foundation, in 1993 we decided to assemble and maintain a list of the 500 most powerful computer systems in the world. Since 1993, our list has been compiled twice a year with the help of high performance computer experts, computational scientists, manufacturers, and the Internet community in general.

In the present list, which is called the TOP500, we publish computers ranked by their performance on the LINPACK benchmark. The main objective of the TOP500 list is to provide a ranking of general purpose systems that are in common use for high-end applications.

The list is available for free at [www.top500.org](http://www.top500.org), and you can create additional sub-lists and statistics from the TOP500 database on your own. At the Apple store, you can also download a free TOP500 app including all TOP500 lists for iPhones and iPads.

## 2.1 The LINPACK Benchmark

As a yardstick of performance, we use the best performance (Rmax) measured by the LINPACK benchmark. We chose LINPACK because it is widely used in the supercomputer world and performance figures are available for almost all relevant systems.

*Jack Dongarra, University of Tennessee*, introduced the LINPACK benchmark, which is based on solving a dense system of linear equations. For the TOP500, we use this version of the benchmark that allows the user to scale the size of the problem. From June 1993 to June 2012, we had to cover more than eleven orders of magnitude with respect to the performance of the systems in all lists. This best LINPACK performance does not reflect the *overall performance* of a system, which no single figure ever could, but it reflects the *performance of a dedicated certain system for solving a dense system of linear equations*. Since the problem is very regular, the performance achieved is quite high, and the performance figures reflect the peak performance very well.

In an attempt to obtain uniformity across all computers in performance reporting, the algorithm used in solving the system of equations in the benchmark suite must conform to LU factorization with partial pivoting. In particular, the operation count for the algorithm must be  $\frac{2}{3}n^3 + O(n^2)$  double precision floating point operations. This excludes the use of fast matrix multiply algorithms like 'Strassen's Method' or algorithms which compute a solution in a precision lower than full precision (64 bit floating point arithmetic) and refine the solution using an iterative approach.

## 2.2 TOP500 Authors

*Hans Werner Meuer* and *Erich Strohmaier*, the two founding authors, launched the TOP500 project in spring 1993 at the University of Mannheim, Germany. Since the beginning, it has been their task to compile all TOP500 lists, to do the daily project work, especially hosting and maintaining the TOP500 Web server, and to secure funding for the project. Another TOP500 author from the very beginning is *Jack Dongarra, University of Tennessee*. He is in charge of further developing LINPACK with respect to the TOP500, e.g., the 'High Performance LINPACK', and developing a reduced LINPACK to keep the runtime manageable for future TOP500 lists. Since 1993 (officially since 2000), the fourth TOP500 author is *Horst Simon, Lawrence Berkeley National Laboratory*. His job is to promote the TOP500 project worldwide, especially within governmental organizations.

## 2.3 The 39<sup>th</sup> TOP500 List since 1993

It is the 20<sup>th</sup> year of the TOP500 list, and after releasing the 39<sup>th</sup> list recently at ISC'12 in Hamburg we will celebrate the 20<sup>th</sup> anniversary of the TOP500 this year: The 40<sup>th</sup> list will be published at SC12 in November in Salt Lake City. At that time, we will have continuously published 40 TOP500 lists with a total of 20,000 systems over 20 years. In the worldwide HPC community the TOP500 is accepted as the only tool for tracking and evaluating the HPC market. One of the main reasons why the TOP500 list is so successful is the fact that we foster competition at different levels: competition between countries, between manufacturers, and between sites.

## 2.4 The 39<sup>th</sup> TOP10 List since 1993

#	Site	Manufacturer	Computer	Country	Cores	Rmax [Pfllops]	Power [MW]
1	Lawrence Livermore National Laboratory	IBM	<b>Sequoia</b> BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	1,572,864	16.3	7.89
2	RIKEN Advanced Institute for Computational Science	Fujitsu	<b>K Computer</b> SPARC64 VIIIfx 2.0GHz, Tofu Interconnect	Japan	795,024	10.5	12.66
3	Argonne National Laboratory	IBM	<b>Mira</b> BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	786,432	8.16	3.95
4	Leibniz Rechenzentrum	IBM	<b>SuperMUC</b> iDataPlex DX360M4, Xeon E5 8C 2.7GHz, Infiniband FDR	Germany	147,456	2.90	3.52
5	National SuperComputer Center in Tianjin	NUDT	<b>Tianhe-1A</b> NUDT TH MPP, Xeon 6C, NVidia, FT-1000 8C	China	186,368	2.57	4.04
6	Oak Ridge National Laboratory	Cray	<b>Jaguar</b> Cray XK6, Opteron 16C 2.2GHz, Gemini, NVIDIA 2090	USA	298,592	1.94	5.14
7	CINECA	IBM	<b>Fermi</b> BlueGene/Q, Power BQC 16C 1.6GHz, Custom	Italy	163,840	1.73	0.82
8	Forschungszentrum Jülich (FZJ)	IBM	<b>JuQUEEN</b> BlueGene/Q, Power BQC 16C 1.6GHz, Custom	Germany	131,072	1.38	0.66
9	Commissariat à l'Energie Atomique CEA/TGCC-GENCI	Bull	<b>Curie thin nodes</b> Bullx B510, Xeon E5 8C 2.7GHz, Infiniband QDR	France	77,184	1.36	2.25
10	National Supercomputing Centre in Shenzhen	Dawning	<b>Nebulae</b> TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU	China	120,640	1.27	2.58

Fig. 1

The current TOP10 systems are completely different than the ones of the 38<sup>th</sup> list while the 37<sup>th</sup> list of June 2011 and the 38<sup>th</sup> list of November 2011 were identical with the exception of the performance of the #1 system, the K Computer. For the first time since November 2009, a U.S. supercomputer sits atop the TOP500 list of the world's top supercomputers. Named Sequoia, the IBM BlueGene/Q system installed at the Department of Energy's Lawrence Livermore National Laboratory achieved an impressive 16.32 Petaflop/s on the LINPACK benchmark using 1,572,864 cores and is 50% more powerful than the K Computer. Sequoia is also one of the most energy efficient systems on the list.

The K Computer, the number one system on the 37<sup>th</sup> and 38<sup>th</sup> lists, was the first system in the world to reach the 10 Petaflop/s threshold with a best LINPACK performance of Rmax= 10.51 Petaflop/s by increasing the number of cores to roughly 800,000. The K Computer currently holds the #2 position on the 39<sup>th</sup> TOP500 list.

In the TOP10, there are three systems from Asia and three from the U.S. Europe is represented with four systems, the No. 4 system at LRZ, Germany, the No. 7 system at CINECA in Italy, the No. 8 system at Forschungszentrum Jülich, Germany, and the No. 9 system at GENCI, France. Cray's Jaguar at Oak Ridge National Laboratory, the former No. 1 system in the world, is now No. 6. It is interesting

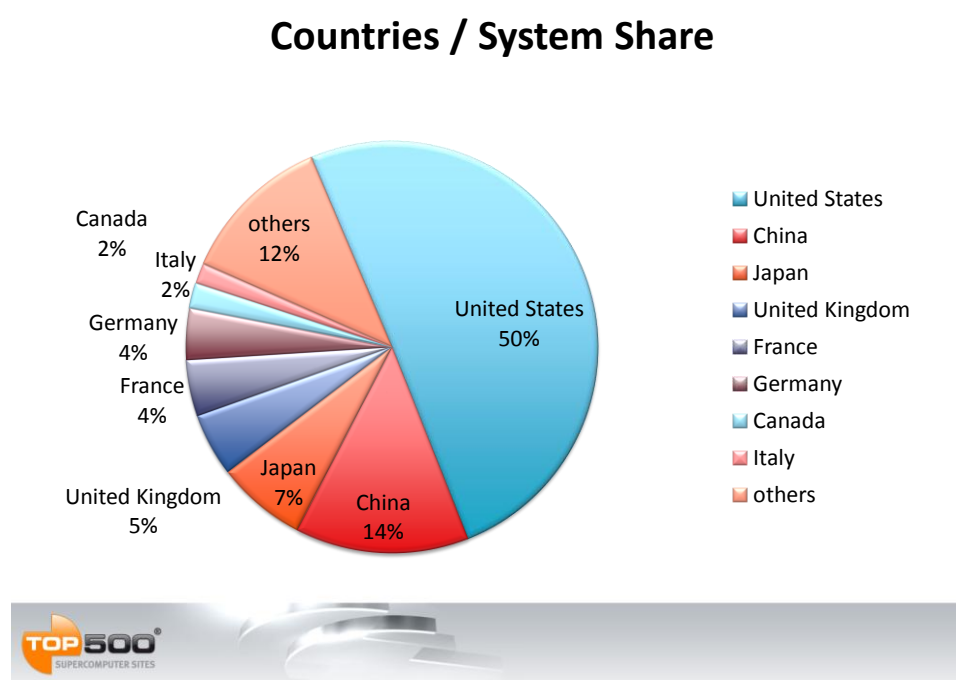
to see that IBM has five systems in the TOP10. The number of Petaflop/s systems has now increased to the point where there are now 20 Petaflop/s systems on the current TOP500 list.

The new Mira supercomputer, an IBM BlueGene/Q system at Argonne National Laboratory in Illinois, debuted at No. 3, with 8.15 Petaflop/s on the LINPACK benchmark using 786,432 cores. Another U.S. system in the TOP10 is the Cray Jaguar at Oak Ridge National Laboratory in Tennessee, which was the top U.S. system on the previous list and now clocks in at No. 6.

Italy makes its debut in the TOP10 with an IBM Bluegene/Q system installed at CINECA. The system has a LINPACK performance of 1.69 Petaflop/s. In all, five of the TOP10 supercomputers are IBM BlueGene/Q systems. France is represented in the TOP10 with a homegrown Bull supercomputer. China, which briefly took the No. 1 and No.3 spots in November 2010, has two systems in the TOP10, with Tianhe-1A at NUDT in No. 5 and Nebulae at the National Supercomputing Centre in Shenzhen No. 10.

Interestingly, the K Computer has the highest power consumption ever in the TOP500 list, i.e., 12.659 MW. Even the new No. 1 system, Sequoia, has less power consumption with 7.841 MW.

## 2.5 Competition between Countries in 39<sup>th</sup> TOP500 List

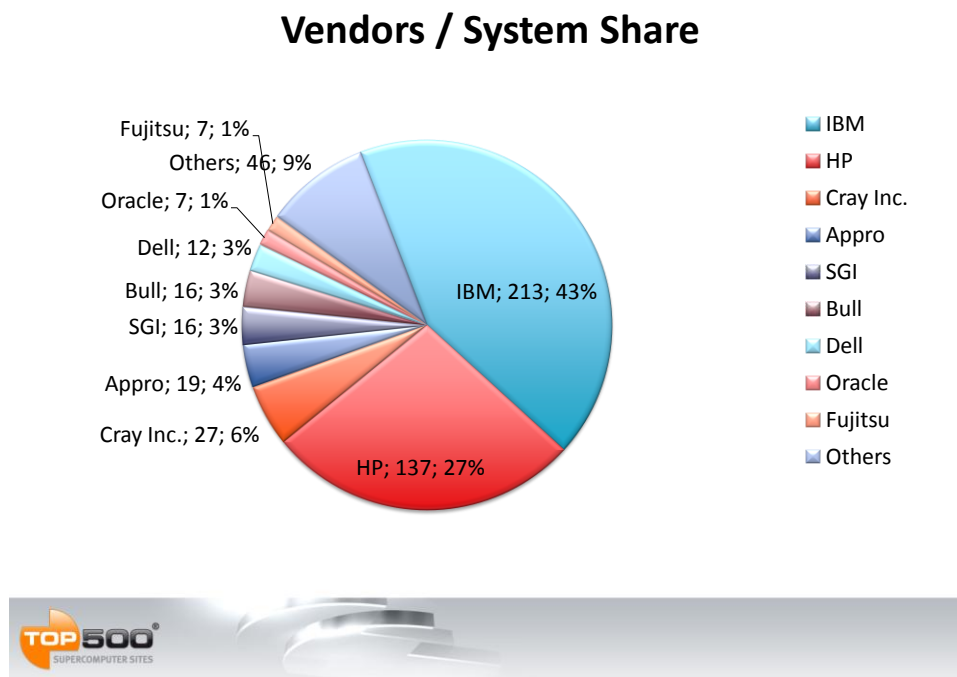


**Fig. 2**

The U.S. is the clear leader (which has been the case since 1993) with a systems share of more than 50%, with 252 of the 500 systems (down from 263 in the previous list). In the first TOP500 list, the U.S. had already twice as many installed systems compared to Japan, so the 'Japanese supercomputer danger' of the early 1990s turned out to be a tempest in a teapot. Surprisingly, China holds the number two position with 68 systems, still very far behind the U.S., but also far ahead of the number three nation, Japan, which has 35 systems. It seems Japan is recovering after a long, disappointing time for manufacturing supercomputer systems. They have convincingly demonstrated that they are still able to build the most powerful system in the world. We have to wait and see if and how they will increase their 7% share, with Fujitsu now offering commercial versions of the K Computer.

The European share (106 systems) is still lower than the Asian share (122 systems). The U.K. is currently ahead of France and Germany, with 25, 22, and 20 systems respectively, but these countries have continually traded places over the past years. However, France is likely to take the lead in the next years, as it is the only country in Europe that builds supercomputers.

## 2.6 Competition between Manufacturers in 39<sup>th</sup> TOP500 List



In the 39<sup>th</sup> list, IBM has the clear lead with a 43% share, and Hewlett-Packard holds the second place with 27%. The other manufacturers do not hold a significant share. Cray, the leader of the first TOP500 list, is now down to 6% from 41% in 1993, but with a share of only 3% a couple of years ago, Cray's situation was even worse. Cray now specializes in very powerful high-end supercomputers, and they are the market leader in the TOP50 with a share of roughly 14%. IBM is represented evenly in all parts of the TOP500 list, while the runner-up, Hewlett-Packard, has only one system in the TOP50, the No. 14 Tsubane 2.0, together with NEC. No other manufacturer in the list has a higher share than 3%.

## 2.7 Competition between Sites

**Fig. 4** shows the 20 most powerful sites since 1993 through 39 TOP500 lists. Thus the percentage in the right-hand column is the site's relative contribution over all 39 TOP500 lists.

In this list, the USA leads with thirteen sites before Japan with four centers (20%). The fact that the U.S. has the five most powerful sites in the world also shows their dominance as a consumer and producer of HPC systems. Europe is represented by Germany (Forschungszentrum Jülich) at position nine, by the U.K. (ECMWF) at position twelve and by France (CEA) at position twenty. Please note that ECMWF is a European and not purely a U.K. site.

## TOP20 Sites since 1993

Rank	Sites	Overall Share (1993-2012)
1	Lawrence Livermore National Laboratory	5.50%
2	Los Alamos National Laboratory	3.30%
3	Sandia National Laboratories	2.61%
4	Government/USA	2.28%
5	DOE/SC/Oak Ridge National Laboratory	2.18%
6	Japan Agency for Marine -Earth Science and Technology	1.60%
7	NASA/Ames Research Center/NAS	1.47%
8	National Aerospace Laboratory of Japan	1.29%
9	Forschungszentrum Juelich (FZJ)	1.25%
10	NCSA / National Center for Supercomputer Applications/ USA	1.07%
11	DOE/SC/LBNL/NERSC	1.07%
12	ECMWF	0.96%
13	Naval Oceanographic Office - NAVO DSRC	0.94%
14	RIKEN Advanced Institute for Computational Science (AICS)	0.94%
15	Semiconductor Company/USA	0.94%
16	University of Tokyo	0.93%
17	Pittsburgh Supercomputing Center	0.88%
18	DoD ERDC DSRC Supercomputer Resource Center/ USA	0.85%
19	IBM Thomas J. Watson Research Center	0.79%
20	Commissariat a l'Energie Atomique (CEA)	0.71%



Fig. 4

## 2.8 Performance Development

### Performance Development

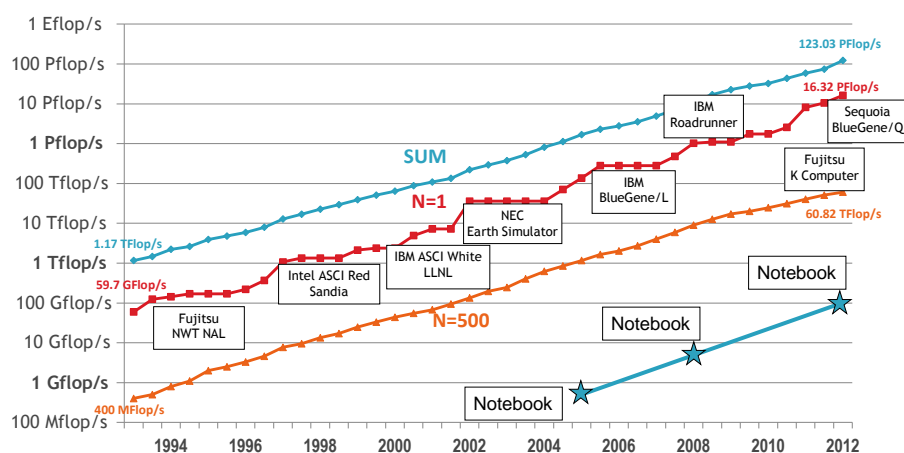


Fig. 5

The Total performance of all the systems on the list has increased by more than 50% since November 2011, reaching 123.03 Petaflop/s. The combined performance of the last list was 74.2 Petaflop/s. In all, 20 of the supercomputers on the newest list reached performance levels of 1 Petaflop/s or more. The No. 500 system on the list notched a performance level of 60.8 Teraflop/s, which was enough to reach No. 330 just seven months ago.



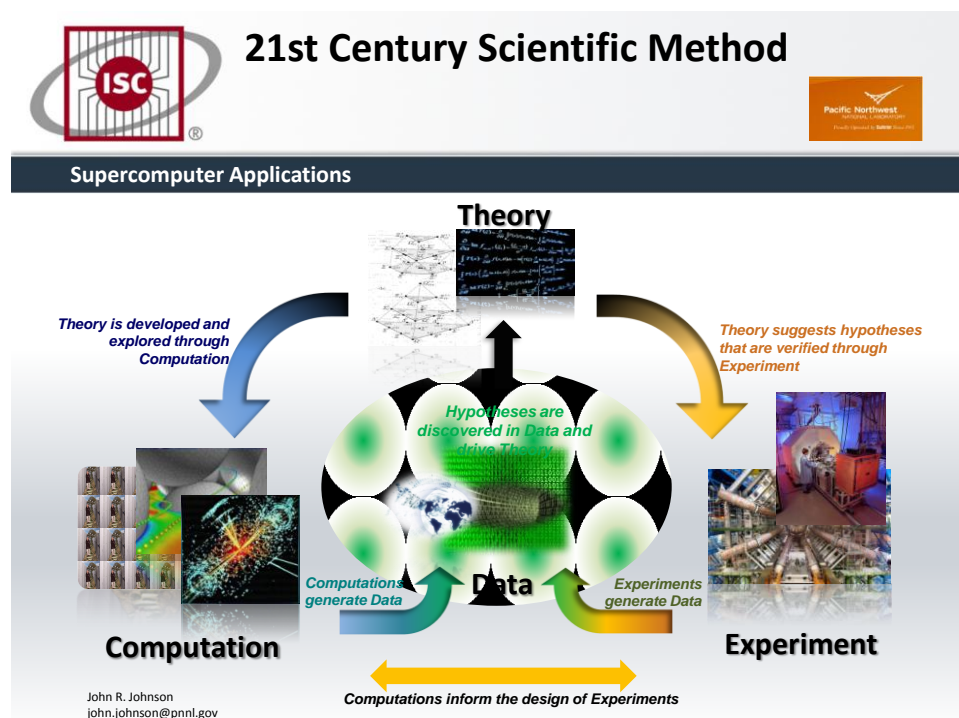
**Fig. 5** shows the performance measured over the last 19 years at position  $N=500$  (entry level performance), at the leading position ( $N=1$ ) and the sum of all 500 systems. The scale on the left-hand side is a logarithmic scale. All curves show an exponential growth. The curves at positions one and 500 are quite different: At number one, we typically see a step function. Once a system has made number one, it remains on the next couple of TOP500 lists. This was true for the Numerical Wind Tunnel (NWT), Intel's ASCI Red, and also for the NEC Earth Simulator, which was ranked No. 1 from June 2002 through June 2004. IBM's Blue Gene/L also held this position at different stages of expansion. If we include a powerful notebook in this figure, we notice that its performance reached 70 GigaFlop/s at the beginning of 2012, which is a six-fold increase in performance within four years.

### 3 Supercomputer Applications in Science and Industry

In the 1980s, Physics Nobel Laureate Kenneth G. Wilson (US, 1982) emphasized that supercomputer simulations added a third pillar of scientific methods to the traditional pillars of experiment and theory.

Supercomputer simulations can be found in various application areas. They are used for highly calculation-intensive tasks such as problems involving quantum physics, weather forecasting, climate research, molecular modeling (computing the structures and properties of chemical compounds, biological macromolecules, polymers, and crystals), physical simulations (such as simulation of airplanes in wind tunnels, cosmology, simulation of the detonation of nuclear weapons, research into nuclear fusion) and as a new application area, data mining or big data. The problems known as the grand challenge problems are problems whose full solution requires semi-infinite computing resources. In this chapter, we will take a closer look at some of these problems.

**Fig. 6** of John R. Johnson, Pacific Northwest National Laboratory, gives an excellent overview of how the three pillars influence each other. It should be mentioned that the experiment in the bottom right corner shows CERN's Large Hadron Collider.



**Fig. 6**

#### 3.1 The Oil and Gas Challenge

The oil and gas industry was one of the first industries to deploy supercomputers. As early as in the 1980s, simulations on supercomputers were used in the exploration of new oil fields. It was a well-known fact that drilling a hole into an oil field without any success was as expensive as buying a supercomputer. It was in 1986, at the very first Mannheim supercomputer seminar, that Werner Butscher, of Prakla Seismos at that time, gave a talk on oil exploration and supercomputers. **Fig. 7** shows the Seismic Marine Data Collection. A vessel with an air gun pulls a string of hydrophones, which collect the reflected shock waves from the vessel. Even at that time, the total data volume for a single exploration was about 1,000 magnetic tapes.

In the 1980s, Prakla Seismos was acquired by oil exploration giant Schlumberger of Houston, Texas. The oil exploration business is still a very important business, with tremendous requirements for HPC power. Today (according to R. Fertl/Schlumberger), Terabytes of data will be produced by single exploration experiments. While the improvement in supercomputing performance between 1986 and 2008 was about  $10^6$ , modern HPC systems have reached a sustained performance for seismic applications of roughly 2,000 times that of 1986. The reason is that nowadays highly sophisticated algorithms (developed in the past 20 years) on HPC systems are necessary to detect new oil reservoirs.



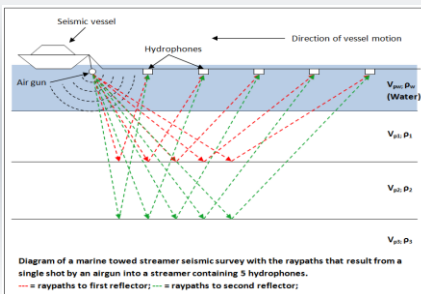
## Early Application Example Oil Exploration and Supercomputers

### Seismic Marine Data Collection

1986

Survey:  
Shot Points:  
Hydrophones:  
Total Data Volume:  
Run Time

Presentation of W. Butscher (Prakla Seismos AG) at ISC'86.  
First Supercomputer Seminar at Mannheim University:  
225 square km  
200.000  
48  
~ 1000 large magnetic tapes  
Cyber 205 Vector Computer: > 100 Hours



Today:

Streamer Lines:  
Length of streamer:  
Receivers/streamer:  
Sampling Rate:  
Data Volume/Survey:  
Modern HPC-System:

Schlumberger (Houston);  
R. Fertl  
8 – 12  
~ 8 km  
2500  
2 ms  
10 TeraBytes  
Factor 2000

**Supercomputer Performance Improvement**  
**2008/1986 : ~  $10^6$**

Fig. 7

The oil and gas industry is facing new challenges: It will be much harder to extract oil and gas from decreasing resources. It is a must for this industry to speed up the energy exploration process by accelerating simulation and analysis of potential resources. Clusters and grids are by far the most economical approach for delivering large-scale systems – it also means that new seismic algorithms have to be developed to cope with the inherent parallelism of these cluster systems.

## 3.2 Weather Forecast

From the beginning of the supercomputer age, weather forecasts and climate simulations have played a dominant role in the development of more powerful HPC (High Performance Computing) systems.

Numerical weather prediction uses mathematical models to simulate the atmosphere and oceans and to predict the weather forecast based on current weather conditions. Since around 1950, it has been possible to produce realistic results for numerical weather predictions using computer simulations.

The mathematical models used can generate either short-term weather forecasts or longer-term climate predictions; the latter are heavily used for understanding and projecting climate change. Weather forecasting and climate simulations demand vast numbers of data points for improving

accuracy of the numerical algorithms. The complexity of the numerical calculations and time constraints of these forecasting models are directly related to the performance of used HPC systems and require the most powerful supercomputers in the world.

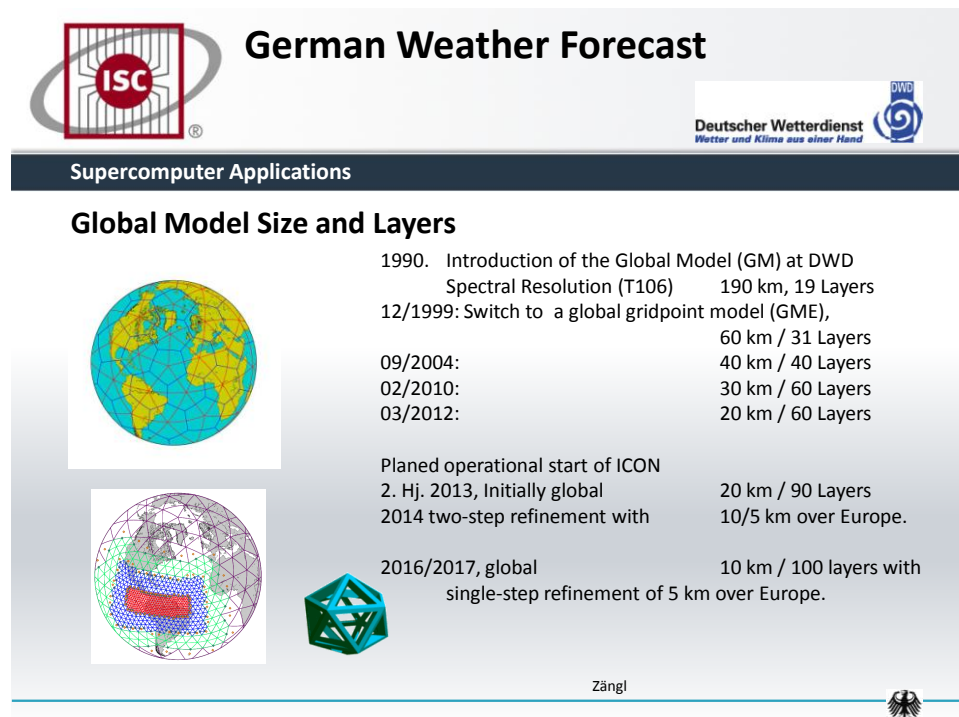
Using these HPC systems, the forecast period of numerical weather simulations, with more than 50% probability, is now in the range of four to six days. Besides the accuracy of the numerical algorithms, the weather simulations are closely related to the quality of weather observations used as input to the forecasts. The reason is that the partial differential equations, underlying the numerical simulations, show a chaotic behavior. It means that small changes in the input data may influence the final numerical atmospheric solution tremendously. Therefore statistical techniques are required to predict atmospheric simulations as precisely as possible.

One important goal for numerical weather prediction is local mesh refinements to enhance local and small-scale accuracy.

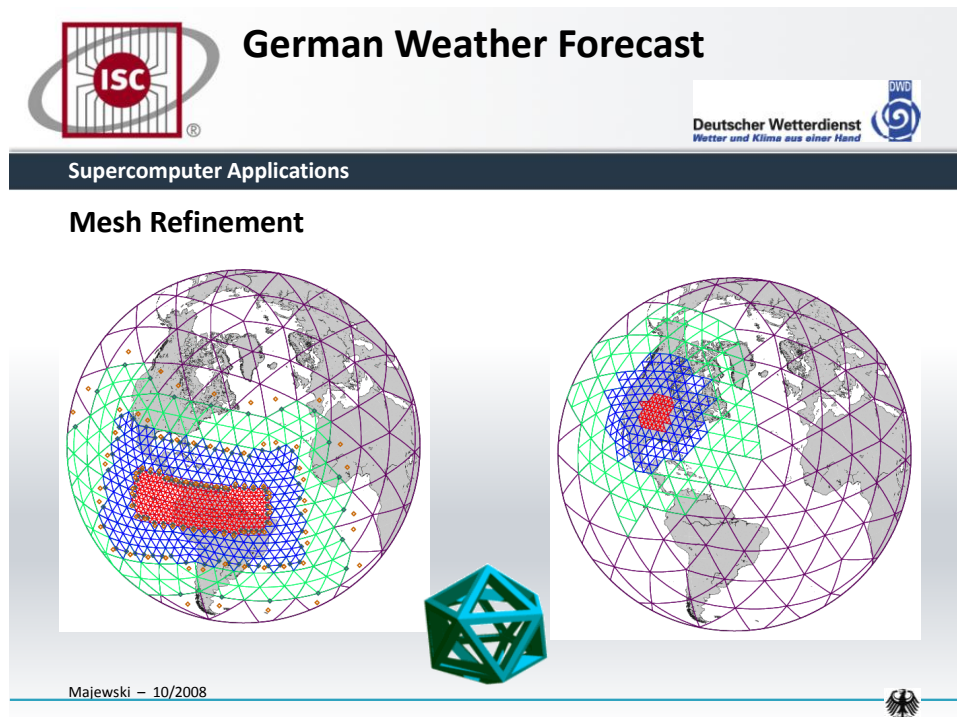
The following figure (**Fig. 8**), courtesy of Deutscher Wetterdienst, Günther Zängl/DWD, shows the mesh refinement development process. In 1990, the numerical grid had a horizontal resolution of 190 km and 19 layers in the vertical direction. In 2012, the corresponding figures are 20 km horizontally and 60 layers vertically. With the planned ICON Project (ICON: ICOsahedral Non-hydrostatic General Circulation Model), the goal is to develop a dynamic model for weather forecasting and climate modeling.

In 2016/2017, the resolution is projected to be 10 km and 100 layers with a spatial refinement over Europe of 5 km in the horizontal direction.

The 3D icosahedron element shown plays a dominant role in the mesh structure and is the basis for the local mesh refinement.



**Fig. 8**



**Fig. 9**

**Fig. 9** shows an exponential growth over time, thus underscoring the validity of Moore's Law. It should be mentioned that the sustained performance of real-life applications such as the weather code, exactly follow Moore's Law.

### 3.3 Numerical Simulation in Aircraft Aerodynamics

Trying to understand the behavior of fluids is one of the most fascinating areas of science. Over four decades, Computational Fluid Dynamics (CFD) continues to be one of the major drivers for advances in HPC.

Numerical simulation of aerodynamic flow is a key element in aerodynamic development work. Not only is the design and optimization of external shapes based on the predictive quality of CFD nowadays, but so is the prediction of all aerodynamic data and the complete flight envelope of an aircraft. The tremendous progress of the past decades in physical modeling, numerical algorithms and computer power is now the basis for developing all current and next-generation aircraft. But this is not the end of the story: CFD now goes into multi-disciplinary simulations (e.g., CFD, structural mechanics), which help aircraft designers optimize next-generation products in a fully integrated way.

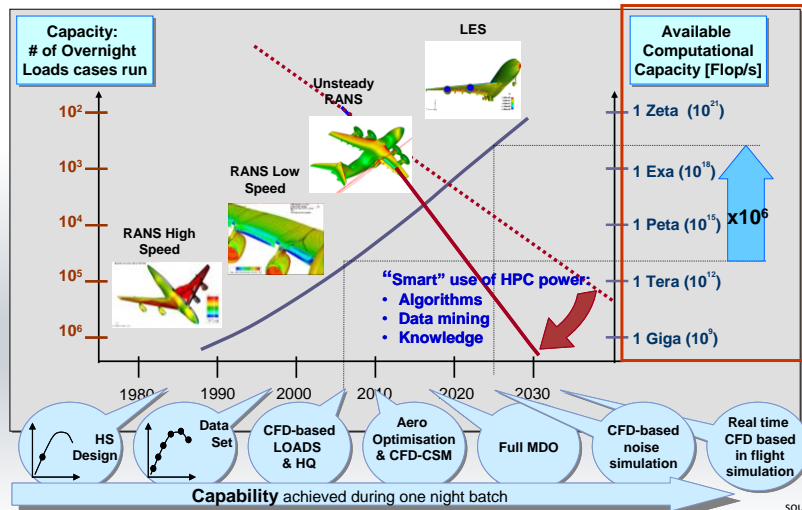
Therefore numerical simulations in aircraft aerodynamics are vital to the whole aircraft industry. **Fig. 10** gives different points of view on the available and required computational capacities at DLR, which co-operates with Airbus (Norbert Kroll/DLR).

It shows that in 2006, the pure CFD solution needed 10 Teraflop/s per run, and, in 2025, a full LES simulation will need 50 Exaflop/s sustained performance. Interestingly, the improvements do not only come from the use of larger HPC systems but also from new algorithms/data mining and better knowledge. The goal for 2030 is to have a real-time based in-flight simulation but this will require a Zettaflop/s supercomputer.



## Numerical Simulation in Aircraft Aerodynamics - Expectations

### Supercomputer Applications



source: Airbus

**Fig. 10**

RANS: Reynolds-Averaged Navier-Stokes equations that are equations of motion for fluid flow where Laminar flow occurs at low Reynolds numbers, and turbulent flow at high Reynolds numbers.

CFD: Computational Fluid Dynamics

CSM: Computational structural mechanics

MDO: Multi-disciplinary design optimization (i.e. fluid dynamics and structural mechanics).

LES: Large Eddy Simulation

### 3.4 Simulating the Galaxy Population in Dark Energy Universes: The Millennium-XXL Project

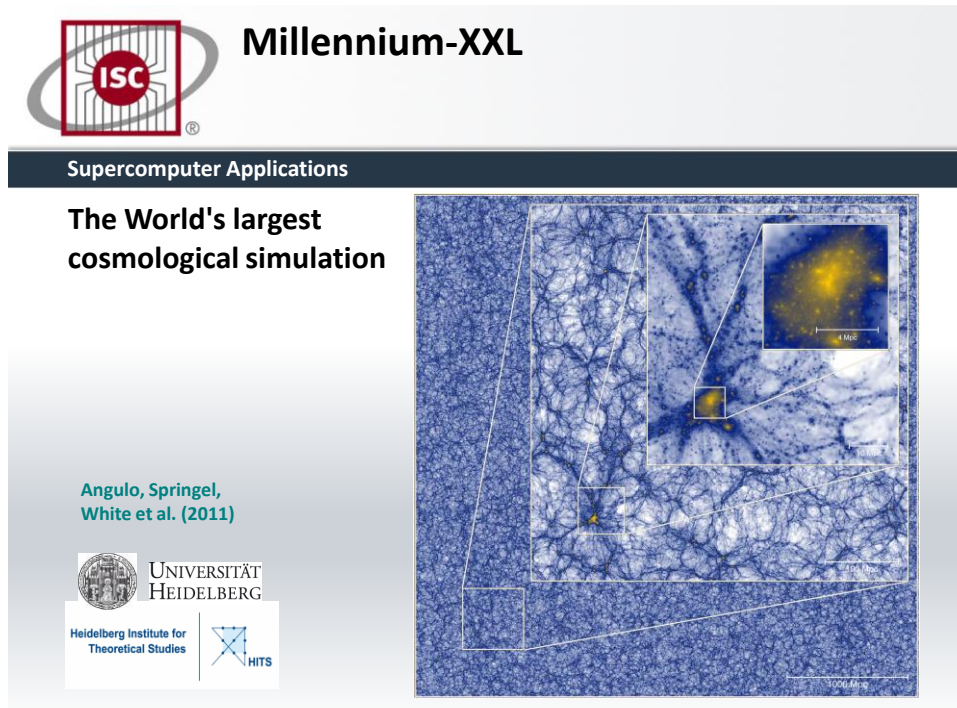
The dark matter density field in the Millennium-XXL simulation, shown on different scales. The image in the background shows the cosmic large-scale structure in a thin slice of thickness 25 Mpc through the simulation box, which has a side-length of 4.1 Gpc (corresponding to 13 billion light years). The filamentary distribution of matter along a 'Cosmic Web' becomes evident in the sequence of zooms that focus on one large galaxy cluster shown in the top right. Such galaxy clusters contain between 1,000 to 2,000 galaxies and are the largest visualized objects in the Universe. Each inset in the composite picture enlarges the structures by a factor of 8. The Millennium-XXL simulation has a dynamic range of 300,000 per dimension, allowing it to simultaneously resolve the internal structure of collapsed objects and to give an accurate account of the quasi-linear and mildly non-linear evolution of large-scale structures in the LCDM cosmology. This is required for the calibration of future precision experiments that try to constrain the nature of dark energy and dark matter, such as the European EUCLID mission of ESA [5].

**Fig. 11** shows the calculation solved for the gravitational interactions between more than 300 billion particles over the equivalent of more than 13 billion years.

Carrying out this computation proved to be a formidable challenge even on today's most powerful supercomputers. The simulation required the equivalent of 300 years of CPU time (compared to a powerful PC system) and used more than twelve thousand computer cores and 30 TB of RAM on the Juropa Machine at Forschungszentrum Jülich, Germany, one of the 15 most powerful computers in the world at the time of execution. The simulation generated more



than 100 TB of data products.



**Fig. 11**

Parsec (pc) means 3.26 light years

Mega parsecs (Mpc) is equivalent to  $10^6$  parsecs

It should be mentioned that at GENCI in France, where they use the new supercomputer CURIE, the computer model simulations for structuring the entire observable universe, from the Big Bang until today, is being performed with 550 billion particles involved. The entire project at GENCI will use more than 30 million hours (about 3,500 years) of computing time on a single CPU. But since nearly all cores of CURIE are used, the runtime should be manageable. More than 150 Petabytes of data (the equivalent of 30 million DVDs) will be generated during the simulations. Thanks to an advanced and innovative data reduction process developed by the researchers, the amount of useful stored data can now be reduced to 1 Petabyte.

### 3.5 Big Data

New fields are contributing to the ever increasing information explosion. Take the LHC at CERN as an example, producing Terabytes of data per day that have to be examined with the help of HPC systems to find yet undetected particles. Or consider the massive amount of data produced by social networks. To extract necessary information out of this bulk of data makes looking for a needle in a haystack seem simple. These applications form a new application field for HPC systems, called big data and the applications used for them are often called data mining or data intensive computing.

The real challenges for the analysis of big data are the underlying datasets with their unstructured behavior. It is well-known that data-intensive sciences have different requirements to HPC systems than those of the traditional computational sciences. Nevertheless it is an open question among experts as to whether the processing of large amounts of unstructured datasets can be efficiently executed on existing parallel processing systems. Many experts think that new HPC architectures are needed to support data mining efficiently and to scale with the growing demand. A first step in this direction will be the new HPC architecture from Flash Gordon at San Diego.

Fig. 12 shows some big data applications.

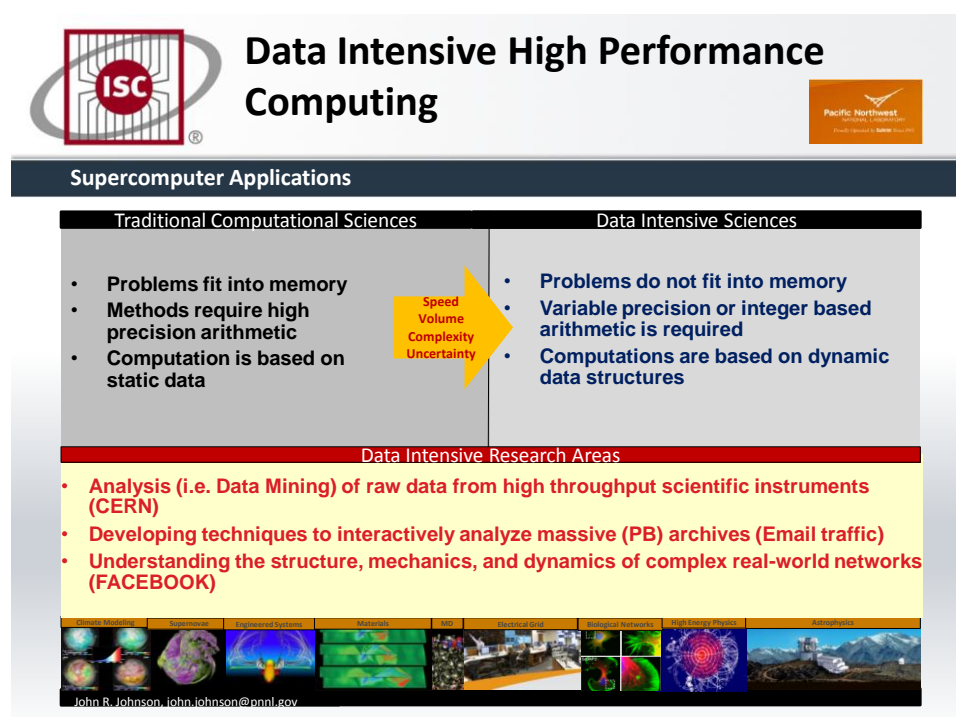


Fig. 12

### 3.6 The Blue Brain Project

The last application presented in this paper comes from computational life sciences, a strategic discipline at the frontier between molecular biology and computer science, impacting medicine and biotechnology and with widespread implications for all of society.

The **Blue Brain Project** [6] was launched at the École Polytechnique Fédérale de Lausanne (EPFL, Switzerland) in May 2005 to study the brain's architectural and functional principles, and Henry Markram has been leading the project from the beginning. The simulation runs on an IBM Blue Gene supercomputer, which is not just an artificial neural network but relies on a biologically realistic model of neurons. It is hoped that this project will eventually shed light on the nature of consciousness. **Fig. 13** illustrates the complexity of the problem as outlined by the technical leader of the project, Felix Schürmann/EPFL.

By 2005, the first single cellular model was completed, and by 2008, the first artificial cellular neocortical column of 10,000 cells was built, which can be considered the smallest functional unit of the neocortex (the part of the brain thought to be responsible for higher functions such as conscious thought). By July 2011, a cellular mesocircuit of 100 neocortical columns with a million cells in total was built. A cellular rat brain is planned for 2014 with 100 mesocircuits totaling a hundred million cells. Finally, a cellular human brain is predicted to be possible by 2023, the equivalent to 1,000 rat brains with a total of a hundred billion cells.

If the simulation is to include the molecular level, which is desirable since it allows studying the effects of gene expression, then molecular dynamics methods are required (called 'Reaction-Diffusion' in **Fig. 13**), which will immediately increase the performance and memory requirements by an order of 100-1,000. Brain simulation tends to be the most challenging supercomputer project in



the future. Henry Markram has applied for a European Flagship Project with a total budget of 1 billion Euros for the next 10 years.

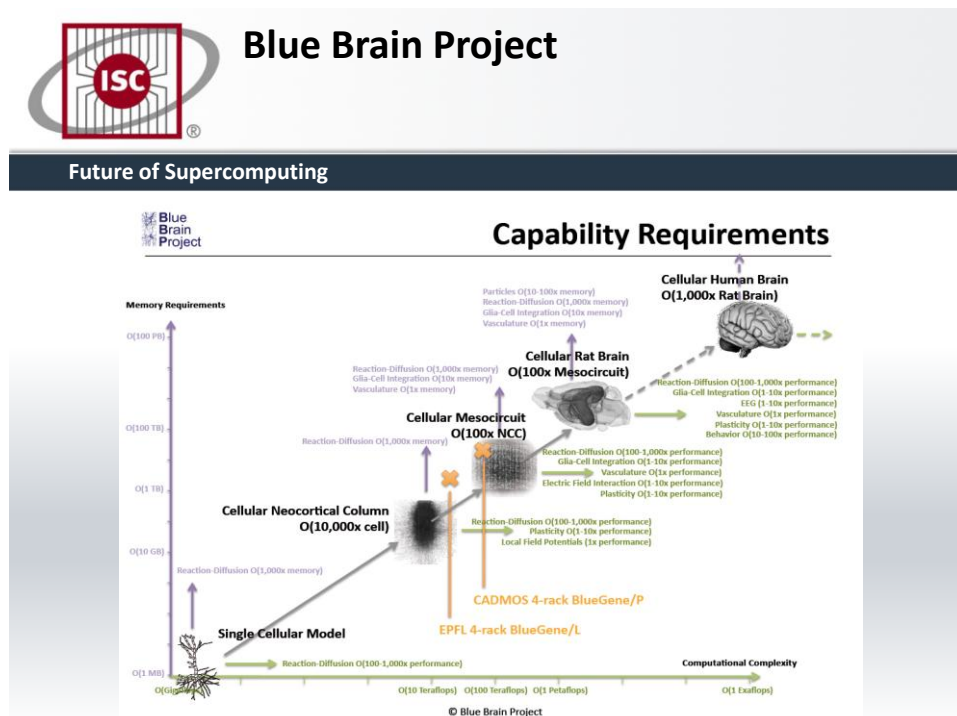


Fig. 13

## 4 Alternative Performance Metrics

Despite some criticism of its suitability, the LINPACK benchmark is the most common measurement for supercomputers. It has been the yardstick of the successful TOP500 lists for 20 years now, but it surely needs to adapt to the new Petascale situation as the execution time of extremely large LINPACK runs exceeds 20 hours. HPC centers will have more and more problems justifying why they run such long LINPACK benchmarks.

The pros of LINPACK are:

- only one figure of merit
- simple to define and rank
- allows problem size to change with machine and over time
- opens competition

The cons are:

- emphasizes only 'peak' CPU speed and number of CPUs
- does not stress local memory bandwidth
- does not stress the memory system
- does not stress the internal network
- a single figure of merit cannot reflect the overall performance of an HPC system

It is obvious that there is a need for additional benchmarks to measure the performance of supercomputers, e.g., the HPC Challenge benchmark, the Graph 500 and the Green500, to name just a few.

## 4.1 HPC Challenge Benchmark

At ISC'06, Jack Dongarra – the founder of this benchmark suite – gave a keynote presentation on "HPC Challenge Benchmark and the TOP500" [3].

The HPC Challenge benchmark basically consists of seven tests:

- HPL – the LINPACK TPP benchmark which measures the floating point rate of execution for solving a linear system of equations;
- DGEMM – measures the floating point rate of execution of double precision real matrix-matrix multiplication;
- STREAM – a simple synthetic benchmark program that measures sustainable memory bandwidth (in GB/s) and the corresponding computation rate for simple vector kernel;
- PTRANS (parallel matrix transpose) – exercises the communications where pairs of processors communicate with each other simultaneously. It is a useful test of the total communications capacity of the network;
- RandomAccess – measures the rate of integer random updates of memory (GUPS);
- FFT – measures the floating point rate of execution of double precision complex one-dimensional Discrete Fourier Transform (DFT);
- Communication bandwidth and latency – a set of tests to measure latency and bandwidth of a number of simultaneous communication patterns, based on b\_eff (effective bandwidth benchmark).

Each of the seven tests addresses a different part of a computer system.

We do not have the advantage of a single figure of merit any longer, and the results of the HPC Challenge benchmark are much more complex, which makes comparing different systems more difficult. One way to compare different systems is the Kiviat chart. A Kiviat chart visually displays a set of metrics that provides easy viewing of multiple metrics against minimum and maximum thresholds.

With these charts, it is much more difficult to report on new systems entering the HPC arena than with LINPACK only. Therefore, we will see LINPACK continue as the TOP500 benchmark for some more time. The HPC Challenge benchmark has already become a certain standard when it comes to selecting an HPC system but it needs benchmarking experts to interpret the Kiviat charts.

## 4.2 The GREEN500 List

The GREEN500 list, overseen by Wu-chun Feng and Kirk W. Cameron of Virginia Tech, is another approach to ranking supercomputers ([www.green500.org](http://www.green500.org)). The purpose of the list is to provide a ranking of the most energy-efficient supercomputers in the world.

The inaugural GREEN500 list was announced at SC08 in November 2008. As a complement to the TOP500, the unveiling of the GREEN500 list ushered in a new era where supercomputers can be compared by performance-per-watt.

At SC11 in November 2011, the latest issue of the GREEN500 list was published with 500 entries. The number one system in the TOP500 list of November 2011 – Fujitsu's K Computer – reached a remarkable 32<sup>nd</sup> place, even though its power consumption of 12.659 MW is the highest ever seen in a TOP500 list.

### 4.3 The GRAPH500 List

The GRAPH500 project by Richard Murphy, Sandia National Laboratories, and others is a highly important project addressing the increasingly dominant data-intensive supercomputer applications. Current benchmarks do not provide useful information on the suitability of supercomputing systems for data-intensive applications. A new set of benchmarks is needed in order to guide the design of hardware/software systems intended to support such big data applications.

While the TOP500 addresses number crunching, the GRAPH500 performs data crunching. Graph algorithms are a core part of many analytic workloads. Backed by a steering committee of 50 international experts, GRAPH500 will establish a set of large-scale benchmarks for these applications. The status of the GRAPH500 project:

Three major application kernels

- Concurrent search;
- Optimization (single source shortest path);
- Edge-oriented (maximal independent set);

It addresses five graph-related business areas: cyber security, medical informatics, data enrichment, social networks, and symbolic networks.

The GRAPH500 was announced at ISC'10, and the first list was published at SC10, with nine systems ranked. Further results were published at ISC'11 (29 systems), SC11 (49 systems) and ISC'12 (80 systems) ([www.graph500.org](http://www.graph500.org)).

## 5 The Future of Supercomputing

### Projected Performance Development

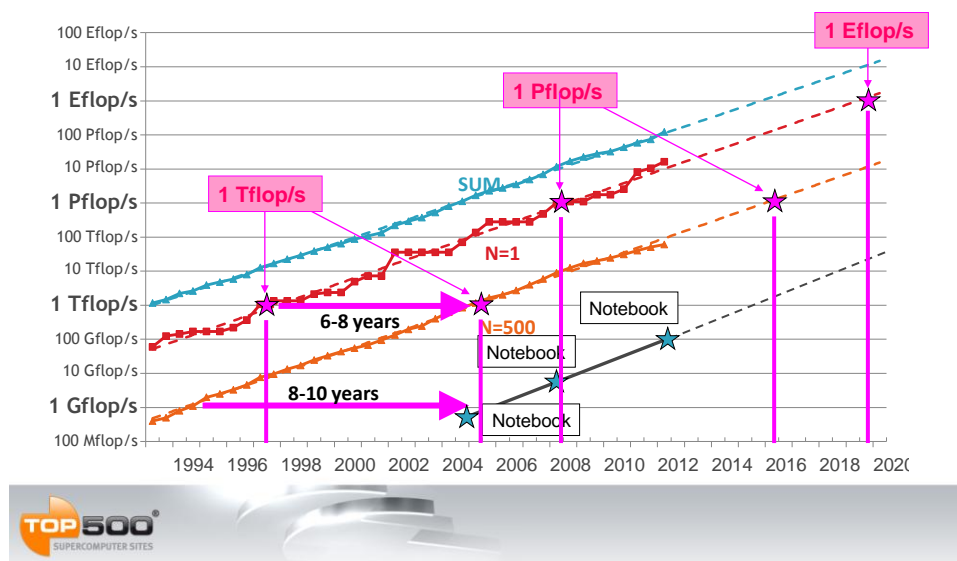


Fig. 14

We have done a projection into the future, based on 39 lists of real data, by a least square fit on the logarithmic scale. A powerful notebook computer, for example, will have a Teraflop/s performance in

the year 2015, 18 years after the first Teraflop/s system, the Intel ASCI Red supercomputer, entered the HPC arena.

Generally, it will take six to eight years for any system to move from No. 1 to No. 500 and eight to ten years to move from No. 500 to notebook level.

The system performance seems to increase by a factor 1,000 every eleven years: the Cray2 broke the Gigaflop/s barrier in 1986 (in pre-TOP500 times); in 1997, Intel ASCI Red broke the Teraflop/s barrier; in 2008, IBM Roadrunner set the Petaflop/s threshold. The latest TOP500 list (June 2012) had 20 Petaflop/s systems; in 2016, the list will show Petaflop/s systems only.

The LINPACK Exaflop/s threshold is expected to be reached by 2019. If we compare the growth in this figure with Moore's Law, we found that our growth is larger. We have a doubling for the No. 1 position in 13.2 months and a doubling for the No. 500 position in 13 months.

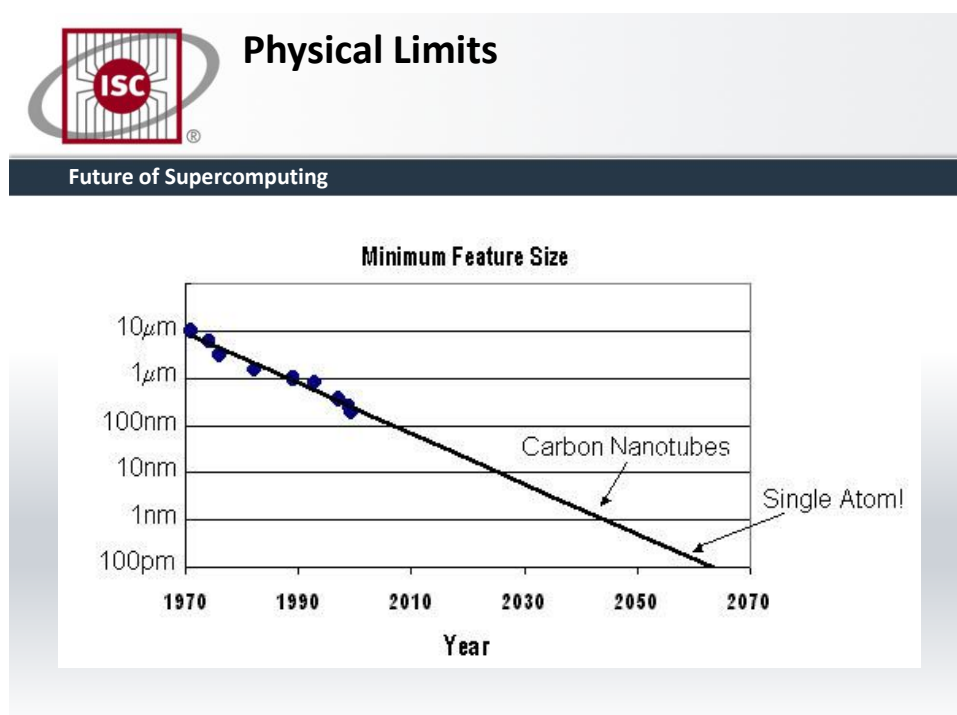


Fig. 15

The question is 'How long will we be able to continue?' The problem is that to increase the number of transistors per chip, the transistors must be smaller and smaller. This means that the manufacturing process must be able to define ever-smaller feature sizes. Obviously, there are physical limits to this process but nobody knows when they will be reached (Fig 15. from [4]). No doubt we will reach the Exascale threshold within this decade, but it needs to be seen if we will make the Zettascale threshold, according to Moore's Law, around 2030 and with silicon technology.

We would like to finish with a remark on quantum computing, one of the future technologies if silicon technology and Moore's Law will end.

'In the past, people have said, maybe it's 50 years away, it's a dream, maybe it'll happen sometime,' said Mark B. Ketchen, manager of the physics of information group at IBM's Thomas J. Watson Research Center in Yorktown Heights, N.Y. 'I used to think it was 50. Now I'm thinking like it's 15 or a little more. It's within reach. It's within our lifetime. It's going to happen.'" (IBM, March 2012)

## 6 Acknowledgement

Hans Werner Meuer cordially thanks Lord Laird of Artigarvan for the honorable invitation to give the Lorraine King Memorial Lecture at the House of Lords, London, on April 18, 2012, and Mr. Kevin Cahill and ComputerWeekly for sponsoring this event.

Both authors would like to thank the following gentlemen (in alphabetical order) for providing their slides for chapter 3, 'Supercomputer Applications in Science and Industry': Werner Butscher, John Johnson, Norbert Kroll, Volker Springel, Felix Schürmann, and Günther Zängl.

We cordially thank Mr. Andreas Stiller, Senior Editor at Heise Zeitschriftenverlag, for providing the Rmax (best LINPACK performance values) of notebooks over the past ten years as used in **Fig. 5** and **Fig. 6**.

Our thanks go to Mr. Jon Bashor, LBNL, for checking the manuscript, and to Mr. Peter Grambitter, Missing Link, for helping to prepare the figures in this paper. Last but not least we would like to thank Ms. Heike Walther, Prometheus GmbH, for her overall support and assistance.

## 7 References

- [1] wiseGEEK: clear answers for common questions/ [www.wisegeek.com](http://www.wisegeek.com)
- [2] Meuer, Hans Werner (2008), Informatik\_Spektrum\_31\_3\_2008, pp 203 - 222, Springer - Verlag 2008
- [3] <http://www.netlib.org/utk/people/JackDongarra//SLIDES/isc-talk-2006.pdf>
- [4] <http://barrett-group.mcgill.ca/tutorials/nanotechnology/nano03.htm>
- [5] By courtesy of Volker Springel
- [6] <http://bluebrain.epfl.ch/>